

Induced Technical Innovation and Medical History: an Evolutionary Approach

June, 1997; revised, March 1998.

Joel Mokyr
Departments of Economics and History
Northwestern University
Email: J-MOKYR@NWU.EDU

Introduction

The motivation for this project is derived from my amazement that changes in human knowledge have been so little analyzed in the economic history literature. For most relevant problems, we tend to assume that knowledge is given and should be regarded, insofar that it is considered at all, a constraint on the maximization problem to be solved. In that approach, knowledge is much like income: for a one-period optimization problem, it is quite warranted to consider income as given and a binding constraint -- but nobody would recommend the same for a study of changes in long-term economic growth. Whereas studies of changes in income are now as numerous as ever, little is being done in the history of knowledge. In part this is because human knowledge is such a slippery concept, and most economists -- including the present author -- do not have the background in philosophy to understand the finer points of epistemology. In part, it is because even if we could agree on a definition what knowledge is, the economics of its creation and historical development violate every axiom of economic goods: it is usually non-rivalrous, often non-excludable, never trades at marginal cost, often lumpy, non-linear, non-convex, non-differentiable and externality-laden, and at times even totally fails to obey the laws of arithmetic. Neoclassical approaches to knowledge growth are therefore unlikely to make much progress. Yet economic history is unthinkable without relating it to what people knew or thought they knew, and the topic is simply too important to be left to the historians of science and technology, especially since so many of them have recently lost interest in knowledge as such and focus increasingly on the social context and political construction of knowledge rather than the thing itself.

One possible avenue to take is to adopt a Darwinian paradigm which regards the evolution of knowledge as the net historical result of blind variation and selective retention. Such an approach has enormous promise and enormous danger. At its worst, it provides empty epistemological boxes to regurgitate old concepts and well-worn facts and observations without adding much insight. Yet the evolutionary approach, when practiced by experts such as David Hull and Robert Richards, has shed considerable light on the history of scientific and engineering knowledge, and whereas it has yet to find much application in economic history, it could become a fruitful approach to a hitherto poorly developed area.

Induced technological change and the history of Medicine

The argument of this paper is simple. In its barest version, it just says that the more we know about a particular subject, the more likely it is that techniques of any kind will be able to adjust to environmental changes and thus generate induced technological change. Many years ago, Nathan Rosenberg (1976) pointed out that for a demand-induced mechanism to work in technological change, the technological capabilities have to exist. In a paper published subsequently, I relied on Rosenberg to question the importance of demand factors in bringing about major episodes of technological progress such as the Industrial Revolution (Mokyr, 1985). The point was not so much that demand played no role, as much as that as a general phenomenon, human preferences for a higher standard of living of some kind were given, and thus wide-ranging episodes of technological progress in which many diverse areas of production were affected seem unlikely to have been the consequences of exogenous demand changes or even the effects of sharp changes in relative prices for whatever reason. In its crudest form, the “necessity is the mother of invention” theory of technological change, this approach succeeds in being at once a cliché and a historical fallacy. Economists and historians alike have treated this folk wisdom with contempt (Mokyr, 1990, p. 151, n. 1).¹

A useful example of this kind of logic is to be found in the history of medicine. The idea that changes in medical technology should be regarded as a special case of technological change seems obvious enough. By medical technology I mean the techniques that prevent, cure, and alleviate the symptoms of disease. Medical technology provides a unusually fruitful ground to study induced technological progress. First, frequent exogenous changes in the environment occurred due to exogenous changes in pathogenic agents and new contacts between people and societies. Human health clearly exists in the least stable environment of any comparable variable. Its history was riddled with autonomous shocks. Throughout recorded history, new diseases appeared apparently ex nihilo and old diseases changed or vanished inexplicably. Until recently, adaptive responses were extremely slow in coming, ineffective, or altogether absent. Second, the demand side was in part biologically and not socially determined. The desire to survive, be disease-free, pain-free, and have one’s children and relatives enjoy the same seems at first glance to be more or less constant over

¹Modern empirical studies of technological advance have often claimed that much innovation is “demand-pull.” For an effective demolition of much of this work, see Mowery and Rosenberg (1979).

history.² It is therefore perhaps surprising that the history of medicine, as viewed from the point of view of the technological historian, shows remarkably little progress of any significance before 1800. Indeed, it could be argued that the ability of mankind to understand, avoid, let alone cure diseases by 1850 was little better than it had been at the time of Galen. The previous century had witnessed huge changes in the deployment of energy, the manipulation of materials, the transportation of goods and people, the transmission and communication of information, and the raising of crops and animals. Yet while the centuries since Vesalius (1514-1564) and Paracelsus (1493-1541) did witness major improvements in the understanding of the human body, these developments had little or no practical medicinal significance. If ever there should have been “demand-induced” innovation, it would have been in the avoidance of physical suffering. Yet the supply side, at least till 1850, would budge but little, particularly as far as infectious disease is concerned.

Moreover, what few improvements there were before that seem to have been not so much induced adaptations to changing circumstances as much as fortuitous events not based on systematic knowledge. Serendipitous discoveries were rarely fully exploited, did not lead to further developments and often ended up being badly applied or forgotten. For instance, Roman physicians discovered a crude form of antibiotics when they applied a mixture of rotting wood and flowers to wounds to prevent infection (Galdston, 1958). One successful medical advance of the more recent premodern age was the discovery of the Cinchona bark (quinine) as an effective cure for malaria (“ague” as it was known at the time) which became widely used in Europe in the last third of the seventeenth century. Yet the medication was applied to other fevers where it was of course ineffective. The smallpox vaccination process was discovered by Edward Jenner in 1798 but no other disease was conquered the same way for almost a century, and even today many diseases have escaped effective immunization. The successful war waged against bubonic plague through tough public policies prevented the spread of the dreaded disease. By the time of the British Industrial Revolution it had entirely vanished from Europe (Biraben, 1975-76; Cipolla, 1981). All the same,

²A more detailed look would of course nuance this picture; people's ability to deal with pain and death is subject to social influences as well as a personal hardening of feelings, and there can be little doubt that in societies in which infant mortality rates were, say 350 per 1000, the pain might be different than in contemporary society where the figures are below 10. All the same, the discomforts of a tooth ache or an allergic attack are in large part physiological, and the instincts of mammals to protect the lives of their young genetically determined.

until the closing decades of the nineteenth century European medical technology remained as ignorant as it was powerless against the bulk of infectious diseases which killed people in the West. Measures effective against the plague failed to produce results for influenza, pneumonia, typhus, or cholera. Mortality rates rose and fell more with exogenous changes in the disease environment than with medical knowledge (Goldstone, 1991).

To make any progress in the understanding of subsequent advances, we need a clearer theory of useful knowledge and its role in economic and social change. Such a theory does not exist, and the economic history of technological change has been written largely in a neoclassical competitive market paradigm or a theoretical vacuum. What I propose to do below is to sketch the bare bones of a framework in which such a theory could one day be constructed, and then show how “induced” innovation can be defined in such a theory. I will then return to the issue of medical technology as a case study of such a theory and try to show how the theoretical concepts can be made operative.

An Evolutionary Theory of Useful Knowledge

The idea that human knowledge can be analyzed using an evolutionary epistemology based on blind variation and selective retention was proposed first by Campbell and has since been restated by a number of scholars in a wide variety of disciplines.³ In previous work, I have outlined the potential of the use of evolutionary biology in the economic history of technological change.⁴ A reasonable criticism of such arguments has been that while models of blind-variation with selective retention are a useful way to look at innovations, they add little direct insight that cannot be gained from standard models. The example of induced innovation in medical technology should be regarded as an attempt to show the potential usefulness of such models. We should not think of such models as written *in analogy with* models in evolutionary biology. Instead, as I have argued elsewhere, both biological and cultural evolution are special cases of a larger class of dynamic models that share certain well-understood properties (Mokyr, 1997).

³The original statement was made in Donald T. Campbell, ([1960], 1987). The most powerful statements are made in Hull (1988) and Richards (1987). For a cogent statement defending the use of this framework in the analysis of technology see especially Vincenti (1990).

⁴For an early version, see Mokyr (1991). For more recent reflections, see Mokyr (1996, 1997).

The fundamental unit on which selection takes place is not a living being or a species as in Darwin's theory, but an epistemological one, the *technique*.⁵ 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, production widely defined (including medical and domestic technology).⁶ In the case of medicine such instructions are reasonably straightforward whether they deal with preventive medicine ("boil your water before drinking it") or curative practice ("stay in bed and drink lots of liquids till the fever has passed"). How are we to understand the Darwinian dynamics proposed by Campbell in such a model?

⁵The notion of a technique is closely related to and inspired by Nelson and Winter's idea of "routines." For an excellent discussion of the "unit of analysis" see Hodgson (1993), pp. 37-51.

⁶Note that these instructions contain a first level of interaction with the environment in that they are *conditional* instructions, so that the actual operations carried out can be made contingent on environmental conditions. The need for induced innovation occurs when the environment changes to a point not accounted for in the technique itself.

One element in this theory is the notion of the relation between an *underlying structure* that constrains but does not entirely determine a *manifested entity*. In biology, the underlying structure is the genotype, which does not respond to the environment, whereas the manifested entity is the phenotype, which does. The relation between the two is more or less understood, although there is still an endless dispute of their respective contributions of the environment and the underlying structure to the phenotype. In the history of technology, I submit, the underlying structure is the set of *useful knowledge* that exists in a society.⁷ This set contains but is not confined to scientific knowledge. It also contains traditions and other strongly autocorrelated knowledge systems which may not get down to the principles of *why* something works but all the same codify it.

The set of useful knowledge needs to be defined with some care. Useful knowledge is defined as the union of all the knowledge possessed by individuals that can conceivably be applied to production in its widest sense (including household activities) in a given society. This knowledge is confined to the natural world: knowledge about the epistemological philosophy in Biblical texts, say, does not count while knowledge of the planets of Jupiter does. There is a certain arbitrariness about this, and in some grey areas the line may not be as sharp as we like. All the same, because techniques always and everywhere involve the manipulation of natural regularities, this seems a natural definition. Such useful knowledge does not have to be actually applied or even be, in some definable sense, correct. Knowledge could well be a set of untested beliefs and prejudices that posterity will eventually reject.

Knowledge can reside in people's minds and in storage devices with greater or lesser accessibility. Leaving out non-human storage devices, let there be n members of society and let each individual in a society possess technical knowledge S_i . Let $N_i(S_i)$ define the number of "useful" pieces of information contained in the set S_i . We can then define the total knowledge of society as the union of all the individual knowledge of members of society.

(1)

OMEGA $\sim = \sim$ BIGCUP $S_{\text{sub } i} = S_{\text{sub } i} \chi S_{\text{sub } j} \sim$ for \sim FORALL $\sim i \neq j$ and

⁷The term is used more or less in this sense by Kuznets (1965), pp. 85-87. Kuznets confines his set to "tested" knowledge that is potentially useful in economic production. In what follows below, this definition is far too restrictive. An enormous number of techniques actually in use, from bloodletting to crop rotations to modern slim-down diets were based on knowledge that was untested and often demonstrably ineffective.

(1)

$$\text{PHI} \sim \sim N(\text{OMEGA})$$

M represents the total number of pieces of useful knowledge possessed by society. When a single individual produces an innovation (that is, discovers something hitherto unknown about nature), we observe unequivocally an increase in Σ . In a biological sense, Σ can be thought of as the gene pool. Σ is a union over n members, although it is very likely that the knowledge of many individuals is redundant in that their knowledge is wholly subsumed in that of others (so that their removal from society does not reduce Σ).⁸ To define diffusion, we simply look at the intersection of S_i and S_j , which is the amount of knowledge that two individuals share, $S_i \cap S_j$.

The set Σ maps into a second set, the manifested entity, which I will call the feasible techniques set δ . This set defines what society *can* do, but not what it *will* do. The mapping function, in essence, translates the knowledge about natural facts and regularities into “how to” blueprints that actually can manipulate nature into improving the human condition in some way (that is, produce a good or service). The outcome is then evaluated by a set of selection criteria that determine whether this particular technique will be actually used or not, in a fashion similar to selection criteria that pick living specimens and decide which will be selected for survival and reproduction and thus continue to exist in the gene pool. The analogy is inexact and to some extent forced: while genes are a mechanism of inheritance and will vanish as soon as the species is extinct, knowledge can continue to exist even if the techniques it implies are no longer chosen. All the same, the bare outline is quite similar in that the dualism between the underlying structure and the manifest entity is maintained. Above all, selection can only pick entities from existing material; variants that cannot be constructed from existing potential will not appear, no matter how beneficial or desirable.

⁸ In general, increasing Σ will lead to higher M , but because adding to the amount of useful knowledge might also make some previously useful knowledge obsolete, the relation between the two is complex. For the purpose here, this ambiguity is not fatal. Note that both in genetics and in technological knowledge only a small fraction of actual existing potentially useful information is actually “switched on.” The human gene uses only about 1 percent of the DNA; the rest seems to fulfill no obvious function, but changes in it may at some point in the future become useful.

To stick with the example chosen for this paper, an example of an underlying structure in medical technology is the humoral theory of disease which viewed all diseases as resulting from imbalances between the four basic bodily fluids, blood, yellow bile, black bile, and phlegm. This theory, propounded by the Hippocratic School of Medicine and part of the Galenian canon, implied that certain techniques be used by physicians on their patients, the best-known of them being bleeding and purging patients suffering from fever. The technique was then examined against alternatives and chosen by physicians as their main weapon against infectious disease for many centuries. Note that there are a number of distinct stages in the translation of knowledge into practices. One of them is the mapping itself, which may differ from place to place and over time. Another is the selection criterion by which the efficacy of a technique is tested. Interestingly enough, even when the belief in the humoral theory underlying bleeding practices waned, the technique stubbornly continued to be used until deep into the nineteenth century. The example also illustrates the many difficulties that such a view implies for optimistic scenarios that are based on the hopeful but ahistorical assumption that rational selection processes will eventually yield an outcome that we may recognize as “efficient.” That, however, is another story.

Operationalizing Σ is, of course, quite hard, since it skirts questions such as “who knows that which is known” and “how easy is the access to this knowledge.” What matters here is that Σ maps into sets of instructions each of which constitutes a technique which jointly make up Σ . Each technique has “traits” that define it. Call these $T_1 \dots T_n$. We can think of those, say as “output quality” and “costs of production” although many different product attributes may matter here. Secondly, each time a technique is “used” it “lives” and a specimen has been “selected.” For each technique j , we may then define μ_j , which is a count of how many times this technique is used, a bit

(2)

$$\mu_j \sim \text{SUP}^*(j) \sim \text{S}(T_1^j, T_2^j, \dots, T_n^j; V)$$

like the size of a population. The fitness equation then defines the basic motion of the system:

where μ^* is some equilibrium level of usage and μ

$\dot{\mu} = f(\mu - \mu^*)$, $f > 0$ is the change in the frequency of use. For any V (the environmental parameters, assumed exogenous), 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 . Each technique is defined as a point in this space. We can then define the curve ZZ' in fig. 1 which defines the condition of fixed fitness (μ




Figure 1

$\dot{\mu} = 0$). An exogenous deterioration in the environment (possibly due to changes in complementary or rival techniques, changes in preferences, or some other autonomous effect) would be depicted as an

outward shift of ZZ' . In addition to the techniques in use, given by the area $*$ in fig. 1, there is a larger set of all feasible techniques δ within which $*$ is wholly contained. The techniques that are in δ but not in $*$ are techniques that are feasible but not selected by society.⁹

Selection of techniques thus occurs at two levels: not all techniques in δ are picked to be in $*$. In fact only a small minority of all feasible 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 with each other, and in the long run, assuming competition is sufficiently stringent, only the ones that are at E_0 actually maintain their numbers. Such an equilibrium, however, is not a prediction of the model. A lot depends on the actual degree of competition, and without more information it is not possible to know if other techniques in δ survive or go extinct. Moreover, points like E_0 are not necessarily unique: the set δ need not be shaped neatly as in fig. 1 and could have more than one tangency with ZZ' . Multiple equilibria and path dependence are, of course, standard fare to both students of economic evolution and the history of technology.

Beyond the level of the selection of techniques within δ , there is a higher level of selection at the level of knowledge in Σ , precisely of the type that evolutionary epistemology addresses (Hull, 1988, ch. 12). It might well be asked why selection is necessary here, since in general knowledge need not displace previous knowledge. If storage costs are sufficiently low, knowledge could just accumulate. The principle of superfecundity (more specimens are born than can be accommodated which lies at the center of Darwinian thought), does not apply here *strictu sensu*. In practice, however, new knowledge often replaces existing knowledge that must be discarded. Thus, accepting the work of Lavoisier meant that one had to abandon phlogiston theory. Such obsolescence is not the same as extinction in the living world, however, and at times knowledge that was believed obsolete can be resurrected. Knowledge can thus be active or dormant, depending on whether \mathbf{i} maps onto δ , regardless of how widely it is accepted.¹⁰

⁹The shape of δ and $*$ as neat and compact shapes is of course not required: they could well be highly irregularly shaped with multiple tangency points with ZZ' corresponding to multiple techniques in use with similar features serving similar purposes.

¹⁰An 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


If the social costs of retaining information are essentially nil, a case could be made for the technological equivalent of biodiversity, that is, preserving seemingly useless old forms of knowledge. Much of this depends on the technology of information storage and the access costs. In effect, that is what postmodernist history of science is trying to do. Rather than reject or accept older forms of knowledge as “correct” or “false” it tries to see them in their social context, and has as much interest in innovations that ended up dead alleys as in those that led further progress and more successful forms. Yet unlike some of the more extremist versions of postmodernist history of knowledge, such agnosticism is not *invariably* useful: we need to reject some knowledge in favor of other and a nihilist view that all knowledge about nature is just as interesting as any other knowledge (Kitcher, 1993). Even Bruno Latour would not want to be treated by a physician holding on to the humoral theory of disease, or the belief in the spontaneous generation of micro-organisms, or the notion that ulcers are caused by stress. Knowledge will be rejected if it is widely believed to be false *and* does not map into any useful technique. It is also often suppressed or delayed by the adherents of an alternative and incompatible set of knowledge, who have more political power.¹¹ Unfortunately, we can not always tell one form from the other. Unlike genetic information, however, it rarely becomes *irreversibly extinct*. If we wanted, we could revive medieval herbal treatments or François Broussais’s (1772-1838) notorious use of leaches based on his absurd theories that all diseases originated from the digestive tract just as we can build Roman Catapults and Chinese waterclocks.¹²

The way subsets of Σ are selected for is largely by persuasion. Selectors who advocate some subset of Σ try to make others see the same. Persuasion contains a large number of rhetorical means such as the proof of mathematical theorems, statistical analysis, experimentation, and other forms of induction. It also contains authoritarian obiter dicta, threats, education, propaganda, and political manipulation. Consequently in many cases useful knowledge came into existence but failed to be

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 Jones (1993), p. 223.

¹¹An interesting recent case is the discovery that peptic ulcers are caused by *Helicobacterium Pylori* (and not by stress), made by the Australian physician Barry Marshall in 1983 and ignored for close to a decade by skeptical opponents or those with vested interests in the status quo. See for instance “Why Doctors aren’t curing Ulcers,” *Fortune*, June 8, 1997, p. 100-07.

¹²Indeed, the increased demand for leaches in certain surgical purposes (especially in the re-attachment of severed limbs) illustrates this point.



“activated,” that is, failed to persuade enough people to end up being translated into the set of feasible techniques.¹³ Active knowledge is not at all the same as “true” knowledge. Much of the knowledge set of the past may be recognized today as “false” and yet have mapped into a useful technique. One can navigate a ship using stars even under a Ptolemaic geocentric astronomy and improve the techniques of iron production on the basis of phlogiston chemistry. Medical procedures based on Galenian theory could be effective. In short, whether knowledge is “correct” or not seems to matter less than whether it mapped into techniques. Selection thus occurs at both ends of the mapping $\Sigma \equiv 8$, although the selection process is different.

Unlike the natural selection mechanism defined by Charles Darwin, in knowledge systems the selection process is not anonymous and decentralized but conscious and deliberate. Techniques are chosen willfully by individuals who are trying to attain certain objectives and tested according to pre-specified criteria. However, for a technique to succeed by being chosen is *not* to say that it actually maximizes those objective functions, let alone a social welfare function. In that regard, the historical success of a particular technique may be quite different than what an ex post assessment of its fitness would imply. This is particularly true for medical techniques, but there are many other instances of technological choices in which the selection process chooses a technique that

¹³Thus the germ theory was first proposed by Girolamo Fracastoro in 1546 and proposed repeatedly without having influence on the practice of medicine until late in the nineteenth century. The sad case of Ignaz Semmelweiss who discovered the need for sterilization of medical tools through the connection between the contamination of physicians and the death rates in maternity wards due to Puerperal Fever, yet whose work was ridiculed and ignored for twenty years is another case in point.

seems inferior or -- more often -- rejects a technique that ex post was more efficient. Resistance to innovation is one of the more interesting features of evolutionary systems.¹⁴ A technique's fitness may thus be judged by two criteria: its success in actually being selected and the way it fulfilled the function for which it was intended. The convergence or divergence of these two criteria are still being debated.

Technological adaptability and Induced Innovation

Despite the rather confining definitions of information and other simplifications, this setup allows us to make some simple distinctions in an evolutionary framework. For one thing, it allows us to define adaptation and adaptability which is crucial to the idea of induced innovation. To start with, most techniques have a certain amount of built-in flexibility. The reason is that like genetic instructions embedded in the DNA, technical instructions often take the form of if-then statements, conditional on a variety of previously experienced or easily predictable contingencies. This allows the technique to adapt at a local level, but such "phenotypical adaptations" are providing only limited flexibility.

What really counts is adaptation to something unexpected requiring a change in technique. Consider fig. 2. Suppose we are looking at a case we described earlier, in which the knowledge of society in medicine is very limited, but it has discovered that a certain technique works. Denote this technique by E_0 . I will call this a *singleton technique*. Such techniques are based on knowledge that does not extend beyond "such and such a technique works." Under the assumptions stated, the singleton E_0 is the only technique in the feasible set. The accidental discovery of a medicinal herb would be a good example of a singleton. The knowledge set Σ contains the knowledge that such and such an herb is effective against a certain disease. This maps into the feasible techniques set as a single set of instructions: if patients exhibit this symptom, give the patients the medication. The knowledge does not contain any pharmacological basis of the herb, any information about the disease's etiology or any hint of its modus operandi against the agent causing the disease. A similar structure is true for most premodern production. In agriculture, fields were fertilized without any

¹⁴The 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, 1998).

underlying knowledge of organic chemistry. Steel was produced for centuries by blacksmiths who had no knowledge about the relation between carbon content and the physical qualities of iron. Much of what we call production was carried out on the basis of “standard operating procedures” passed from generation to generation. The “underlying structure” of knowledge was little more than purely pragmatic: “such and such works.” This kind of knowledge was acquired by trial and error and passed on from master to apprentice. When the criteria for efficacy were unclear or the testing procedures flawed, ineffective techniques could be adopted and survive for centuries.

The knowledge base of a controversial technique may be one determinant of its success in the political arena. Joseph Lister, who revived Semmelweis's discarded antiseptic techniques could rely on the recent discoveries of Pasteur to defend his insights. Less well-known is the rather stiff resistance against smallpox vaccination by a host of enemies, who exploited the fact that the causation of the disease was unknown and no one had any clue as to the modus operandi of Jenner's successful vaccination technique. It was even disputed whether the disease was contagious, much less understood how it was transmitted. This was particularly true for the United States where a new epidemic of smallpox flared up in the 1870s after the disease had all but disappeared following the successful vaccination campaigns in the 1810s and 1820s. Many states repealed their mandatory vaccination programs, and for decades the feared disease returned with often devastating effects (Kaufman, 1967).

Another example is the conquest of scurvy. The importance of fresh fruit in the prevention of scurvy had been realized even before James Lind published his *Treatise on Scurvy* in 1746. The Dutch East India Company kept citrus trees on the Cape of Good Hope in the middle of the seventeenth century yet despite the obvious effectiveness of the remedy, the idea obviously did not catch on and the idea “kept on being rediscovered and lost” (Porter, 1995, p. 228). Lind's ideas assisted Captain Cook in keeping his crew scurvy-free. Yet modern scholarship has established that Cook's efforts only confused the understanding of the disease and delayed rather than hastened the solution (Carpenter, 1983, p. 83). On his voyage, Cook, determined to eradicate the disease, tried a number of different things, and it was difficult to attribute the disappearance of the disease to a specific measure. Even after the curative properties of lemon and lime juice were recognized, it was still thought that the disease itself was caused by breathing foul air in the ships's living quarters. In

1795, Gilbert Blane made the use of oranges and lemons in the British Navy mandatory and scurvy was dramatically reduced. Yet precisely because this was a singleton technique, its persuasive force remained weak. Consequently, while controlled on shipboard, scurvy remained a serious problem on land: it survived in jails and poorhouses, and made a serious appearance during the Famine of 1845-48. It was still endemic during the Crimean and U.S. Civil Wars and the Russian army during the first World War. Infantile scurvy was prevalent among wealthier families in which weaning occurred at a relatively early age. The discovery of the germ theory led to decades of futile search for a causative microorganism. Only after the seminal papers by Holst and Fröhlich after 1907 did it become clear that certain diseases were *not* caused by infectious agents but by nutritional deficiencies, and only in 1928-32 was the crucial ingredient isolated (French, 1993).

If the environment changes exogenously in any form as indicated by a change in the slope of ZZ' , the primitive production system in which the feasible set is the singleton E_0 cannot adapt and stays at E_0 to its detriment. A singleton technique means that there will be little or no adaptability in the system, so that minor environmental changes can cause very significant losses in fitness. A dramatic example is the appearance of bubonic plague in Europe. While in the very long run adaptations were developed, it took close to three centuries for the disease to disappear.¹⁵ More fortunate were the Europeans of the nineteenth century, who were visited by the Cholera. While ignorance of contagious disease in 1829 (the date of the first European appearance of the disease) was almost as deep as in 1348, scientific method had progressed and rational investigation based on established procedures was quite different. By the 1850s the mode of transmission of cholera was understood (before its etiology, to be sure), and eventually the disease disappeared again from Europe. The germ theory of disease, arguably one of the most significant increases in Σ in history, mapped into thousands of small and large techniques to avoid infection both in the domain of public health and in that of household technology, long before antibiotics were developed (Mokyr and Stein, 1997). The increase in microbiological and immunological knowledge in the past decades has been pivotal in our ability to deal with HIV; it is hard to imagine how any medical adaptation would

¹⁵Even more devastating were the appearance of European diseases such as smallpox and measles in the American Continent after Columbus annihilated most of the native population. The appearance of syphilis in Europe in 1494 (in all probability imported from the New World) had at first devastating effects, but the disease changed its nature in later years and became less fatal.

have been possible had HIV appeared, as Cholera did, *ab nihilo*, in the middle of the nineteenth century. Adaptation to a new disease does not require necessarily a cure or a vaccine: in the case of infectious disease the critical piece of knowledge that is required is the understanding of the mode of transmission. Once those are understood, even imperfectly, preventive techniques may have been enough to deal with the disease.¹⁶ One common view of the main change in the process of technological change is that since the late nineteenth century, engineering and the “knowledge of production” have been far more closely connected to science than previously (Copp and Zanella, - 1993). What this implies is that the δ “around” modern techniques actually in use is much larger, that the Σ is far more capable of producing new techniques, and that the mapping from Σ to δ is far more flexible and capable of producing novelty “on demand.” Yet the “trial and error” and “try every bottle on the shelf” modes in invention have not disappeared, especially in pharmaceutical and biological technology in which many of the underlying processes are very complex and poorly understood.

Flexibility and adaptability, then, have three dimensions. First, the larger the set of feasible techniques δ , the more adaptability the production economy has. If the environment changes from ZZ' to WW' , the system can adapt (as from E_0 to E_1). This is precisely what we normally mean by substitution. The technique E_1 may have “existed” at the time that WW' was in force, in the sense that it was part of the feasible set δ but not “selected.” In other words, the society was *capable of* producing E_1 given its knowledge base but chose not to. When circumstances changed, it adjusted by moving along the frontier of δ . Some interpretations of technological progress in economic history are based entirely on this concept: Ester Boserup (1981) has argued that essentially new knowledge is unimportant and what governs the technique in use is population density. When population increases, society will find labor intensive techniques that will keep living standards from falling.

Second, environmental change may also lead to new searching over Σ and new mappings into δ creating techniques that were previous unknown even if they could be known given that the

¹⁶An example is the idea that diseases were transmitted by vectors. For centuries it had been believed that the association of swamps with malaria was caused by the “bad air” that emanated from standing water. The work of Patrick Manson, Ronald Ross, and G.B. Grassi demonstrated the culpability of the *anopheles* mosquito in the 1890s and in 1909 Charles Nicholl discovered the louse vector of typhus, five years before the causative germ itself was isolated. These discoveries were decisive in persuading households how such diseases were contracted and thus could be successfully avoided.

knowledge base for them existed. In some sense this approach redefines the traditional distinction between substitution and induced technological change.¹⁷ Substitution in the standard microeconomic approach concerns a choice between *known* techniques. Induced technological change leads to the emergence of new techniques based on the existing stock of useful knowledge. What really is known is not just a set of blueprints that firms and individuals can pick and choose from freely, but an underlying knowledge set, far more complex and multidimensional. As long as the knowledge set exists, it is possible for society to adapt to changing environment and innovate as needed by mapping from an existing subset of Σ onto a new segment of δ .

The existence of such knowledge is not sufficient for the mapping to occur; nor is an exogenous stimulus due to a change in the environment a necessary condition. The mapping function is one of the more difficult concepts in the history of technological change. How accessible to the searcher is knowledge available to somebody else or stored somewhere? Does existing knowledge translate itself into a new useful technique when the need arises? Historical cases in which that happened, especially in the last century, can be found -- one thinks inevitably of the German Fritz Haber's invention of the nitrogen-fixing process in the face of his beloved motherland's needs to produce fertilizers and explosives during a naval blockade. But exceptions are difficult to explain and are numerous enough to doubt any regularity. There are examples of the mapping of existing knowledge into a technique occurring without any obvious stimulus. One is the invention of spectacles around 1280. The basic elements of the knowledge set -- knowledge of glassmaking and the observation that lenses change the refraction of light -- had existed in Roman times.¹⁸ It seems hard to believe that a sudden change in demand occurred in the thirteenth century: the physiological changes that cause the need for eyeglasses are more or less constant. The development of printing or, more accurate, of moveable and interchangeable type, by Gutenberg runs into a similar dilemma. The famously serendipitous elements in the invention of antibiotics illustrate the precariousness of

¹⁷The 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.

¹⁸Seneca had already observed that letters were enlarged and made more distinct when viewed through a glass globe. Alhazen, who lived around 1000 AD studied the reflection of light from curved mirrors and spheres, yet spectacles were invented in Italy only toward the end of the thirteenth century.

any easy conclusions about the mapping from Σ to \mathcal{T} . All the same, it is clear that without the concept of pathogenic bacteria, the widespread development and adoption of antibiotic techniques would have been absurd. In areas in which the knowledge base developed slower because it was more complex, such as in viral, autoimmune and psychiatric diseases, progress toward effective cure was much slower and the techniques used are far more inclined to be singleton techniques based on trial and error or serendipitous finds.

One obvious factor in the mapping function is the *accessibility* of knowledge. Knowledge may exist, either in someone's mind or in some storage device, but a great deal depends on the ability of those who perceive the need for it to find and access it. Another factor is the *technical capability* of society to design and build techniques that its knowledge base suggests might be feasible. In many cases a particular technique is imagined, envisaged, and even designed, but a critical component or complement is missing which makes it impossible to exploit it. A good example is the measurement of longitude, which has recently been popularized thanks to Sobel's (1995) excellent little book. It had been understood how longitude could be measured (by the use of two accurate clocks), but it turned out to be very difficult to construct marine chronometers of sufficient accuracy until the technical difficulties were cracked by John Harrison in the middle of the eighteenth century. Much like longitude, the exploitation of fusion energy in our time seems to elude us despite the knowledge that such energy is possible in principle and probably in practice. The practical application, however, has not materialized. Similarly, President Clinton's recent announcement that AIDS will be either curable or wholly preventable within ten years is based on a sense that the solution is within reach and that only a few elements are missing before the puzzle is wholly solved.

Third, we can think of *induced knowledge change* differing from induced technical change in that changes occur in the knowledge set Σ rather than in the set of feasible techniques \mathcal{T} . It might be thought that this is a distinction without a difference. The extension of the knowledge base Σ itself is the underlying force, believed to propel human progress. Again, there exists a grey area where the distinctions are blurry. Yet some insights can be gained from it.

Knowledge growth does not develop wholly exogenously. It responds to outside stimuli and search processes and can in that sense be said to be "induced." Scientists do not pick topics at

random, they work on problems they feel that other scientists or some patron may be interested in. Knowledge is thus constrained by its own past: the direction of change depends on the state of the world at each moment. In that regard, knowledge can be said to follow an evolutionary path like a Markov chain in which normally innovations are incremental rather than revolutionary. While there is some disagreement among evolutionary theorists as to the likelihood of very rapid, discontinuous evolutionary changes, even “saltationists” realize that there are limits to the amount of change that can occur per unit of time. In that sense, all learning is “local.” To be sure, some innovations are less localized than others, and at times we observe the birth of something that is radically novel in that it represents knowledge that simply was not available before. It may be argued ad infinitum to what extent Pasteur’s famous refutation of spontaneous generation or the appearance of *The Origins of Species* were “local” or “radical” innovations. We can all agree, however, that they were not likely to have been produced in the age of Thomas Aquinas. Yet the “need” or “demand” for them existed as much in 1720 as it did in 1860. People were just as sick and arguably as curious about the development of living beings. The germ theory, of course, had been proposed earlier but lost out in the battle of persuasion.¹⁹ Darwin’s insight, while shared with Alfred Russel Wallace, simply had not occurred to anyone before, triggering T.H.Huxley’s famous response “how very foolish not to have thought of this.”

This framework allows us to classify all additions to knowledge into three broad classes: new knowledge that will be selected under the current environment; new knowledge that is potentially useful if circumstances change but currently is neutral; and detrimental mutations that will never be useful. In fig. 3, which is an elaboration of fig. 2, I depict three different increments to knowledge (note that the diagram is drawn for convenience in the space of techniques, but that we really should think of it more in terms of Σ than in terms of δ). The increment \forall is a classic favorable mutation or invention, in which the traits of the entity “improve”, allowing it to become fitter and to augment its numbers. The selection mechanism will choose any technique displaying the traits in this set. The increment \exists is useless in the current configuration. As long as T_1 and T_2 are positive traits (that is, have positive partials with respect to the objective function) nothing based on the information in it

¹⁹Jacob Henle, the main proponent of the germ theory in the 1840s was regarded as “fighting a rearguard action in defense of an obsolete idea.” In that regard his student Robert Koch was more successful. Cf. Rosen (1993), p. 277.

will ever be selected. The mutations in $($ are neutral in that they do not affect the phenotype. Yet they could come in handy when the environment changes in such a way as to favor the need for more T_2 as depicted by the curve WW' . Such environmental changes could be changes in factor prices or resource availability, but should be interpreted as including the appearance of a new trait T_3 complementary to T_2 which could “activate” the region $($. A great deal of scientific and mathematical

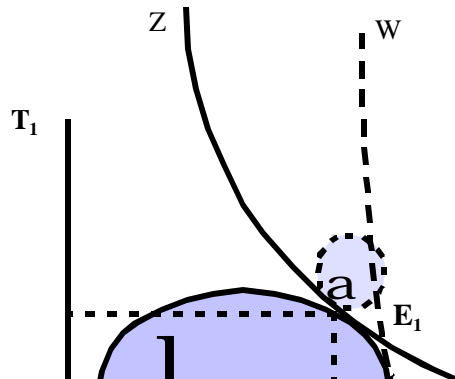


Figure 3

knowledge that seemed useless at the time became useful much later when complementary knowledge became available. Boolean algebra and Hellenistic astronomy, to cite just two examples, eventually became indispensable to technological developments. The “lucky” economy is the one which develops the neutral knowledge $($ and is able to access it when circumstances change to move ZZ' to WW' .

Path-dependence implies the existence of multiple equilibria, and in the history of medicine this means essentially the development of alternative medical paradigms that emerged more or less along independent trajectories and whose ultimate form was shaped as a result of the history of the system as much as by the objectives that practitioners tried to attain and the constraints on them (David, 1997). The importance of “alternative” medicine, from homeopathy to Christian Science is

evidence that multiple approaches to disease are still enormously influential. Tens of millions of Americans resort regularly to one form or another of alternative medicine. The historical development of Chinese and Western medical knowledge along separate historical paths has produced different techniques and thus trajectories for change. Some of those trajectories have better options in dealing with changes in the environment than others. For instance, it seems reasonably uncontroversial that in dealing with the AIDS infection, modern virology and immunology are better equipped than acupuncture or chiropractors. But in other syndromes that seem to have emerged or worsened in recent years, from skeletal pain caused by exercise to lower back pains to sleeping disorders to Tunnel-Karpel syndrome, this superiority seems less secure. In medical conditions that are psychosomatic in nature, a set of techniques based on the placebo effect, including the hypnotic effects of witch doctors, the touch of holy men, or the couch of the psycho-analyst maybe better suited to deal with outside shocks than the chemical and surgical interventions of modern medicine. Hence the sheer size of the entire body of medical knowledge, including various forms of non-standard medicine provide more adaptability and ability to respond to unexpected shocks.

How do changes in Σ affect what people do and how they produce? Despite dissimilarities with living systems,, the process of generating innovations I am proposing here shares some important features with the way living beings do. It may well be the case that a substantial amount of knowledge is “induced,” that is, created with a specific purpose in mind. However, the vast majority of all human knowledge, like DNA, is non-coding or “junk” in the sense that it does not apply directly to production. Most scientific (let alone other forms of) knowledge has no applications and does not affect production technology right away although it may be “stored” and in rare cases called into action when there is a change in the environment or when another complementary invention comes along. Thus most additions to Σ , like mutations, are predominantly “neutral” and are not affected by the selection criterion one way or another, but may become useful when the environment changes and calls for adaptation (Stebbins, 1982, p. 76). The activation of such previously inert material may be the evolutionary equivalent of what economists think of as “induced innovation.” General knowledge, too, is being created at a rate much faster than technological knowledge, but if it finds no application in production, it is not a part of technological evolution and might be regarded as “neutral.” In this important sense, neutral change creates a powerful element of contingency in

history: if the knowledge had happened to accumulate for no special reason at some point in the past, it would come in handy when there is an autonomous change in circumstances. In that very sense historians of technology might well adopt the term proposed by Motoo Kimura, the proponent of neutral drift in evolutionary biology, who has suggested replacing “survival of the fittest” by “survival of the luckiest” (Kimura, 1992, pp. 225-230).

Above all, however, there seems to be no obvious social mechanism that predicts endogenous changes in the size of the knowledge set Σ any more than we can predict what direction changes in the gene pool will take. Instead, relative prices and other focusing devices in Rosenberg’s phrase may determine *the direction* in which the search for new knowledge takes place, and thus in some cases determine the successful increase in the set of useful knowledge. In that sense they are like a steering mechanism; but it is not the steering mechanism that makes a vehicle move forward and is the primary determinant of its speed. The expansion of the useful knowledge set does not respond neatly to changes in demand because in some sense this demand is always there. This is not just true for medical knowledge even if there the case can be made most persuasively. We can point again and again to societies that “needed” knowledge but did not get it simply because it was out of their reach, or because they looked in the wrong places, or because it did not occur to them to look at all. Classical civilization, an iron-using and seafaring society made practically no changes in the primitive processes of ironmaking, shipbuilding and navigation extant around 400 B.C. for the next half millennium. Ironcasting which had been known to the Chinese since the 2nd century BC reached Europe in the fourteenth century. There are obviously many brakes and obstacles to the expansion of knowledge and even more so to their mapping onto the useful techniques set.

Concluding Remarks

It seems therefore that the ability of the knowledge set to respond to environmental changes depends in large part on the nature of the existing knowledge itself and on the conditions that determine how conducive the knowledge base and society are to its expansion. The more variegated and diverse the knowledge set, the more likely it is to be able to create a response to outside shocks. Useful knowledge is to a large extent cumulative, depending mostly on storage devices and the cost of access. That does not mean that its growth is always harmonious. Thus “modern” medicine regards alternative forms with suspicion bordering on the disdain, but so do alternative forms each other.

New knowledge often encounters stiff resistance from existing forms, when vested interests who have heavily invested in the latter fear the rapid depreciation of their specific human capital. This is equally true of other natural sciences, but there modern science often -- if not always -- creates the conditions to test different paradigms against each other. All the same, the diversity of the forms of knowledge means a great deal of flexibility and ability to deal with different shocks and needs.

In short, then, past knowledge has developed more or less by its own rules and thus any induced responses to environmental changes at the level of Σ were not very significant. Even in our time, I would insist, it is still true that changes in useful knowledge have remained to a great extent an exogenous variable (though clearly less so than in the past) and that any attempt to endogenize -- let alone predict -- it is foolhardy. True, if society faces a well-defined and clear-cut problem, it can allocate more resources toward increasing the knowledge set. The “research” part of research and development, unique to the twentieth century, is precisely the conscious attempt to expand Σ in a “needed” direction. It may, therefore, be the case that the modern age will develop a dynamic of knowledge completely different from previous ages. Yet the precise form this will take is still quite unclear. The success of “project-Manhattan” endeavors is hardly the rule. The relatively modest gains in the war against cancer, declared with pomp by Richard Nixon, attests to the reality of real knowledge constraints. No amount of real resources devoted to medical research would have helped European society in 1348 to solve the riddle of the Black Death. To be sure, the war against HIV is conducted more effectively thanks to the great breakthroughs of our own time, information processing, molecular biology, and genetic engineering. Yet these breakthroughs themselves can hardly be regarded primarily as demand-induced as most of their uses only became obvious *ex post*.

More seriously, to carry out such research, the problem has to be formulated correctly and for that prior knowledge is required. Between 1875 and 1890, bacteriologists focused their research on the discovery of pathogenic bacteria and tackled them at the rate of about an organism every two years. But this program required a prior knowledge that such a search would indeed have a reasonable probability to yield results. Consider the question of why in our age we devote so few resources in our search for substance that would halt or reverse the aging process, the ultimate Faustian dream. It can hardly be argued that no demand exists for such a substance, or that the demand for it has not increased steadily with the rise in the average age of the population. Yet the resources devoted to

such research are rather modest simply because few scientists believe it likely that such a substance can be found. In the past such searches, despite being well-focused, have often failed: alchemy is perhaps the most striking example. Alchemy was based on the false analogy of the change in the physical properties of compounds to changes in elements. It may well be that the search for an AIDS vaccine is based on the false analogy between a highly stable virus (Polio) to one that is genetically unstable (HIV) and that no vaccine is feasible at all -- or that no vaccine is feasible at current knowledge of the molecular processes involved. Rosenberg (1976, p. 51) cites Henry IV, pt I in which Glendower says that he "can call spirits from the vasty deep" and is met by the deadly response of Hotspur: "why, so can I, or so can any man; but will they come when you do call for them?"

REFERENCES

- Bauer, M (1995). *Resistance to New Technology*. Cambridge: Cambridge University Press.
- Boserup, E (1981). *Population and Technological Change*. Chicago, University of Chicago Press.
- Campbell, DT [1960] (1987). "Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes." In Gerard Radnitzky and W.W. Bartley III (Eds) *Evolutionary Epistemology, Rationality, and the Sociology of Knowledge*, 91-114. La Salle, Ill.: Open Court, orig. published in 1960.
- Carpenter, K (1986) *The History of Scurvy and Vitamin C*. Cambridge: Cambridge University Press.
- Copp, NH , and Zanella, AW (1993). *Discovery, Innovation, and Risk*. Cambridge, MA: MIT Press.
- David, PA. (1997). "Path Dependence and the Quest for Historical Economics". University of Oxford, Discussion Papers in Economic and Social History, No. 20.
- French, R (1993) "Scurvy." In Kenneth F. Kiple, (ed.), *The Cambridge World History of Human Disease*. Cambridge: Cambridge University Press.
- Galdston I (1958) , ed., *The Impact of the Antibiotics on Medicine and Society*. New York : International Universities Press.
- Goldstone, JA (1991) "The Causes of Long Waves in Early Modern Economic history" in Joel Mokyr (ed) *The Vital One: Essays in Honor of Jonathan R.T. Hughes*. Greenwich, CT., JAI Press, pp. 51-92.
- Hodgson, G (1993), *Economics and Evolution*, Cambridge: Polity Press.
- Hull, DL (1988). *Science as a Progress*. Chicago: University of Chicago Press.
- Jones, S (1993). *The Language of Genes*, New York: Anchor Books.
- Kaufman, M "The American Anti-vaccinationist and their Arguments," *Bulletin of the History of Medicine*, Vol. 41, No. 5 (1967): 463-78.
- Kimura, M (1992). "Neutralism," in E Fox Keller and E Lloyd, eds., *Keywords in Evolutionary Biology*, Cambridge, MA: Harvard University Press.
- Kitcher, P (1993). *The Advancement of Science: Science without Legend, Objectivity without Illusions*. New York: Oxford University Press.
- Kuznets, S (1965). *Economic Growth and Structure*, New York: W.W. Norton.
- Mokyr J (1985). "Demand vs. Supply in the Industrial Revolution." In Joel Mokyr (ed) *The Economics of the Industrial Revolution*. Totowa, N.J.: Rowman and Allanheld (orig. published, 1976).
- Mokyr J (1990). *The Lever of Riches*, New York, Oxford University Press.

- Mokyr, J (1994). "Progress and Inertia in Technological Change" in John James and Mark Thomas, eds., *Capitalism in Context: Essays in honor of R.M. Hartwell*, Chicago: University of Chicago Press.
- Mokyr, J (1996). "Evolution and Technological Change: a new Metaphor for Economic History?" in Robert Fox, ed., *Technological Change*. London: Harwood publishers.
- Mokyr, J (1997) "Innovation and Selection in Evolutionary Models of Technology: Some Definitional Issues". Presented to the Conference on Evolutionary Models in Technology, Wallingford, England, Jan. 9-12, revised June 1997.
- Mokyr J and Stein R (1996). "Science, Health and Household Technology: the Effect of the Pasteur Revolution on Consumer Demand," "Science, Health and Household Technology: the Effect of the Pasteur Revolution on Consumer Demand," in RJ Gordon and T Bresnahan, (Eds.) *The Economics of New Goods*, Chicago: University of Chicago Press and NBER.
- Mowery, D and Rosenberg, N (1979). "The Influence of Market Demand upon Innovation." *Research Policy* 8: 102-53.
- Porter, R (1995). "The Eighteenth century" in Lawrence Konrad et al (Eds.) *The Western Medical Tradition, 800 BC to AD 1800*. Cambridge: Cambridge University Press.
- Richards, RJ (1987). *Darwin and the Emergence of Evolutionary Theories of Mind and Behavior*. Chicago: The University of Chicago Press.
- Rosen, G (1993) *A History of Public Health*, new ed., Baltimore: Johns Hopkins University Press.
- Rosenberg, N. (1976). "Science, Invention and Economic growth" in *Perspectives on Technology*, Cambridge: Cambridge University Press, orig. pub. in 1974.
- Sale, K (1995). *Rebels Against the Future: the Luddites and their War on the Industrial Revolution*. Reading, MA., Addison Wesley.
- Sobel, D (1995). *Longitude*. Harmondsworth: Penguin Books.
- Stebbins, GL (1982). *Darwin to DNA, Molecules to Humanity*. San Francisco: W.H. Freeman.
- Vincenti, W (1990). *What Engineers Know and How They Know It*. Baltimore: Johns Hopkins Press.