

Science, Technology, and Knowledge:

What Historians can learn from an evolutionary approach.

Joel Mokyr, Departments of Economics and History

Northwestern University

Department of Economics

J-Mokyr@nwu.edu

Presented to the Conference on The Evolution of Science, Santa Fe, May 16, 1998.

This paper is part of my forthcoming *Neither Chance nor Necessity: Evolutionary Models and Economic History* (Princeton University Press, 1998).

Final version, June 1998. This paper has been much improved thanks to the insightful suggestions by Paul A. David, Richard R. Nelson, and Ulrich Witt.

Introduction

How are we to think of the *fundamental* causes of technological change? I should like to argue here -- but not demonstrate -- that standard neoclassical economic models perform poorly in explaining technological progress. In a recent paper, Edward Prescott (1997) has noted the failures of standard theory to explain the huge differences in incomes and productivity. Instead of criticizing standard analysis again, I propose to experiment with an alternative. The argument is that economic growth depends crucially on the growth in human knowledge, both in terms of *what* is known and *who* does the knowing. What follows is an attempt to create a framework for an evolutionary approach to the economic history of technological change. It should be regarded as a

In the paper below, I propose to sketch out the rough outlines of a model of knowledge in economics based not on standard neoclassical analysis but on evolutionary dynamics. The idea that knowledge can be analyzed using an evolutionary epistemology based on blind variation and selective retention was proposed first by Donald Campbell and has since been restated by a number of scholars in a wide variety of disciplines.¹

In previous work, I have outlined the potential of the use of evolutionary biology in the economic history of technological change.² A reasonable criticism of such arguments has been that whereas models of blind-variation with selective retention may be an instructive way to look at innovations, they add little direct insight that cannot be gained from standard models. In other papers I have provided examples of insights that come more naturally from evolutionary models, although most such results can also be teased from standard models, properly specified. For instance, economics' knee-jerk response is to regard technological diversity as a source of inefficiency: if a product under very similar circumstances is made in different ways, our first suspicion as economists is that at least one of the producers is doing something wrong. Alternatively, if the system is in equilibrium, economics suggests that "very similar" circumstances still differ in some overlooked detail that explains the difference. An evolutionary perspective tends to regard variability as a source of innovation and long-run successful performance. Moreover, by some standards we could define the biological processes as "wasteful" without much evidence that mechanisms that make the use of scarce resources more efficient will eventually eliminate (Wesson, 1991, p. 94).

Part of the confusion in the use of Darwinian models in the Economics of knowledge and technology comes from imprecision regarding the unit of analysis. The unit I am interested here is not a living being, but an epistemological one, the *technique*. The technique is in its bare essentials nothing but a set of instructions, if-then statements (often nested) that describe how to manipulate nature for our benefit, that is to say, for production widely defined. Focussing on a *technique* underlines the main interest of students of the history of technology. Much of Economics (and of Darwinian Biology) is written as a story of the struggle between competing members of one species. But the interesting action in technology is not so much with people struggling against each other but struggling to control their natural environment. Techniques are thus tools to play a game against nature and to produce a "material culture" from an unyielding and often niggardly environment.

¹The original statement was made in Campbell, 1987. Among the most powerful elaborations are Hull, 1988 and Richards, For a cogent statement defending the use of this framework in the analysis of technology see especially Vincenti, 1990.

² See Mokyr, 1991, 1996, and 1998.

Below, I will first lay out the groundwork for an evolutionary analysis of technological knowledge, and briefly examine how it changes (or does not) over time.

A few definitions.

To start off, we need a definition of what essential elements constitute a Darwinian model. It will surely come as no surprise that there is little consensus on the matter amongst biologists or evolutionary theorists on the matter. Darwinian models encompass a larger set than just the evolution of living beings and population dynamics whence it first originated. Darwin himself recognized the applicability of random variation with selective retention to changes in language. Douglass North (1990) has suggested a similar approach to the development of economic institutions, Richard Dawkins (1976) to the realm of ideas (“memes”), Cavalli-Sforza and Feldman (1981), and Boyd and Richerson (1985) to culture, Donald Elliot (1985) to the analysis of Law, and Daniel Dennett (1995) to practically everything. The biological reproduction of living things in this scheme of things turns out to be a rather special case of a broad set of such dynamic models. The main idea of a Darwinian model is a system of self-reproducing units that changes over time. A Darwinian model must contain in my view of three fundamental elements (Mokyr, 1998).

One is the notion of the relation between an *underlying structure* that constrains but does not wholly determine a *manifested entity*. In biology, the underlying structure is the genotype, the manifested identity is the phenotype. The relation between the two is well understood, although there is still an endless dispute of their respective contribution to observed traits. In the history of technology, I submit, the underlying structure is the set of *useful knowledge* that exists in a society. The idea that such a set can be defined dates back to Simon Kuznets (1965). It contains all “knowledge” and beliefs about the natural world that might potentially be manipulated. Such knowledge includes the cataloguing and identification of natural phenomena, including regularities and relationships between them. This set, which I will call **S**, contains but is not confined to consensus scientific knowledge. It also contains beliefs, traditions, superstitions, and other knowledge systems which may not get down to the principles of *why* something works but all the same codify it. It certainly contains discredited and erroneous theories or theories that will subsequently be refuted. A good example might be the humoral theory of disease or the Ptolemaic description of the Universe. As long as such beliefs are held by *somebody* they must be included in **S** which is the union of all such beliefs. In the next section, I will discuss some of the properties of this set.

The critical point is that the elements in this set maps into a second set, the manifested entity, which I will call the feasible techniques set **8**. This set defines what society can do, but not what it *will* do. It differs from **S** in that it defines our power over and not just our knowledge of nature. In some way, the dichotomy between **S** and **8** is symmetric to Ryle's (1949, see also Loasby, 1996) distinction between knowledge “that” and knowledge “how.” Some of the knowledge “how” can be readily written down and codified, such as in engineering textbooks and cookbooks, but some of it is clearly tacit, instinctive, even subconscious. The distinction involves a great deal of subtlety: as Rule points out, it is quite possible for a single individual to perform tasks without consciously recalling or even understanding the underlying natural principles and phenomena involved, and conversely for brilliant theorists to be incompetent on the shopfloor. For the process of economic production at the social level, however, matters are different. Some knowledge has to exist

somewhere for a technique to be effective. Thus, the rider need not understand the mechanical principles of bicycle riding to ride a bicycle, but if this knowledge -- possibly in a rudimentary and incomplete form -- exists nowhere at all, bicycles would not be built.³ New techniques may well be generated by people who do not actually practice it.⁴ The nature of the process of economic production is such that the knowledge can exist in one place or in one carrier and the technique it maps exists elsewhere. To be sure, as I will argue below, elements in **S**, that is, techniques, could exist with little or no base in **S**, yet all such techniques are enormously limited and clearly would not form the basis of sustained technological progress.

Each technique is a set of instructions how to manipulate nature to attain a desirable outcome.⁵ The outcome is then evaluated by a set of selection criteria that determine whether this particular technique will be actually used again or not, in similar way to the fashion in which selection criteria pick living specimens and decide which will be selected for survival and reproduction and thus continue to exist in the gene pool. Thus the humoral theory mapped into techniques such as bloodletting and the Ptolemaic theory implied certain rules about navigation and the determination of latitude.

The analogy between the mapping of genotype to phenotype and the mapping of underlying knowledge to technical practice in use is inexact and to some extent forced: while genomes will vanish as soon as the species is extinct, knowledge can continue to exist even if the techniques it implies are no longer chosen. In other words, in all living beings the composition of the gene pool depends on the selection on the phenotypes because the genes cannot exist outside the living beings which are the “vehicles” or “carriers” of the genetic information. In this respect, knowledge is quite different. The properties of **S** and the distribution of knowledge do not depend principally on which techniques are selected because knowledge about nature can be carried by storage devices. Indeed, a technique can be “extinct” (not used by anyone) and yet still in **S** (e.g. explained in old engineering textbooks). This conclusion needs to be modified by the existence of *tacit* knowledge which is

³As Ryle points out (1949, p. 49), a man could not be a good surgeon without knowing anything about medical science. The intelligence involved in putting prescriptions into practice is not identical with that involved in grasping the prescriptions.

⁴Adam Smith (1776 [1776], p. 14) notes that “Many improvements have been made ... by those who are called philosophers or men of speculation, whose trade is not to do any thing but to observe every thing.”

⁵It might be noted that this definition does not define any role for artifacts, often thought to be the main unit of selection in evolutionary models of technology (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 “weaving using a flying shuttle.” The flying shuttle could be thought of as “the artifact” yet it was no more the artifact than the loom itself, the “pickers” 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. Many artifacts are meaningless without specific instructions. A shovel, say, can be used to increase agricultural output, to protect one from a flood or to dispose of a corpse. Moreover, as an artifact a shovel confuses two techniques: how to make shovels and how to use them for specific purposes. Concrete artifacts can be seen to be awkward units of selection once we realize that techniques sometimes involve instructions of “how to” that may be most useful and yet involve few or no artifacts such as avoiding exposure to bacteria by washing one's hands before eating, training animals, or how to navigate using the stars.

defined here as knowledge that can only be carried by the technique itself. If all knowledge were tacit, the distribution of knowledge would depend on the selection on 8.

Second, any Darwinian model must be a dynamic system of change over time, a stochastic process of some definable characteristics. As Nelson (1995) stresses, these models are in a class that is more or less half-way between deterministic and purely random dynamic systems, what he calls “somewhat random variation.” At each moment, the state of the world is given and around it there are stochastic innovations that could be wholly random but do not have to be. These innovations create a fortuitous variation that defines the options for the system to move to. In this kind of a model techniques “reproduce” from period to period and thus “carry” the knowledge embodied in them over time. A technique, in this view, uses human agents to reproduce itself to make another technique much like, as in Samuel Butler’s famous quip, a chicken is the way an egg produces another egg. It seems plausible, for instance, to think of it as a Markov Chain in which the state in t is dependent entirely on the state in $t-1$ and earlier history does not matter since it is entirely encapsulated in the state at t . “Extinction” might then be thought of as an absorbing barrier, but as long as the underlying knowledge (or some crucial component of it) has not been lost, the technique can be regenerated.

How do techniques reproduce themselves? The most obvious mechanism is through repetition: if a truck driver follows the instructions how to get from Cincinnati to Kansas City, and then does so again, the technique has reproduced. If the agent changes, as in the long run has to be the case, learning and imitation take place. But as long as we insist that the technique itself is the unit of selection, the identity or characteristics of the agent is not the main subject of discussion. There are many ways to drive from Cincinnati to Kansas City and among those certain specific routes are selected and others are not. Because the number of uses of each specific technique changes over time (as a function of certain characteristics), evolutionary processes belong to a special group of Markov chains known as “branching processes.” In these models each unit reproduces a number k of offspring, where k is a random variable.

Again, the evolutionary dynamics differ in some important way between living beings and techniques. In living beings, persistent change occurs only through gene mutation (essentially random change) and direction occurs through selection on the living beings which carry them. In technical knowledge systems there are two stochastic process at work: the useful knowledge reproduces itself over time with possible “mutations” (discoveries about natural phenomena). The techniques also reproduce themselves and there, too, there can be change, say, through experience and learning by doing. The two stochastic processes are clearly related, with feedback going in both directions. Such feedbacks do *not* occur in living beings, where Lamarckian feedback mechanisms from phenotype to genotype are ruled out. This two-pronged stochastic process is depicted in fig. 1. This feedback should be thought of as occurring through the impact of new techniques on knowledge, e.g. the invention of telescopes affected knowledge of astronomy or the impact of early steam engines on the development of theoretical physics.⁶

⁶The impact of technology on natural knowledge is stressed by Nathan Rosenberg (1982), though Rosenberg confines his essay to “science.” Yet many elements in S are made possible through better techniques that we would not think of as “science” including for example the European discoveries of the fifteenth century, made possible by better ship-building and navigational techniques.

Third, there is a property of superfecundity in the system, that is, there are more entities than can be accommodated, so there must be some selection in the system. This selection process is what provides the entire system with its historical direction by determining the likelihood that a certain technique will be actually used. The nature of superfecundity in epistemological systems is different than in Darwinian biology, where entities reproduce at a rate that is faster than can be accommodated by available resources. In the world of technology it essentially means that there are far more conceivable ways to skin a cat than there are cats and more ways to drive from i to j than can be accommodated. Selection at the level of technique in use is thus essential. As Nelson has pointed out (1995, p. 55), the theory's power depends on its ability to specify what precisely these selection criteria are. The answer must be that they are historically contingent: For instance, some societies prefer such as nineteenth century America emphasize price and efficiency above all, others may select against mass production and prefer individually manufactured custom made products wherever possible. Unlike Darwinian models, however, selection is not a metaphor for an invisible hand kind of mechanism that operates in a decentralized and unconscious manner: there are actually conscious units, firms and households, that do the selecting.⁷ Rather than the unit of selection, then, in this kind of model economic agents are the *selectors* themselves.

The importance of selection is critical to Darwinian models because, as noted, selection lends direction to the historical purpose of change. Selection and innovation are indispensable parts of the historical process. Without innovation, selection occurs on existing entities only, and whatever variation exists at any time will eventually disappear. Without selection, innovation either disappears entirely (if a new technique never has a chance of being chosen, that is, implemented) or becomes completely chaotic (if any and all new technique proposed is adopted).


Darwinian models of technology need therefore to specify the exact mode of selection that is operating on the choice set \mathcal{S} . But because innovation consists of more or less local variation on *existing* states of the world past selections govern the range of variability on which selection can occur, yielding the path-dependent property of these models. By the standard definition of path dependence, this means that the final outcome depends both on the special characteristics of each technique and the historical path which partially is contingent (David, 1997a). A multiplicity of conceivable outcomes with the actual result often determined by historical contingency is thus part and parcel of the process. Yet it is not quite the same as the standard problems that occur in economics with multiple equilibria and the need to refine them. As Witt (1997a) points out, the process of evolutionary change is *unending*, that is, unforeseeable mutations always can and do occur to destabilize an existing state of the world. Such mutations occur either in a given technique itself or in other techniques that are complementary or rivalrous, thus changing the environment faced by this mutation.

What these models do not give us is why new ideas occur at all. If \mathcal{S} expands and hence makes new techniques possible, this just pushes the question into another corner. Economists have long believed that innovation is produced in the system, by agents devoting resources to the search


⁷It should be noted that the combination of selection and the particular dynamic structure defined before imply that selection is “myopic” even when it is perfectly rational conditional on what is known at the time. That is, a particular choice may seem rational but that choice places the system on a trajectory that eventually leads to less desirable outcomes. For more details, see Mokyr, 1992.

for innovation. To the extent that this occurs, human agents not only select from the available techniques but also help in bringing about the diversity in the menu from which the choosing takes place. Yet such black boxes have been roundly criticized not just for not specifying why and how some of those searches end up being successful and others fail, but more deeply that they never explain why some people think *ex ante* that such a search will be more or less fruitful so that they can decide to which search they should allocate these resources and how many. Before the twentieth century, in any event, the process of R&D, such as it was, was far too haphazard and success far serendipitous to merit the term “production.”

There are many issues in technological history that can be re-explored using evolutionary ideas and terminology. For instance, the question of whether nature makes leaps: does technological change occur in a gradual manner as Leibniz, Alfred Marshall and the neo-Darwinian phylogenetic gradualist orthodoxy in evolutionary biology hold, or can it move in bounds and leaps as Eldredge and Gould insist? The debate parallels those in economic history between scholars who believe in the Industrial Revolution and the great discontinuity it constituted and those who would deny this. Are Darwinian models of natural selection sufficient to explain the course of history as the ultra-Darwinians such as Dennett and Dawkins claim, or do we need additional inputs from chaos theory, self-organization theory, or something yet unsuspected? Clearly these debates mirror those between the scholars who insist on multiple equilibria and hence path-dependence such as Brian Arthur (1994) and Paul David (1992, 1997), and their opponents such as Liebowitz and Margolis (1995) who feel that the rational market -- if only left alone -- will get it right every time and hence there is no need to worry about path dependence. What exactly is the relation between the environment



Standard Darwinian Model



Evolutionary Epistemological Model of Technology

Figure 1

and technological selection, and can we distinguish meaningfully between adaptation and innovation? Are traits invariably explicable in terms of their functions, as the fundamentalist adaptationists claim? To what extent can evolutionary theory help us understand important issues in invention: the notion of exaptation, for instance, which refers to cases in which an entity was selected for one trait but eventually ended up

carrying out a related but different function.⁸ What kind of environment facilitates innovation? What is the nature of resistance to innovation and how is it overcome if it is? It is also crucial to re-explore the connection between the history of natural knowledge which provided part of the “underlying structure” (in biology: the genotype), and the “manifested entity,” the technique (in biology: the phenotype). Some concepts of biology do come in handy in describing this mapping function such as the concept of epistasis (more than one piece of knowledge is required to generate a technique) or pleiotropy (one piece of knowledge generates multiple techniques). Concepts like “niches,” “recombination,” and “extinction” come naturally to the historian of technology, although they may not mean the same thing they do in biology. It is also possible to use an evolutionary framework to distinguish between autonomous and “induced” technological change (Mokyr, 1998c). There are interesting questions around the “Panglossian” point that “whatever was, was good.” Is it possible to rank outcomes using some kind of criterion, and assess to what extent historical outcomes were “desirable”? Finally, we may ask in what sense can we think of “progress?” Biologists have long argued the point, and the consensus today is that there is no obvious way in which progress can be defined in the evolution of living beings. Is the same true in the history of knowledge?

Knowledge and Technique

Of all living creatures, *homo sapiens* alone seems to have a conscious knowledge of the natural world. Such knowledge consists first and foremost of the observation and cataloguing of natural phenomena and secondly of discerning natural regularities and causal mechanisms that link these phenomena. A particular subset of this knowledge is what we call “modern science.” The megaset S defined above contains the following information: consider all knowledge about the natural world contained either in people's minds or in storage devices such as magazines, books, hard disks and so on. The union of these pieces of knowledge is the total set of S . A discovery, scientific or not, is the addition of a piece of knowledge to S that was not there before.⁹

This definition immediately raises many questions. For one thing, it is unclear where exactly to draw the border between knowledge about the natural world and knowledge about other matters. There are grey areas, particularly in the social and cognitive sciences. Clearly, demarcation lines inevitably are arbitrary and a matter of taste: if managers use insights from psychology or sociology to organize their workers in a specific way, would this be a mapping of the kind described? Secondly, the “existence” of knowledge in the sense that somebody knows it still does not mean that it can be accessed with ease by those who do not know it. It also forces us to define “society” with some care. Some natural phenomena may be known in one society and unknown in another. If these societies are totally disjoint (such as pre-Columbian America and medieval Europe) this is not much of a problem. In other cases, however, this becomes more ambiguous. When is knowledge truly common knowledge in the sense that anyone can costlessly acquire it? Thirdly, useful knowledge may be in

⁸The term was first suggested by Gould and Vrba (1982). Historians of technology have of course long pointed to the phonograph and the digital computer as examples of such instances in economic history. In his recent book Dyson (1997, pp. 81-82) draws the parallel explicitly: “feathers must have had some other purpose before they were used to fly...the same thing happened to ENIAC: a mechanism developed for ballistics was expropriated for something else.”

⁹Geographical discoveries are a prime example of a discovery that is not scientific (in the standard sense of the word) that maps knowledge of the physical world directly into a techniques space.

dispute in the sense that two elements in *S* maybe mutually incompatible. When this occurs, both elements remain in *S* in the sense that they are “known” but obviously they need not be symmetrical.

Is there selection at the level of *S*? Because in knowledge systems storage of information does not have to occur exclusively in the manifest entities that carry them, it is unclear what precisely would be meant by superfecundity. Only if there is some form of congestion of knowledge caused by storage cost will society shed some pieces of knowledge as it acquires and selects better ones. Whereas this is surely true for an individual with a finite memory, it is less obvious for society's knowledge being the union of all knowledge stored up in memory banks, libraries, hard disks and so on. Yet through most of human history before the Age of the Gigabyte, such congestion was a reality: books were hugely expensive before the invention of printing, and while their price of course fell, they were still quite costly by the time of the Industrial Revolution. Other forms of storage outside human memory banks such as drawings, artefacts in musea, were all costly. Some selection may therefore have even occurred at that level. Moreover, in many techniques knowledge drawn from storage devices may not suffice to fully define the set of instructions. So-called “tacit knowledge,” a term coined by Michael Polanyi, plays a role in production. This part of the “how-to” instructions that form a technique must be learned from experience or from another person, and therefore the part of *S* underlying it is subject to congestion and thus selection. More sophisticated technology transmission may reduce the ratio of tacit to formal knowledge that can be transmitted through non-human devices. What previously required an instructor or actual on-the-job experience can now be viewed on videos or simulated on computers.

Yet selection and extinction in *S* are very different from natural selection. While of course certain views of nature are incompatible with each other so that some theories are rejected if others are accepted, they do not necessarily become *extinct* in the technical sense of being inaccessible. Thus the humoral theory of disease is still understood today, but no longer serves as a source for prescriptive techniques in modern medicine. Scientific theories that are “accepted” will normally be the ones that are mapped onto the techniques in use, whereas the ones that are rejected will be dormant, known only to historians of knowledge or stored in library books. Accepting the work of Lavoisier meant that one had to abandon phlogiston theory but not destroy any trace of it. Copernican and Ptolemaic views of the Universe reside together in *S*, though of course not in the same position.

Insofar that there are such incompatibilities between different subsets of *S*, people have to choose between them and will do so if they can, in some sense, rank them (Durlauf, 1997). We need to distinguish between knowledge that is “widely accepted” (consensus) and knowledge widely believed to be false. Absolutes are of course useless here, since people believing in Creationist science and members of flat earth societies must be allowed for even if we can define a consensus from which they beg to differ. Many beliefs about nature remain “in dispute” at each point of time, and clearly those who adhere to a particular subset of *S* will be ones that map this knowledge into techniques. Those who believe that mosquitoes can transmit the HIV virus, will take any precaution possible to avoid being bitten.

Yet of course each societies has rules by which it tests knowledge about nature. The *tightness* of a piece of knowledge, defined as the ease with which others can be persuaded to accept it, depends on the consensus-forming mechanisms and rhetorical tools that are admissible in distinguishing between it and its alternatives. Mathematical logic, experimental and statistical evidence, consistence with pre-existing authority (“Aristotle says”) and many other criteria determine whether a piece of

knowledge will be accepted. These standards are invariably socially set within paradigms: what constitutes logical “proof”? what is an acceptable power of a statistical test? do we always insist on double blindness when testing a new compound? how many times need an experiment be replicated before the consensus deems it sound? If a piece of knowledge is not very tight, as Durlauf shows, choosing between competing pieces of knowledge may well become a function of imitation, persuasion, and fashion.¹⁰ When knowledge is “rejected” for whatever reason, it can become dormant in the sense that it is no longer mapped into \mathcal{S} . How and why such selection occurs is of course the central issue in the History of Science with emphasis varying between scholars to what extent evolutionary success is correlated with some measure of “progress” or “truth” (Kitcher, 1993).

The set \mathcal{S} can also be divided into “active” and “dormant” knowledge, which are mapped or not onto \mathcal{S} , a bit like “coding” and “non-coding” DNA. Again, a grey area makes such distinctions somewhat tricky. Advances in paleontology or improved understanding about the distances of other galaxies or the properties of black holes are at first glance dormant forever, but many dormant sections of \mathcal{S} can become active given a change in the environment.¹¹ This is, I submit, what we must mean when we speak of induced technological change (Mokyr, 1998c). We can thus subdivide \mathcal{S} into four different cells in a little two by two table: whether knowledge is active or not, and whether it is accepted or not. The cell “not accepted” but “active” is far from empty, not only because some people may not share the consensus but because the very essence of prayer, magic, and miracles is to beg exceptions of nature's regularities rather than exploit them.

The exact relationship between \mathcal{S} and \mathcal{S} varied a lot across different areas of knowledge as well as over time. In general, techniques can have narrow or wide bases in \mathcal{S} . In some cases, new techniques are discovered serendipitously without anyone having a clue as to *why* this works. Such techniques may be termed *singleton* techniques: the only element in \mathcal{S} that corresponds with it is the knowledge that such and such works. The mapping is then of course trivial. Historically, singleton techniques were very common. Farmers sowed clover on arable fields because they knew it raised subsequent fertility. They did not know that clover increased soil nitrogen content, much less how (through a symbiotic nitrogen-fixing bacterium). Before 1800, knowledge of biological processes and organic chemistry hardly existed, and most techniques in use were singletons. Such is not the case, say, for water power, where a body of formal knowledge of hydraulic science was worked out in the eighteenth century. For a more detailed discussion of singleton techniques see Mokyr (1998c).

¹⁰One example to illustrate this principle: A Scottish physician by the name of John Brown (1735-88) revolutionized the medicine of his age with Brownianism, a system which postulated that all diseases were the result of over- or underexcitement of the neuromuscular system by the environment. Brown was no enthusiast for bleeding, and instead treated all his patients with mixtures of opium, alcohol, and highly seasoned foods. His popularity was international: Benjamin Rush brought his system to America, and in 1802 his controversial views elicited a riot among medical students in Göttingen, requiring troops to quell it. A medical revolutionary in an age of radical changes, his influence is a good example of the difficulty which contemporaries had selecting amongst alternative techniques and the enormous possibilities for failure in this area (Brown was asserted to have killed more people than the French Revolution and the Napoleonic Wars combined).

¹¹By “environment” here I mean not only the physical environment in which the technique operates but also the development of complementary or rival techniques which may lead to the activation of previously dormant knowledge. Indeed, such processes are what constitutes “adaptation” in all evolutionary processes.

Over time, the interaction between S and δ is central to the development of knowledge yet there is nothing deterministic or obvious about the relation. Some natural discoveries are mapped fairly immediately into some kind of practice. Others remain dormant for long periods, then are activated, often combined with other discoveries. Pure “inventions” (expansions of δ) are trivially contained in S by realizing that the knowledge that such and such works is part of S . Yet S confines and constrains technological practices far more than is realized: the insight that hot air expands and thus has lighter specific weight is necessary to build a hot air balloon, just as the knowledge that carbon content determines the quality of steel is necessary to make steel in a Bessemer converter (though not to make steel altogether). At the same time new knowledge in S is produced, but it is not produced randomly: it responds in some sense to perceived needs and explicit curiosities, and it enhances knowledge through tools and instruments fashioned by techniques. Moreover, technological practices operate as a focusing device for researchers. When a technique exists with a narrow base in S , it is natural to direct research into that direction: why does aspirin work? why is an overshot waterwheel more efficient than an undershot wheel?

Three features of this set-up are worth underlining. First, as long as there are *some* selecting agents that do not share the consensus that some element in S is false, mapping may continue and obsolete techniques could survive in niches. Particularly when traits are not very tight, as is the case in processes dealing with living beings such as agriculture and medicine, techniques will survive even when their base in S has been rejected by the preponderance of professional opinion and might be resisted even if scientific opinion supports them. Second, when an element in S has been rejected, it is *not* eliminated but survives in storage devices such as books on the History of Science. It happens -- if rarely -- that knowledge believed to be obsolete is resurrected.¹² Third, the bulk of S is dormant because there is no known way in which this knowledge can be used to map into anything useful (say, the speed at which the planet Uranus circles around the sun). Environmental changes can activate such dormant knowledge when, for example, complementary knowledge emerges that turns previously unexpressed knowledge into something that can be mapped into δ .

One possibility of dealing with this issue is to define an *access cost* function. This function determines how costly it is for an individual to access information from a storage device or from another individual. The access function defines the relevant cost as the cost to an individual in society to acquire this knowledge. The average access cost would be the average cost paid by all individuals to acquire the knowledge. More relevant for most useful questions is the marginal access cost, that is, the minimum cost for an individual who does not yet have this information. A moment reflection will make clear why this is so: it is very expensive for the *average* member of a society to have access to the Schrödinger wave equations, yet it is “accessible” for advanced students of quantum mechanics at low cost. If the rest of society “needs” to know something it will clearly go to someone for whom this cost is as low as possible to find out. Such people would then be defined as experts, and much of the way knowledge has been used has relied on such experts. The access cost is determined by three basic factors:

¹²It is more common for techniques that have fallen into disuse to be brought back. An example is the treatment of malaria which, because of the constant mutation of both the mosquitoes and the plasmodium parasite around all medications aimed at them, has become increasingly difficult to treat. It is reported that physicians in their desperation are returning to quinine and even to an extract of wormwood used in China many centuries ago. See Steve Jones, *The Language of Genes*, New York: Anchor Books, 1993, p. 223.

1. The degree of diffusion of knowledge. This can be easily defined as the ratio between the intersections of knowledge that contain this element and the total number of individuals in society. The number of people who know that the earth is not in the middle of the solar system or that pneumonia is caused by bacterium is probably not much less than the number of people in society; the number of people who know the laws of thermodynamics is much smaller relatively. However, the *total* size of S is also important: if the total amount of useful knowledge is large, the likelihood is that people will specialize increasingly, meaning that access requires a communications technology. Of course, diffusion is to a large extent a function of the nature of the knowledge itself and of the education system of the country, literacy rates and so on. Perhaps what matters most for technological success is the amount of contact between those who possess the knowledge in S or have access to it, and those dealing with the practical problems of production such as power generation, farming and medicine.
2. The technology of information retrieval. Because so much natural information is not “known” to people but can be looked up in storage devices, the technology of doing so is quite crucial. Information needs to be organized and classified into taxonomies, transformed and translated into different languages and jargons, and communicated effectively to those who want it. Innovations such as writing, alphabetization, the indexing of books, encyclopedias, library technology, “how-to manuals,” hard disks, and now the internet all have major effects on the access cost function and thus affect the efficiency by which the knowledge in S can be mapped into techniques in δ .¹³ It should be noted, however, that these techniques largely affect average and not marginal access costs. On the other hand, marginal access costs can be affected by allocating resources to maintain experts, specialized storage devices (say, well-organized technical journals and data bases).
3. As Dasgupta and (1988) and David (1997b) point out, modern science is “open” in the sense that it purports to produce a public good while minimizing access costs. Scientific findings are placed in the public realm once they are complete, and often before thus deliberately reducing access costs. Before the Scientific Revolution, science was more proprietary and some parts of (like alchemy and medical knowledge) actively kept secret. Clearly, the more open science is, the lower the access cost.

The access cost function is closely related to the mapping of S into δ and to the ability of any society to “search” over what it knows to actual instructions to do things. The lower the marginal access costs, the more society can search over the knowledge it has to generate useful techniques.

Turning to the δ space, as noted, the main evolutionary feature is that there are many techniques in δ and the selectors have to choose. I see three different types of selection here. One is the standard neoclassical mechanisms: techniques are selected according to whether they maximize some kind of objective function. This includes supply considerations (which techniques work at all and do

¹³Headrick (1998) has maintained that such access technologies underwent a major improvement in the eighteenth century with the growth of encyclopedias, technical handbooks and similar storage devices.

so more efficiently than others?), externalities (techniques may have strategic complementarities or incompatibilities with other techniques), and demand considerations (what does the market want?). But it will, if all works well, produce economically efficient solutions if they are the only ones to work. A second type of selection is what may be called hysteresis. In any Markov chain we can build in as much inertia and irreversibilities as we want. In biological evolution a camel cannot change into a zebra once it discovers that zebras are more suitable to a given environment. In technological evolution such changes can and do occur, but they do so at a cost, often quite high. A third type of selection occurs at a social level, much like the social constructivists tell us. At many levels, political power and lobbying, motivated by self interest, beliefs, prejudice, and fear enter upon the selection process and direct it in one direction or another. Such a model could of course be incorporated into the objective functions and reduced to “what does the market want” but it seems instructive at least to distinguish between criteria that relate to the actual functioning of a technique and other characteristics. For instance, non-Western societies might reject a technique developed in the West not because of any of its intrinsic characteristics but because it originated in the West. At a lower level, selectors may choose a technique through trust, conformism, tradition and so on. These are not necessarily just information-saving devices, but may reflect a variety of preferences, from resignation to the dictates of nature to religious beliefs that certain techniques are bad (e.g. Christian Science).

In selecting techniques, *tightness* can be defined when a technique can be easily and unambiguously compared to its alternatives so that selection is fairly easy. In many areas, however, it is in fact far from easy to decide what “works better” not only because the traits we are most interested in are imperfectly observed, but because there could be differences of opinion on how to weight different traits. We have no difficulty in picking a laser printer over a dot matrix, or a push button tone telephone over pulse dialing one, but we are far less sure whether to go to an acupuncturist for a lower back pain or whether irradiated vegetables are really safe. How was a seventeenth century farmer to decide whether to sow clover or turnips as part of his crop rotation? When neither the knowledge in **S** nor the techniques in **8** are very tight, techniques with different bases in **S** can coexist: Freudian psychiatry and anti-Freudian psychiatry are one example. When we cannot easily tell “what works better,” persuasion, intimidation, emulation, and inertia play important roles. Some examples of techniques which have historically been notoriously “untight” are birth control techniques (such as the rhythm technique and various potions and herbs that were supposed to make a woman infertile or abort); techniques to elicit truth from witnesses, including modern technological devices such as lie detectors; psychiatry and the use of psychiatric medicine such as electric shock treatments and psychotropic drugs; the feeding of babies. Needless to say, a great deal of knowledge at any time may have subsequently been refuted and yet at an earlier time was accepted and played a major role in mapping into techniques. Ptolemaic astronomy was used in the voyages of discovery, phlogiston theory was relied on by the metallurgists of the eighteenth century, and the caloric theory of heat in the development of early steam engines.

The tightness of **S** serves as a criterion for **8**: we no longer use bleeding in the case of infectious disease, not because double-blind experiments have shown it to be ineffective but because the entire theory on which it rested has been rejected. In turn, the tightness of techniques in **8** can at times serve as a criterion for competing segments of **S** when the standard criteria for knowledge are somehow inadequate. There is a classic tale of one of the great bacteriologists and sanitary scientist, Max von Pettenkofer, who objected to the germ theory of disease and tried to disprove it by publicly

drinking a glass of water contaminated with cholera germs and showing he would not be infected. While Pettenkofer survived this attempt, the rhetorical logic clearly defeated his purpose: the germ theory became widely accepted precisely because its recommendations worked. Behring's discovery of the diphtheria vaccine in 1890 clearly helped persuade many of the correctness of the germ theory and the discovery of the culpable germ.¹⁴ In water power, organic chemistry, electrical engineering, and the care for newborns, success as measured by the traits of a technique could be used to persuade the community in favor of one segment of S. Maybe, as Einstein said, Quantum Physics is wrong because God does not play dice with the universe, but the huge application of quantum physics in electronics persuaded the world that he might be mistaken.¹⁵ Strictly speaking, this is incorrect. Bad science sometimes maps into good techniques (e.g. phlogiston theory which guided the great pre-Lavoisier metallurgists such as Tobern Bergman and René Reaumur, and the miasma theory which implied many effective nineteenth century sanitary practices).¹⁶ Moreover, such practical criteria do not work in fields such as cosmology; the triumph of Copernican over Ptolemaic astronomy or even the acceptance of Darwinian evolution had little to do with the successful techniques eventually derived from them. But in the history of natural knowledge, including science, technical success was a good rhetorical device. Indeed, difficult and controversial as it has been to define “progress” in science (Kitcher, 1993), one could argue that we could measure the pudding by the eating: if the traits of a technique are quantifiable in terms of human welfare of one kind or another, the underlying knowledge can be said to be progressive if it leads to advances in techniques that improve human welfare. This is obviously an “impersonal non-epistemic goal” of science, using Kitcher's (p. 73) terminology: knowledge in the service of community which wants to know something about molecular processes in bacteria so as to avoid tooth decay and food poisoning. Its attainment is the greatest positive sum game in History. Yet it is not the only criterion for progress in knowledge, since it would exclude advances in knowledge of paleontology or cosmology that have no direct applications, and is perfectly consistent with the emergence and perpetuation of false beliefs.

Change and Inertia in evolutionary systems.

Below I develop a few simple tools and then use them to distinguish between technological adaptation, technological mutation, and technological speciation. These terms are of course borrowed from biology, but they will mean somewhat different things, and throughout this discussion we need to keep reminding ourselves that the fields are really intrinsically different in many dimensions.

First, consider the selection mechanism on techniques. Suppose that a technique has only two “traits” that define it. Call these T_1 and T_2 . The traits determine whether a technique will be selected. The standard economic approach identifies costs as the main trait that is of importance here. Yet

¹⁴As late as 1888 one of the leading American pediatricians, Lewis Smith, held strongly to the view that diphtheria was caused by children inhaling sewer gas (Rosen, 1993, p. 265).

¹⁵Other examples of this phenomenon, cited by Rosenberg, 1982, p. 157 are Lord Rayleigh's denial of heavier-than-air flight a few years before Kitty Hawk, and Rutherford's denial of releasing energy from the nucleus of an atom in the 1930s.

¹⁶Indeed, Pettenkofer and his colleague the great pathologist Rudolph Virchow, without any benefit of bacteriology which they rejected, made Berlin and Munich, respectively into healthy cities (Ackerknecht, 1982, p. 213).

clearly there must have been many others that explain why techniques have been selected. Some of those traits are readily observable: user friendliness, esthetic qualities, compatibility with other elements in a technological system. Others can only be observed over time from experience: durability, reliability, the existence of “hidden” costs not obvious at first (say, asbestos, cfc's), and hence mistakes are unavoidable. Many can only be observed from careful statistical analysis of large samples, larger than the user can generate so that knowledge of the T's is a public good. For instance, is one of the traits of broccoli that it reduces the probability of contracting colon cancer? In other cases techniques are inherently untight because the traits are difficult to observe for anyone, and so a great deal of rhetoric is expended on persuading users. The specific traits of one brand of cheap beer, or the advantages of heavy use of bactericidal soap are examples in point. In the past, with accurate testing being far more difficult, such choices were far more difficult than today. Pithy proverbs in the “an apple a day” tradition might have been enough to secure selection. “Good enough” may have been able to survive for extended periods with “best” simply because the comparison required data and statistical analysis beyond the capabilities of the selectors. The amazingly durability of medical practices such as bleeding and purging can also be explained this way.¹⁷ Techniques may have survived through sheer inertia or spread through imitation simply because there was no easy way of determining if they worked “better” or not.

Secondly, each time a technique is “used” it “lives” and a specimen has been “selected.” The environment, V , is defined as anything external to technique j that is not part of it, including but not confined to the T's of other techniques. For each technique, we may then define $\mu(j)$, which is a count of how many times technique j is used, a bit like the size of a population. For any set of traits and an environment, we can define an equilibrium level of μ^*

$$(1) \quad \mu(j) = S(T_1^j, T_2^j \dots T_n^j; V)$$

and $\dot{\mu} = f(\mu^* - \mu)$. For any V , there are combinations of the T's which define $\dot{\mu} = 0$. Assume for simplicity that $M\mu/MT_1 > 0$ (the trait is favorable) and that $M^2\mu/MT_1^2 < 0$ (diminishing returns) and the same holds for T_2 . We can then define the curve ZZ' in fig. 2 which defines the condition of fixed fitness ($\dot{\mu} = 0$).¹⁸ A deterioration in the environment (possibly due to changes in complementary or rival techniques, or a change in preferences) would be depicted as an outward shift of ZZ' . Each technique in use is defined as a point in this space. In addition to the techniques actually in use, given

¹⁷For a fascinating example of this in the context of contraceptive technology, see David and Sanderson (1997).

¹⁸This is akin to an equilibrium condition in standard models and even more than in those models, evolutionary models do not presume that the system is always (or ever) in equilibrium or even close to it. All it does is to determine the laws of motion of μ .




Figure 2

by the area $*$ in fig. 2, there is a larger set of all feasible techniques δ within which $*$ is wholly contained. The relationship between δ and total human knowledge is precisely the relationship between the underlying structure or “genotype” and the manifest entity or “phenotype” but of all the possible entities only a small subset get selected.

Technological selection thus occurs at two levels: not all techniques in δ are picked to be in $*$. In fact only a small minority of all feasible and known ways of making a pencil, shipping a package from Chicago to New York, or treating a patient suffering from Pneumonia are in actual use at any point of time. Secondly, techniques in use themselves are competing over the scarce resource of a finite number of usage, and in the long run, assuming competition is stringent, only the ones that are at E_0 actually maintain their numbers.

In many cases new and potentially useful knowledge emerged but failed to be expressed, that is, failed to persuade enough people to end up being translated into the accepted subset of S and thus to be mapped into δ . Thus the germ theory of disease was first proposed by Girolamo Fracastoro in the sixteenth century and proposed repeatedly without having influence on the practice of medicine until late in the nineteenth century. Conversely, much of the knowledge set in the past may be recognized today as “false” and yet has historically mapped into useful techniques. One can navigate a ship using the stars even using a Ptolemaic geocentric astronomy and design a steam engine based on the caloric theory of heat. Medical procedures based on Galenian medical theory could at times be effective. Selection, in short, occurs at both ends of the mapping function but whereas selection

in \mathcal{B} is at least in some cases performance-tested, selection in \mathcal{S} requires tools of persuasion, some rigorous, some purely rhetorical.

A glance at figure 2 demonstrates how evolutionary concepts can be invoked to understand technical change as a historical process. The set \mathcal{B} contains all feasible techniques in a given society, and illustrates how *adaptation* occurs as a consequence of changes in the environment: if the frontier ZZ' shifts around (say, to WW') and its tangency moves elsewhere, society can pick another technique somewhere in \mathcal{B} and thus adapt to the changed environment. The new technique may or may not be in \mathcal{C} . It can also go back and search over \mathcal{S} to see if there is a way of endogenously expanding \mathcal{B} . There is no guarantee than such searches lead to results; indeed, normally they do not.¹⁹ For evolution to work, it has to find the answer to changing needs within the *existing* raw material. The differences between a search over \mathcal{B} and one over \mathcal{S} is roughly equivalent to phenotypical and genotypical adaptation in living beings, always remembering that in living beings the “search” is purely metaphorical, the result of an invisible hand working through differential survival and reproduction, whereas in knowledge systems the “survival” and “differential reproduction” of bits and pieces of knowledge is often the result of willful and intentional actions of agents.²⁰ In knowledge systems, but *not* in living beings, it is conceivable for the “gene pool” \mathcal{S} to respond to some extent to changing needs. For example, pure research might be partially motivated by possible applications. Most major historical scientific breakthroughs were clearly not of that nature, but it is possible (especially in modern time) that some expansion in the knowledge base, for instance in molecular biology or materials science, is driven in part by a well defined need. The set-up here makes clarifies one distinction: new elements in \mathcal{B} can occur with no change in \mathcal{S} , using existing knowledge only, or it can take advantage of an expansion in \mathcal{S} . In a limiting case, \mathcal{B} can expand on its own (by purely fortuitous discovery) with the new opportunities duly recorded in \mathcal{S} even if they do not any knowledge of the natural regularities involved.

Localized learning is more common than technical changes in a completely novel region (David, 1975), so that we can envisage changes in \mathcal{B} to take normally the form of pimples and protuberances appearing on the frontier of \mathcal{B} . Yet localized learning runs into diminishing returns and dead-ends, and for sustained technological advance to occur, bold and radical departures need to take place. I have referred to such departures as “macro-inventions,” a term that describes less the impact of an invention on the economy as a whole as much as the relation of knowledge incorporated in the invention to the rest of the knowledge currently in existence and in use.

The idea of macroinventions is akin -- but not identical to -- the notion of speciation in biology. Speciation is the emergence of a new category of life that is distinct from everything that existed before. Such distinctions are often hard to make, because of the grey areas between the

¹⁹Steven Pinker (1997, p. 206) cites Lamarck as saying “new needs which establish a necessity for some part really bring about the existence of that part as a result of efforts” and adds, “if only it were so -- if wishes were horses, beggars would ride.”

²⁰ In some sense this approach redefines the traditional distinction between substitution and induced technological change. The evolutionary equivalent of this distinction is, roughly speaking, natural selection from a given set of heterogenous traits as opposed to the emergence of new phenotypes from a given gene pool, in which fortuitous new combinations -- if they emerge -- are picked up by selective forces. Whereas the former is a more or less deterministic and predictable process, the latter remains a matter of contingency.

categories. This is even true in biology: the distinction between species is purely based on reproductive isolation which is not a clean binary variable. In any event biological distinctions are rather more arbitrary in higher genera. Yet in most dynamic systems satisfying the criteria I outlined above we do recognize a different type when we see it and we have an intuitive notion of the distance even if this is hard to specify. Thus we realize that zebras are different from horses but that the difference between them is less than between a zebra and a cockroach. DNA analysis can nowadays quantify these metrics, but they were intuitively clear long before. Similarly, Catalan is different from Portuguese but closer to it than to Urdu. In the history of technology we can readily distinguish such categories although there is an inevitable area of inaccuracy and subjective judgment in such distinctions. Yet few would quibble with the statement that a four-stroke engine is different from diesel yet closer to it than to a toothbrush.

My point is that macroinventions are inventions that start the emergence of a new “technological species” or “paradigm”. Insofar that the notion of these groups or classes is arbitrary, the distinction between macro- and microinventions which I first proposed in 1990 has been criticized as arbitrary. While correct, this does not obviate their usefulness. After all, historical analysis cannot proceed unless we try to find similarities and distinctions between phenomena that lend themselves to that. We could add another dimension to the idea of macroinventions using the ideas above. We could think of a macroinvention as mapping from a hitherto inert or possibly novel part of S . The steam engine relied on the newly discovered principles of atmospheric pressure, and the electric telegraph on the laws of electricity and magnetism worked out by Oersted and others. Some macroinventions were, of course, singletons. Jenner’s smallpox vaccination process is a case in point but it, too, established an element in S that was quite remote from existing knowledge.

One useful way to think about the economic history of technological progress is to think of it in terms of evolutionary trajectories that begin through a sudden novelty or macroinvention, which then are continuously improved and refined through a multitude of microinventions. Those refinements eventually run into diminishing returns and asymptote off, at which point stasis is likely to occur until “punctuated” by a new macroinvention. It could be said that microinventions occur within an existing technological paradigm and are part of “normal technological change” whereas macroinventions require a stepping outside of accepted practice and design, an act of technological rebellion and heresy. It is not my contention that *every* technological tale we can tell can necessarily be reduced to this simple dynamic story. There are times, for instance, that the macroinvention proceeds through a few discrete stages. In some cases, such as the sailing ship and the water wheel, refinements were resumed and accelerated for a while after the technique had seemingly asymptoted off. But by and large, whether we are looking at power technology, chemical engineering, information processing, medicine, the metallurgical industry, or even textiles, this type of dynamic seems a reasonable characterization of the history of useful knowledge. That raises the stakes in understanding where macroinventions come from. Rather than attacking this question head-on, I propose to examine a somewhat more manageable one: under what conditions and in what kind of environment are major departures from and rebellions against existing useful knowledge more likely to occur, all other things equal.

Innovation and Environment

Evolution is inherently conservative. In all evolutionary systems, technological systems included, there is considerable inertia and constraints on change (Nelson, 1995, p. 54). In evolutionary models of technology, because of the dynamic structure of evolution in which knowledge depends on past knowledge, technical innovations (that is, additions to \mathcal{S}) are likely to be extensions and modifications of existing techniques. Most inventors and technological historians have thought of innovation as largely a set of new combinations of existing knowledge. Usher (1929, p. 11) felt that “invention finds its distinctive feature in the constructive assimilation of pre-existing elements into new syntheses” but in the revised (1954) edition of his book admitted that strategic inventions “comprise both old and new elements” (p. 68). Weitzman (1995) sets up a model in which invention is essentially nothing but a recombination of existing technology. Yet this seems confusing at second thought: it is not the techniques *themselves* (much less the artefacts in which they are embedded) that are combined as much as the knowledge underlying them. This may appear a distinction without a difference, but in the history of technology when recombinations occur, it is not the techniques that are combined but the knowledge underlying them. Only if the technique is a singleton are these processes undistinguishable. It should be noted that complementarity rarely involves “similarity” and the elements to be recombined often come from what Adam Smith already called “the most distant and dissimilar objects” -- in my terminology totally different subsets of \mathcal{S} , for instance software and hardware, or the electrical components of an automobile with the transmission.

For whatever reason, some evolutionary systems change rapidly and frequently while others remain in stasis for very long periods. In biology we observe periods of very rapid change known sometimes as “adaptive radiation.” It might be tempting to think of exogenous agents, such as “mutagens” that somehow affect the rate at which novelty occurs. In biology, such mutagens have been well-identified as chemical and physical agents that disrupt the DNA copying processes. But in knowledge systems the creative process is quite different, and it is far more difficult to identify such “mutagens.” While concepts such as mutation and recombination can perhaps be identified, the process is quite different. One property shared by all evolutionary systems, however, is that their rate of change depends not just on their ability to generate innovations as such but on the ability of those innovations to be selected and become part of the set of manifested entities.

This is particularly true when it comes to the creation of new entire groups (or classes) of entities. The ruling paradigm, based on extensive evidence in evolutionary biology, states that speciation is most likely to occur in relatively small, isolated populations. This is Ernst Mayr’s (1970) concept of *geographic (allopatric) speciation*, in which speciation occurs when a subset of a population is isolated from the main body and reproduces with each other, eventually and gradually producing genetic variability. This kind of phenomenon has no counterpart in cultural and informational evolution because concepts of homozygosity and heterozygosity (and hence recessive and dominant genes) have no equivalent outside of biology. At that level of abstraction, arguing from analogy is plainly false. But it is important to realize that the genetic structure of living beings is subject to inertive mechanisms, which all evolutionary systems need to have unless they are to slide into chaotic mode. These inertive mechanisms are set up to resist change; without them the system would clearly become unstable and likely to end up in what Stuart Kauffman (1995) has called the hypercritical region in which change becomes uncontrollable and unrestrained. In biology the

resistance shows up first in the absence (or extreme rarity) of anything that resembles a Lamarckian mechanism. A genotype is set upon meiosis. If Lamarckian change could occur, the rate of change of an evolutionary system would accelerate and stability would be unthinkable.²¹ With the Weissmanian constraints change is very rare, and resistance to change is built in at any stage. As Mayr (1991, pp. 160-161) has explained, “Just exactly what controls this cohesion is still largely unknown, but its existence is abundantly documented...during the pre-Cambrian period, when the cohesion of eukaryote genotype was still very loose, seventy or more morphological types (phyla) formed. Throughout evolution there has been a tendency for a progressive “congealing” of the genotype so that deviation from a long-established morphological type has become more and more difficult.” While such genetic cohesion has of course not precluded the well-known adaptive radiations which created different species, these explosions of variety are little more than ad hoc variations on a single *bauplan* or structural type. This cohesion, as Mayr emphasizes, while not wholly understood, is essential to the development of the world of living species: the key to success is to strike a compromise between excessive conservatism and excessive malleability. Evolutionary systems, whether biological or other, that are too conservative will end up in complete stasis; too much receptivity to change will result in chaos (Kauffman, 1995, p. 73).

Such resistance also exists in knowledge systems. Technological resistance has a number of different sources and mechanisms but it is a property of *all* evolutionary systems, including cultural ones. Consider language: neologisms, grammatical errors, and spelling mistakes are weeded out mercilessly by the red pencils of English teachers and copy-editors. Yet new words and usages, forms of spelling and even grammatical rules do eventually make it through or languages would remain immutable over long periods. It is just that only the tiniest fraction of them ever have a chance, and of those another very tiny fraction gets selected. In technology it is a direct consequence of superfecundity in the set \mathcal{S} : a lot of new ways to carry out a particular production are “proposed” or “occur to individuals” but unless the vast majority of such suggestions are rejected, the cost of continuous experimentation and change would become infinite and the system would turn into complete chaos. Even for unequivocally superior techniques, however, resistance is likely because given the finiteness of the number of techniques in use, they are likely to replace existing techniques. In knowledge systems, existing techniques are embodied in agents using them, and these agents operate as intentional and rational agents. One can think therefore readily of situations in which these agents will sustain losses if the new techniques are adopted and they are likely to resist. Even at the level of \mathcal{S} it is conceivable to think of cases in which resistance to innovation occurred because of “vested interests” in certain paradigms which through our mapping functions leads to conservatism in techniques as well. Yet when there are few direct interests at stake, and when persuasion devices such as mathematical proof, statistical significance, and experimental evidence are well-developed and widely accepted, resistance to new knowledge about nature that meets those rhetorical tests tends to be short-lived and moribund. While both Copernicus and Darwin ran into well-determined and well-organized resistance, the outcome was never in doubt. This is much less true for new techniques. In part, techniques could be untight, in the sense that it is difficult to establish their superiority. But there is far more involved.

²¹I am indebted to my colleague David Hull for this insight.

As a first approximation, the struggle against nature is not a social activity. One can imagine a solitary individual with a certain knowledge available to her, who maps from it to solve the material problems of survival. While almost immediately one can think of social dimensions of this game, the control and manipulation of nature in what we call “production” is in the first instance not a social activity but an environmental one. Once we admit that in organized groups technological activity becomes something that involves others as well, there are new opportunities but also new constraints on individual action. In part, technology becomes something of a social consensus, “the way we do things” and the individual must face up to the fact that when she wants to deviate from this norm, some measure of consent and cooperation from others has to be secured.

Every act of major technological innovation, then, is an act of *rebellion* not just against conventional wisdom but against existing practices and vested interests, and thus will normally lead to some kind of resistance.²² In what follows below I will discuss the latter in an attempt to assess the historical sources of resistance to technological innovation. In the history of technology we can distinguish a number of different sources of resistance. None of these have exact counterparts in evolutionary biology nor should we expect there to be any; what matters is that there is resistance to change.

1. *Economically motivated resistance*: groups or individuals with a stake in the incumbent technology may resist change because a switch to a newcomer may benefit other groups at their expense. Workers in danger of losing their jobs, facing changes in their work environment, or fearing that their human capital will depreciate are one example of this, but many others can be mentioned as well.
2. *Ideologically motivated resistance*: these include various sources of political resistance that are not fueled by direct economic motivation: technophobia, neophobia, a sense that meddling too much with the creation and nature is in some way sinful, or a high degree of risk aversion with particular high cost function on low-probability catastrophic events. Much of the resistance to nuclear reactors and cloning can be read this way, as do attitudes such as “we should not play God,” or “if it ain’t broke, don’t fix it.” The most obvious way such resistance takes, then, is as an ideology of *conformism* in which deviancy -- whether technological, political, religious or ethnic -- is actively discouraged.
3. *Epistemological resistance*: The flip side of the relationship discussed above between knowledge of nature and its manipulation implies that techniques might be resisted when they are seemingly contradicted by an element in *S* that currently enjoys “accepted” status., especially when it does not have a strong base in *S* itself. Thus when quinine was first introduced into Europe, it was resisted for a number of reasons, at least one of them being that it did not mesh with accepted Galenian practice (Duran-Reynals, 1946 pp. 45-53). Similarly, Dr. Barry Marshall's suggestion in the 1980s that peptic ulcers were caused by bacteria was resisted because “accepted” knowledge suggested that bacteria could not survive in the acid stomach lining. Such resistance can be overcome, and often is when the techniques are tight so that results can be readily demonstrated, as was the case with smallpox inoculation. In most cases in the history of technology the “proof of the pudding was in the eating” and

²²The literature on the subject has been growing rapidly in recent years. For a recent useful collection, see Bauer, 1995; A one-sided and popularized account is Sale, 1995. See also Mokyr, 1994, 1998b.

simple observation and experimentation were enough to persuade skeptics that even if an invention flew in the face of accepted knowledge it worked better and too bad for accepted knowledge. But when techniques are techniques are untight such as acupuncture, astrology, mind-reading and other techniques not firmly based on an accepted part of *S*, they are still regarded with great scepticism by most selectors even if they are widely used by others. The same is true by the polygraph machine which relies on a questionable foundation in natural knowledge and whose actual effectiveness, much like homeopathic medicine is controversial (Alder, 1998).

4. *Strategic complementarities.* A considerable number of technological breakthroughs in history failed to gain widespread implementation because of the absence of strategic technological complementarities. Without the right tools, the right materials, and the necessary skilled workmanship, good ideas simply could not make it from the drawing table to the prototype and certainly not from the prototype to mass production. The difference between James Watt and Leonardo Da Vinci, both enormously original and creative technological geniuses, was that Watt had first rate instrument makers and cylinder drillers at his disposal. Hot-air ballooning could not become an effective means of transportation until light-weight sources of power could be made that solved the problem of direction; electrical power could not become a widespread means of energy transmission till the problem of cheap generation through self-excitation was resolved.
5. *Systemic resistance.* As long as technology consists of individual components that can be optimized independently, changes in individual techniques depends on those of others only through the price mechanism. In other words, a change in a particular technique will drive up demand for complements and reduce that for substitutes. As long as there are no strong network externalities, it may not matter what happens to other techniques. But such externalities have always existed, even if their extent may have been limited.²³ If the costs and benefits of the adoption of a technique depend on the technique's ability to match with existing components of a given platform, the process of innovation has to take this into consideration. Technological change in a "system" becomes a coordination game which may have multiple stable solutions. Once settled on a solution, it may require a substantial cost advantage for the system to move to a different one (Loch and Huberman, 1997; Witt, 1997b). In our own age, network externalities (broadly defined) place serious limits on the degree and direction technology can change at any given time. The concept of a "lock-in" into an existing outcome is the most extreme case. Normally all that occurs that a finite but considerable cost has to be paid to make the switch. A new technology will need either to fit in with the existing system or be able to create a "gateway" technology that will somehow create a bridge between it and existing components. Software has to be "windows-capable," electrical tools require 115 V, car engines are constrained to gasoline and diesel fuels. As noted, such standardization problems can be overcome, but they impose an effective

²³It is sometimes thought that "technological systems" in T.P. Hughes's celebrated definition did not come into being till the Industrial Revolution (see for instance Edward Tenner's (1997) otherwise brilliant book, p. 13). Yet open field agriculture was clearly a complex system in which individual components such as crop choice could not be optimized independently of the whole. The same is true for the sailing ships, a complex entity in which rigging, masting, hull and steering all depended on each other and jointly determined the parameters of the vessel.

constraint on new techniques, and constitute a source of resistance. At times, such situations lead to ex post inefficient outcomes believed to require government intervention.²⁴

6. *Frequency dependence*. In many cases, the rate of technological change and the rate of adoption depends on the number of users. Frequency dependence is an example of systems with positive feedback (Arthur, 1994) and exhibit certain non-ergodic properties more common in evolutionary systems than in standard economics. Economies of scale (within a firm), external economies (among different firms), and learning by doing effects fall into that category, as do all models with network economies such as communication, such as fax machines (thus blending with systemic properties in category 5). Frequency dependence is also more likely to play a role in technological change when the technique is untight so that the benefits of an innovation are hard to observe. Selectors will look at what their neighbors do and emulate them, trying to save information costs. Social learning and imitation are thus an important source of frequency dependence. In a selection model of literary success, for example, one can easily write down a model in which more books are published than can be read by the population. Readers select books by relying on advice from friends and neighbors; obviously, the more people that have read a book, the more likely it is for that book to be read by even more people. Some books become winners through little more than historical accident. Systems with these properties do not necessarily resist change as such; indeed, some of them do nothing but change; but their change tends to follow a given trajectory and resists moving to an alternative. They are classic examples of path dependence: where the system ends up depends on the particular path it has traveled and not only on its parameters (David, 1997a). Network externalities are often the cause of positive frequency dependence (i.e., the likelihood of fax machines to be purchased depends on how many people already have one), but the two are not conceptually identical. In any case, frequency dependence means that it is often difficult to break into a market with a new product, for instance if there is no name recognition or there are no service networks. Frequency dependence can mean that one technique entirely drives out another (as in the case of VHS and beta) or allow the survival of a technique in a niche (such as two-stroke engines). It normally implies a high entry cost of a new technique or product, that is, resistance to novelty.

In short, resistance to the new exists at various levels, and if innovations are to occur at all, they have to overcome these barriers. Innovation should thus be regarded as a two-stage process: first, will the new technique be permitted to compete at all? If it overcomes this resistance, it will be tested on the merits of its own traits. The question that needs to be asked is not, why is there no more innovation, but why does innovation occur at all -- how does it succeed in overcoming the first stage barriers? There is no single answer to that question, of course. There have been inventions in history which have been so truly overwhelming in their superiority that no effective resistance could be put up. The mechanical clock and moveable type, a quarter of a millennium apart, simply swept Europe

²⁴An interesting example of a network technology that now resists change is the use of Minitel computer terminals in France, heavily subsidized and encouraged by the French government, which a decade ago was regarded as cutting-edge technology but according to *The Economist* "those inflexible Minitels are still in use a decade later, while the rest of the world has embraced the advantages of networking through the Internet. France, once a leader, now lags behind." *The Economist*, May 2, 1997, p. 18.

off its feet. Both of them were “macroinventions” by the standards described above. Among the nineteenth century inventions, the telegraph, aniline purple dyes and x-ray photography were of that nature. These advances were all quite tight: the improvement in the desired traits were easily verified and all but impossible to dispute. They did not need to fit into an existing system. But many other breakthroughs encountered resistance of one form or another. In our own age, nuclear power, high definition TV, and genetic engineering are noted examples (though the roots of resistance to each of those is quite different). The economic historian, stimulated perhaps by other cases in which evolutionary systems overcame resistance and produced sudden spurts of rapid evolutionary change, may ask what kind of environment and what type of community tends to be conducive to such change?

An application: community size and the city state

The concept of a community here is not identical to the political unit. Communities can be subsets of national units or transcend them. A community is the unit in which the fate of a new technology is decided. Are small communities, those that are “geographically isolated” to use Mayr's idea, more likely to overcome resistance than large communities in the same way that small groups are more likely to spawn successful new species? Many models of technological adoption using some kind of non-linear dynamics seem to conclude that in a variety of circumstances size as such matters in the sense that the new technique needs to satisfy a critical mass in order to overcome the inertia in the existing technological status quo (Witt, 1997b; Loch and Huberman, 1997). Is big always better? On a priori grounds, it could be argued either way. There are economies of scale in research and development; large markets create more opportunities to cover the fixed costs of increasing the knowledge base. All the same, I would argue that by and large an analysis of the *sources of resistance* delineated above suggests that relatively small and open units -- properly defined -- would have an edge over larger units in adopting new techniques. It is important to stress that this is strictly a *ceteris paribus* argument; there are many factors involved in the creation and implementation of new useful knowledge and size is only one of them. Moreover, size depends on the nature of the technology that transports people and information, and can thus only be defined in time dependent and relative terms.

The first cause of resistance, the *political lobbying of vested interests*, is clearly size-dependent. The reason is firmly rooted in the logic of collective action: it is easier for small groups to organize than for larger groups because the costs of detecting and punishing free riders goes up with the size of the group. The benefits of innovation, on the other hand, normally are distributed over a large population of consumers, hard to organize for collective action while the costs of technological progress are often concentrated among a comparatively small number of producers in a trade association, guild, or even town. One, rather oversimplified, way of looking at the success of technological progress to overcome resistance is to examine this as a political struggle between consumers and potential losers. The condition for size to be an effective determinant of the efficacy of resistance to technological change is simply whether the cost of organizing consumers falls faster with size than the cost of affected producers organizing.

Perhaps more relevant is the simple fact is that small economies tended, all other things equal, to be more open to the rest of the world than larger units. There are notable exceptions to this rule (North Korea and Myanmar today; Albania before 1992) but the historical experience up to 1945 would point to the United States, Russia and China as large economic units with a comparatively

small proportion of their GNP exported while most small economies tended to be more trade-oriented. In large part, this is simply the consequence of the trivial fact that larger countries are more diverse and thus have less need to depend on foreign trade. It has, however, an unexpected benefit: political restrictions on new techniques are obviously less effective in open economies (Mokyr, 1994b). It is more difficult to keep producers in open economies from using new techniques developed elsewhere, as they are competing with producers not subject to restrictions. Hence it is not size as such that matters here, but its correlate, open-ness. The effect, however, is quite similar.

The *ideological element in the social resistance* to innovation works quite differently. To be sure, here, too open-ness to the rest of the world will make it more difficult for technologically reactionary lobbies to use non-market mechanisms to put obstacles in front of innovation. Small, compact societies are not invariably progressive. Ethnic or religious groups and minorities, even if they were always forced to interact with the world around them, have at times displayed a stubborn ability to stick to their own conservative ideology, of which the Amish are perhaps the most prominent example. The history of the Jews points in the same direction: Jewish traditions have always been backward-looking highly respectful of tradition and precedence, and focused on exegesis of existing wisdom rather than rebel against it as invention demands. The tolerance of its society to the non-conformist and the innovator has been low. It is therefore not surprising that the contributions of Jews to the history of technology before 1850 were negligible, a fact all the more striking because of the high standards of literacy and education prevalent among them. At the same time, very large Empires such as Rome, China, the Ottoman Empire, and Russia have been victims of cultural arrogance spawning a “not-invented-here” attitude and regarding foreigners as Barbarians -- a common form of technological resistance.

Third, *epistemological resistance* may also depend on the type of community. Large communities may have in principle access to a larger body of knowledge, but in practice most empires tended to be less eclectic and open-minded. In part this was because Imperial power was viewed as supported by an entire infrastructure of knowledge and beliefs, in which religion, natural philosophy, and law all formed an integral whole on which the power structure rested. Empires tended to be more suspicious of novelty and the distance between inventors, magicians, heretics, and rebels was often a matter of subjective judgment by rulers and their officers. Such structures also existed in smaller units, of course, and while they may have been more flexible on balance, they too could at time appalling bigotry and intolerance toward heretics and deviants. The difference is that a system of small and by-and-large independent units will be less effective in carrying out the suppression of new knowledge because of inter-unit competition. In a fragmented world of small communities it is usually easier to question authority. Despite powerful defenders, Aristotle, Ptolemy and Galen were eventually dislodged during the Scientific Revolution. Dissenters and innovators could, as a *pis aller* emigrate, and in the European experience this happened a great deal. Moreover, smaller communities were on balance more exposed to different cultures and views of nature, spoke more foreign languages and might have been more willing to examine long-held beliefs if they flew in the face of new information.

Fourth, the *strategic complementarities* effect seems at first glance to work squarely against “smallness.” After all, if someone invents technique A which can only work if complemented with technique B, the larger the community in which the inventor operates the larger the chance that he or she will find the right complementarity or someone who can produce it. While this is a function

not only of size but also of the diversity of the community and its open-ness, by itself this element seems to argue against smallness. But such a notion ignores the fact that information flows are themselves a function of size. Depending on the particular way information is exchanged in a network, size can be an advantage or a disadvantage. While normally there are economies of scale in information networks, it is possible to imagine a small enough community where people interact more intensely and know more about each other even if they have contact with fewer people. Dependent on the technology of communication, the flow of information may depend on size, density, or some combination of the two. The net effect of size, in such an environment, is indeterminate. As I will argue below, relatively compact, dense units such as cities may have been optimally sized for technological advances.

Network externalities, on the other hand, are strongly correlated with size. Indeed, some network technologies tend to help *create* larger communities (think of the effects of railroads and automobiles on market integration). Almost by definition, introducing an incompatible technological element into an existing system is more costly, the larger the system. Especially when the system is defined as a standard that has to be altered, inertia seems unavoidable. If we follow the recent literature in network externalities here, each selector only picks the new technique if enough others do. If “enough” is a ratio rather than an absolute number, relatively compact units would have an advantage in reaching what is essentially a Pareto superior coordination equilibrium. Yet there are many complications here that make the connection more complex and ambiguous. For instance, system incompatibilities can be resolved by creating gateway techniques; the higher the costs of systemic resistance to a novel technique, the higher the payoff to the development of such gateways. Because the adoption of a new technique in network models often creates opportunities for free riders, in small and compact communities may be better situated to deal with them.

Finally, *frequency dependence* seems to point firmly in the direction of advantages of smallness. In general, smaller societies are more flexible when it comes to changes in highly interdependent and coordinated equilibria. If, for instance production costs are a monotonically declining function of aggregate output (cumulative or not), those costs would be higher in small communities, and this would imply that existing techniques are more deeply entrenched and more difficult to displace in large communities.

The testable implication of the above analysis is that the community that may have been optimally structured for minimizing the resistance to new technology may well have been the independent city state as it emerged in Europe in the middle ages and early modern age. I have produced evidence to that effect elsewhere (Mokyr, 1995; 1998d) and there is no need to go into it here. The argument is not an unqualified one. Not all city states were perfect instances of technological creativity and receptivity, and at times larger units displayed impressive ability to overcome the constraints of inertia. The advantage they had has to be balanced against disadvantages of scale, and the inner dynamics of the political economy of technological change often led to the eventual victory of technological reaction even in those communities. By the time of the Industrial Revolution, the optimal size of the community may have increased to include medium-sized, compact, well-integrated nation states such as Britain.

The main point here is that such an emphasis on the institutional environment in which knowledge and techniques occur is not a natural focus for standard models of economics but almost inevitable if one employs an evolutionary perspective.

REFERENCES

- Alder, Ken. 1998. "To Tell the Truth: The Polygraph Exam and the Marketing of American Expertise," unpublished ms., Northwestern University.
- Arthur, W. Brian. 1994. *Increasing Returns and Path Dependence in the Economy*, Ann Arbor, University of Michigan Press.
- Bauer, Martin (ed.). 1995. *Resistance to New Technology*. Cambridge: Cambridge University Press.
- Campbell, Donald T. 1987. "Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes." In *Evolutionary Epistemology, Rationality, and the Sociology of Knowledge* edited by Gerard Radnitzky and W.W. Bartley III, 91-114. La Salle, Ill.: Open Court, orig. pub in 1960.
- Dasgupta, Partha and David, Paul A. 1987. "Information Disclosure and the Economics of Science and Technology," in G. Feiwel, ed., *Arrow and the Ascent of Modern Economic Theory*. New York: NYU Press.
- David, Paul A. 1992. "Path Dependence and Predictability in Dynamic Systems with Local Network Externalities: A Paradigm for Historical Economics", in *Technology and the Wealth of Nations*, D. Foray and C. Freeman, eds., London: Pinter Publishers.
- David, Paul A. 1997a. "Path Dependence and the Quest for Historical Economics". University of Oxford, Discussion Papers in Economic and Social History, No. 20 (November).
- David, Paul A. 1997b. "Reputation and Agency in the Historical Emergence of the Institutions of 'Open' Science." Unpublished ms., Oxford University.
- David, Paul A. and Sanderson, Warren C. 1997. "Making Use of Treacherous Advice: Cognitive Progress, Bayesian Adaptation, and the Tenacity of Unreliable Knowledge" in John N. Drobak and John V.C. Nye, eds., *The Frontiers of the New Institutional Economics*. San Diego: Academic Press.
- Dawkins, Richard. 1976. *The Selfish Gene*. Oxford: Oxford University Press.
- Dennett, Daniel C. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon and Schuster.
- Duran-Reynals, M.L. 1946. *The Fever Bark Tree: the Pageant of Quinine*. New York: Doubleday, pp. 45-53.
- Durlauf, Steven N. 1997. "Reflections on How Economic Reasoning can Contribute to the Study of Science." Santa Fe Institute Working Paper 97-05-043.

Dyson, George B. 1997. *Darwin Among the Machines: The Evolution of Global Intelligence*. Reading, MA: Addison Wesley.

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.

Headrick, Daniel R. 1998. "Information Systems and the History of Technology in the Eighteenth and early nineteenth centuries." Unpublished ms., presented to the Northwestern University Economic history seminar.

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

Kitcher, Philip. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. New York: Oxford University Press.

Krusell, Per and Rios-Rull, José-Victor. 1996. "Vested Interests in a Positive Theory of Stagnation and Growth." *Review of Economic Studies* 63: 301-29.

Kuznets, Simon. 1965. *Economic Growth and Structure*, New York: W.W. Norton.

Liebowitz, S.J. and Margolis, S.E. 1995. "Path Dependency, Lock-In, and History." *Journal of Law, Economics, and Organization*, 11: 205-26.

Loasby, Brian J. 1996. "The Organization of Industry and the Growth of Knowledge." Max Planck Institute for Research into Economic Systems, Lectiones Jenenses, No. 7.

Loch, Christoph H. and Huberman, Bernardo. 1997. "A Punctuated-Equilibrium Model of Technology Diffusion." Unpublished ms.

Mayr, Ernst. 1970. *Population, Species, and Evolution*. Cambridge, MA: the Belknap Press.

Mayr, Ernst. 1991. *One Long Argument: Charles Darwin and the Genesis of Modern Evolutionary Thought*. Cambridge, MA: Harvard University Press.

Mokyr, Joel. 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. New York: Oxford University press.

Mokyr, Joel. 1991. "Evolutionary Biology, Technological Change, and Economic History." *Bulletin of Economic Research*. Vol. 43, No. 2 (April), pp. 127-149.

Mokyr, Joel. 1992. "Is Economic Change Optimal?" *Australian Economic History Review*, Vol. XXXII, No. 1 (March) pp. 3-23.

Mokyr, Joel. 1994a. "Progress and Inertia in Technological Change" in *Capitalism in Context: Essays in honor of R.M. Hartwell*, edited by John James and Mark Thomas. Chicago: University of Chicago Press.

Mokyr, Joel. 1994b. "Cardwell's Law and the Political Economy of Technological Progress," *Research Policy*, Vol. 23 (1994), pp. 561-74.

Mokyr, Joel. 1995. "Urbanization, Technological Progress and Economic History" in Herbert Giersch, ed., *Urban Agglomeration and Economic Growth* Berlin: Springer Verlag, 1995, pp. 3-37.

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

Mokyr, Joel. 1998a. "Innovation and Selection in Evolutionary Models of Technology: Some Definitional Issues" in John Ziman, ed., *Technological Innovation as an Evolutionary Process*. Cambridge: Cambridge University Press, forthcoming.

Mokyr, Joel. 1998b. "The Political Economy of Technological Change: Resistance and Innovation in Economic history" in Maxine Berg and Kristin Bruland, eds., *Technological Revolutions in Europe*, London: Harwood Publishers, forthcoming.

Mokyr, Joel 1998c. "Induced Technical Innovation and Medical History: an Evolutionary Approach," *Journal of Evolutionary Economics*, forthcoming.

Mokyr, Joel 1998d. "Invention and Rebellion: Why do Innovations Occur at all? An Evolutionary Approach." Forthcoming in Elise Brezis and Peter Temin, eds., *The Role of Minorities and Elites in Economic Growth*. Amsterdam: Elsevier.

Nelson, Richard R. 1995. "Recent Evolutionary Theorizing About Economic Change," *Journal of Economic Literature*, Vol. XXXIII (March), pp. 48-90.

North, Douglass C. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge: Cambridge University Press.

Olson, Mancur. 1982. *The Rise and Decline of Nations*. New Haven, CT: Yale University Press.

Pinker, Steven. 1997. *How the Mind Works*. New York, W.W. Norton.

Prescott, Edward. 1997. "Needed: A Theory of Total Factor Productivity." Federal Reserve Bank of Minneapolis, Staff Report 242, December.

Richards, Robert J. 1987. *Darwin and the Emergence of Evolutionary Theories of Mind and Behavior*. Chicago: The University of Chicago Press.

Rosenberg, Nathan. 1982. "How Exogenous is Science?" in *Inside the Black Box: Technology and Economics*. Cambridge: Cambridge University Press.

Rosenberg, Nathan and L.E. Birdzell. 1986. *How the West Grew Rich: The Economic Transformation of the Industrial World*. New York: Basic Books.

Sale, Kirkpatrick. 1995. *Rebels Against the Future: the Luddites and their War on the Industrial Revolution*. Reading, MA., Addison Wesley.

Tenner, Edward. 1997. *Why Things Bite Back: Technology and the Revenge of Unintended Consequences*. New York: Alfred A. Knopf.

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

Weitzman, Martin. 1996. "Hybridizing growth theory." *American Economic Review* 86: 207-13.

Witt, Ulrich. 1997a. "Economics and Darwinism." Max Planck Institute for Research into Economic Systems, Jena, Papers on Economics and Evolution, No. 9705.

Witt, Ulrich. 1997b. "'Lock-in' vs. 'critical masses' -- Industrial Change under Network Externalities." *International Journal of Industrial Organization* 15: 753-773.