



Laboratory of Economics and Management
Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)
Tel. +39-050-883-343 Fax +39-050-883-344
Email: lem@sssup.it Web Page: <http://www.sssup.it/~LEM/>

LEM

Working Paper Series

ON THE UNEVEN EVOLUTION OF HUMAN KNOW-HOW

*Richard R. Nelson**

** Columbia University, New York, US*

2003/25

November 2003

ISSN (online) 2284-0400

ON THE UNEVEN EVOLUTION OF HUMAN KNOW-HOW¹

Richard R. Nelson

Columbia University

6/22/02

I. INTRODUCTION

Economists long have recognized the advance of human know-how as the central driving force behind the remarkable increases in living standards that have been achieved over the past two centuries. Contemporary scholars tend to attribute this understanding to the work of Solow (1957) and other economists working with the new National Product account data in the years after World War II. However, these studies are best regarded as providing supporting quantitative evidence for something that sophisticated economists have known for a long time. Thus, writing during the first industrial revolution, Adam Smith (1776) clearly understood and highlighted the key role of technological advance in lifting *The Wealth of Nations*; recall the pin-making example with which he begins his great book.

I am using the term “know-how” here to denote the wide range of techniques and understandings human societies have acquired over the years that enable them to meet their wants. In my use of the term, it encompasses “technology”, but includes more than that latter

¹The research behind this paper was supported by the Mellon Foundation. An earlier version of this essay was published in the OECD volume Knowledge Management and the Learning Economy (see Nelson, 2000).

term customarily is assumed to cover. One of the principal purposes of this essay is to analyze the nature of human know-how, and how it has been acquired.

While the remarkable advances in know-how are widely recognized, less attention has been paid to the fact that the advance of effective know-how has been extremely uneven across different economic sectors and classes of human needs. Some areas of human know-how today are extraordinarily powerful; consider modern information and computation technologies, or certain fields of modern medicine. On the other hand, certain human illnesses have defied continuing efforts to deal with them better. Breast cancers remain a major scourge. And many broad areas of human activity have seen little progress in know-how. It is not clear that our ability to educate children has advanced much over the last century. Despite a lot of huffing at Business Schools and in books on management, there does not seem to have been much improvement over the years in management know-how. Why?

I offer here a very preliminary exploration of this important puzzle, which will be divided into three parts. First, as I suggested above, it seems important to try to get a grip on the nature of modern human know-how. What are its aspects, and how is it organized? Where is it "located" and how is it applied? I shall argue that human-know is multifaceted and variegated, and stored in different places and forms. Some of it is of the form often thought of as engineering product or process design, relatively well articulated "how it is done" knowledge. However, much is embodied in particular human skills, as contrasted with "blue print like" know-how. Some involves sophisticated understanding of why practice works; some simply understanding from experience that a practice does work. And an important part

of know-how is knowing how to tap into, and coordinate, the various capabilities and efforts that need to be brought together to do a job.

Second, there is the basic question of how humans achieved the tremendously broad and effective body of know-how that we have achieved. I (in accord with many other scholars of technological advance) will propose that cumulative advance of know-how must be understood as a process of "cultural" learning or evolution. That cultural evolutionary process, in turn, involves the coevolution of technique and understanding. In recent times a good part of that understanding has been associated with a field of science or an engineering discipline.

Third, once one recognizes the extremely unbalanced nature of what we have achieved, it is apparent that our cultural learning or evolution system works much better in certain arenas than in others. In section IV I explore the factors that might explain this. Section V is concerned with education as a special case. In the concluding section I reflect on some of the consequences, if I am correct about the key reasons why certain areas of know-how are very difficult to advance.

II. THE NATURE OF HUMAN KNOW-HOW

Many of the important characteristics of human know-how, characteristics that are important to have in mind in reflecting on how know-how advances, and what makes the advance of know-how difficult in certain areas, can be brought into view by considering a particular example of modern advanced know-how: the performance of a surgery on a human

heart.² In the first place, it is important to recognize the variety of particular skills

²I am indebted to Dr. Annetine Gelijns and Dr. Alan Moskowitz for their checking of what I say about heart surgery. I also have consulted the splendid paper on an aspect of heart surgery by Edmondson, Bohmer, and Pisano (2001).

involved, and that effective performance is a group achievement.

Thus the surgeon, who generally is thought of as the key actor, has command over a certain body of practice. So does the anesthesiologist. To a considerable extent these bodies of practice are different. On the other hand, each actor knows "about" the skills of the other. Also, in the performance of an operation there will be a number of assistants involved who have command over certain skills. Some, but not all, of what they do could be done by the surgeon or the anesthesiologist, but it is far less costly to delegate relatively simple tasks to less highly trained and paid people. In general the surgeon serves as orchestra conductor, as well as key player in the operation. However, all the players know at least the broad outlines of the overall operation, and the details of their own roles in it. In general a successful operation requires that all of the roles be performed effectively, and in effective tune with each other.

In the case of heart surgery, like in most modern technologies, much of the technique is embodied in specialized apparatus, substances, and other artifacts. The anesthesiologist works with various substances that have been found to be effective, with pieces of apparatus that deliver those substances, and with a variety of dials and other measuring instruments that enable him or her to monitor what is going on. And the surgeon, of course, also works with a complex of materials and instruments. The embodiment of key aspects of the techniques involved in specialized artifacts should be understood as an extension of the team nature of know-how. Clearly much of that know-how is "upstream" from the locus of immediate action.

Another central characteristic of effective know-how is that it involves both a body of practice or technique, and a body of understanding. Behind the surgeon's command of skilled practice, and the anesthesiologist's, lies a broader body of understanding involving the human

body, of what is involved in the procedures being employed and the conditions of success and failure, and of the various substances and instruments being used. When things are going routinely, that broader body of understanding never may be invoked consciously. But it may play a very important role in holding skilled performance in place, being invoked unconsciously to prevent deviations that could undermine effectiveness or court trouble. And from time to time, in particular when something is seen or occurs that is not quite what is expected, conscious thinking tapping that body of understanding may be essential to effective performance.

Ever since Polanyi (1958) pointed it out, scholars have recognized that some of human know-how is "articulated," in the sense that it can be described and communicated in some form of language, or other symbolic system, while other aspects are "tacit". Thus a good portion of the specialized know-how of the surgeon and the anesthesiologist can control the work of their fingers, but may not be easily explainable in words or other symbols to others, even to other physicians, who however perhaps can learn by watching and trying to imitate. But other parts of their relevant know-how can be expressed in a way that can be understood, at least by other professionals with the same background of tacit knowledge.

These articulated parts of know-how often are written up in texts and treatises. Studying these may be an essential, if not sufficient, part of the way that pre-meds become doctors. And experienced doctors will go to the journals, or the Internet, to find out what is new, and sometimes to refresh their own knowledge. Like extant equipment and materials, texts and libraries provide storage for know-how outside of individual human minds.

While it seems natural to associate "tacit" with the practice or technique aspect of know-how, and articulated with the understanding aspect, I do not think the mapping is all that neat. Although it is clear that much of "technique" is tacit, a cake recipe, or a blueprint, is all "technique", but to a considerable extent is laid out and articulated on paper. Also, a considerable amount of technique is embodied in the artifacts used, and while the anesthesiologist may not be able to explain just how his apparatus works, he almost certainly can identify it by name and explain its use in a way that would enable another doctor to obtain and use it. On the other hand, the surgeon may see and understand that something is not going quite right with the operation, and not be able to explain in words just what he or she sees, or why that seems to signal trouble.

But language, and the ability to lay out know-how in language, clearly is very important in making know-how broadly available--an element of culture, as it were. The know-how of the surgeon and the anesthesiologist is cultural in the sense that much of what they know also is known by other surgeons and physicians, who have gone through similar training programs, use the same equipment, read the same journals, attend the same conferences. There are various mechanisms that facilitate, or even force, sharing of information among anesthesiologists. I do not mean to play down here the tacit aspects of learned skills, which may lie behind very great differences in effective performance, or the efforts of some professionals to keep certain aspects of their technique and understanding privy. But a striking aspect of most broadly important bodies of technique and understanding is that they are broadly shared.

On the other hand, it is clear that the overall know-how needed to perform complex tasks often is very divided. I have highlighted the separate bodies of practice and understanding possessed by the surgeon, and the anesthesiologist. In turn the anesthesiologist may know how to make his equipment work, but little about how to produce or design that equipment. People at the company that sold the machine may know those things, but no one at that company may know all of it. Reflect on whether anybody, or any small group, at Boeing Aircraft Company "knows how" to produce, or design, a modern aircraft, including the essential "details".

Because overall know-how is divided and widely distributed among different individuals and groups, to be effective know-how needs to be brought together and coordinated. For that reason, an extremely important part of know-how is knowledge of the elements that are needed, and of how to coordinate, and manage their combined operation. Much of the know-how possessed by the chief engineer at Boeing is of this sort.

In another paper, Sampat and I used the term "social" technologies to describe this latter kind of know-how, and differentiated social technologies from physical technologies, a term we used to denote what engineers generally mean by technology (Nelson and Sampat, 2001). Under the standard conception, physical technologies are recipe or blueprint-like, characterizing what is to be done, including designation of the particular operations (which may require highly developed skills) and (in some cases quite specialized) materials involved, but do not speak to how the work is to be divided and coordinated. In contrast, what I call social technologies are associated with effective structures of division of labor, and procedures for task coordination, and management.

As with practice and understanding, and tacit and articulated know-how, the physical and social aspects of technologies often are intimately intertwined. Consider the famous Ford mass-production line for Model T cars, or the Toyota method of "lean manufacture." These involve both a set of sequenced physical actions taken by the parties to the process, and a division of labor and a coordinating mechanism so that the actions taken by the particular parties ultimately add up to a finished automobile. Or reflect on the heart surgery example that I gave at the start of this section. Again, one sees a complex mix of physical technologies involved, employed by a team in which each member must do assigned tasks in harmony with what others are doing.

I propose that the human know-how involved in getting complex things done generally involves this mixture of understanding and practice, of articulated and tacit knowledge, of physical and social technologies, that I have described in the particular case of heart surgery. The analysis of how automobiles are produced by Womack et. al. (2000) involves a similar mix of ingredients. Hutchins (1996) describes what is involved in navigating a ship in much the same way that I have described heart surgery, and uses that example as a vehicle for illuminating collective "cognition". Bucciarelli (1994) has arrived at a similar conception in his analysis of what it means to know "how your telephone works".

These kinds of know-how systems have been brought into place, and develop further, through the cumulative actions of many individuals and organizations who have particular objectives in mind. However, the overall system cannot be regarded meaningfully as having been planned. Rather, our know-how systems need to be understood as having evolved, in a sense I now will elaborate.

III. THE COEVOLUTION OF TECHNIQUE AND UNDERSTANDING

Scholars of technological advance, from a wide variety of disciplines, have converged on the proposition that technological advance proceeds through an evolutionary process. (See for example Constant, 1980, Nelson and Winter, 1982, Basalla, 1988, Dosi, 1988, Vincenti, 1990, Mokyr, 1990, Petroski, 1992, Saviotti, 1991, Metcalfe, 1998, Ziman, 2000, Nelson and Nelson, 2002.) The process is evolutionary in the sense that at any time there generally are a wide variety of efforts going on to advance of technology, which to some extent are in competition with each other, as well as with prevailing practice. The winners and losers in this competition are determined through an ex-post selection process.

However, the proposition that technology evolves in the above sense in no way denies, or plays down, the often extremely powerful body of understanding and technique used to guide their efforts, at least in modern times, by those who seek to advance it. Thus Vincenti's discussion (1990) of what is involved in inventing, and problem solving, in aircraft design stresses the scientific knowledge, and the professional experience and technique, that is employed. He provides an extensive catalog of the kinds of complex knowledge that modern aeronautical engineers possess, and discusses in detail how this body of knowledge guides their efforts at design. Included centrally in his story is the body of testing technique and apparatus that designers can use that provides them with information regarding how proposed configurations and problem solutions likely will work in practice, which can be employed without going to an actual full scale test.

However Vincenti, along with other scholars who propose that technological advance is an evolutionary process, argues (and provides the documentation for the argument) that efforts at inventing and technological problem-solving almost always reach beyond the range of options that are perfectly understood, or which can be reliably tested, short of full scale operation in the actual practical environment, and in that sense are somewhat "blind." Therefore, to a considerable extent what works and what does not, and what works better than what, must be learned through actual experience and actual competition.

Most scholars of technological advance also are united by their insistence that the process needs to be understood as "cultural" in the sense, first, that much of the background knowledge needed for inventing is widely held, and second, that anyone attempting to advance a technology almost always is standing on "the shoulders of giants", or more accurately, on the top of a large body of already achieved technique and understanding that has been developed by a large number of predecessors.

David's discussion (1991) of the large number of diverse technological advances, made by different individuals and organizations, that were required to take advantage of the potentialities opened up by the earlier invention of technologies for the generation and distribution of electric power, provides a vivid demonstration of the cumulative and collective nature of technological advance. Rosenberg's historical analysis (1996) of the development of the multiple uses of the laser is another splendid example of the point.

Earlier I argued that technology should be recognized as involving both a body of technique or practice, and a body of understanding or knowledge. In the process of technological advance, both evolve. Or, I would propose more specifically that technique and

understanding coevolve. The development of a particular new product or process generally brings with it a wider body of new understanding that includes, but transcends, the particulars of the new technique. A new understanding, earned through this route, or through efforts more directly aimed to advance understanding, in turn provides clues and opportunities for the further advance of technique.

Since the days of Francis Bacon, the drive to advance technological practice and the pursuit of understanding bearing on that technology have gone hand in hand. Over the last century the linkage has been institutionalized in the development of fields of applied science and engineering disciplines expressly dedicated to providing the understandings useful for advancing practice in a field of technology or an industry. The latter activity is largely the domain of business firms, or other organizations or individuals, who will actually use or sell the products or processes. The advance of understanding, however, is largely the mission of universities and public laboratories (see Rosenberg and Nelson, 1994).

Indeed, at the present time, the lion's share of research going on in American universities is in fields with names like "material science," computer science, electrical engineering, pathology, etc. In today's world science is useful to inventing not so much because of serendipity, but because many fields of modern science are designed to help clear the path for technological progress. In a recent survey (Klevorick et al, 1995), industrial R&D executives were asked to identify the fields of academic science that most contributed to the successes of R&D, and they strongly tended to list fields of the sort mentioned above, as contrasted with, say physics or mathematics.

For me at least, a striking characteristic of fields where technological advance has been rapid is that they all seem to be closely connected to a powerful applied science or engineering discipline (see Rosenberg, 1974, Klevorick et al, 1995, Nelson and Wolff, 1997, Rosenberg, 2001). These bodies of scientific knowledge serve, first, to enlarge and extend the area beyond existing practice that an inventor or problem solver can see relatively clearly, and hence go into without being completely "blind". That is, strong science provides guidance regarding what particular paths are likely to lead to solutions or improvements, and which are likely to be dead ends. In technologies illuminated by strong science, an inventor often can see a good distance beyond current best practice.

Second, the sciences and the engineering disciplines provide powerful ways of experimenting and testing new departures, so that a person who commands these can see relatively quickly and cheaply if they work, or are promising, or problematic. Thus pilot plants play a key role in efforts to develop new chemical process technology. Wind tunnels used to play a similar role in aircraft design. Where scientific and engineering knowledge is strong, these days one can explore and test by building computer models. More generally, strong scientific knowledge not only enables inventors to see promising paths, but also to reliably assess the promise of the path in a timely fashion, and without having to build and test a full scale version in the actual operating environment.

I note, I stress, that these advantages lent by a strong body of understanding do not diminish the importance of learning by doing and using in the advance of a technology. As Vincenti has argued, in the end whether a new design or process is satisfactory, or better than what it aims to replace, can only be determined in on-line experience. I shall argue in the next

section that the capabilities to recognize, generate, evaluate, and duplicate on-line variation is absolutely essential. If these capabilities are strong, cumulative technological advance can proceed even if the body of understanding, the underlying science, is weak.

However, for reasons I have put forth above, a strong science base greatly augments the power and efficiency of efforts to advance a technology. I also note, propose, that when there is a strong body of underlying scientific knowledge, a good share of the work of advancing a technology tends to go on “off-line”, in facilities like industrial R and D laboratories. The power of the underlying sciences means that people who have mastered that body of specialized knowledge are needed to do effective R and D. In general the skills here are very different from those who work “on-line”. And the activities involved in doing R and D tend to be different from those involved in on-line experimentation. The work of advancing the technology thus tends to be specialized both in terms of what is done, and in terms of the personal involved. A considerable degree of such specialization is a hallmark of modern industries where technological advance is rapid.

However, as Vincenti has argued, there is no escaping the need for on-line evaluation, and tinkering. In general, in fields where technological advance is rapid there is an interactive mix between learning by doing and using, and off line R and D. I shall argue in the next section that ability to experiment, and learn from experiments, is key to both aspects of the process.

IV. WHY HAS ACHIEVEMENT BEEN SO UNBALANCED?: SOME SPECULATIONS

I want to focus now on the puzzle of why the advance of human know-how has been so uneven, spectacular in areas like information and communications, and in dealing with certain kinds of human illness, but very limited in other areas, for example education, or rehabilitation of criminals.

One obvious reason why know-how has advanced so much more rapidly in some fields than in others is that more resources have been applied to the effort. Business firms have seen certain kinds of advances as being profitable, but not others. Governments have been willing to put public funds into R and D on certain classes of problems, but there has been little effective political support for public R and D moneys in other areas. Thus if one considers human illness, a major reason why little progress has been made on certain tropical diseases is that drug companies do not see the market in poor tropical countries as promising much profit, and publicly funded efforts have been limited.

But while “demand side” limitations clearly have been important in some cases where the advance of know-how has been very small, as Mowery and Rosenberg (1979) have pointed out, many important human wants remain unmet, even though significant profit could be earned by a person or firm that figured out how to remove the roadblocks to meeting those wants more effectively. There clearly are major differences across sectors and areas of human activity in the ability of society to advance effective know-how. Within medicine, cures have been found for Hodgkin’s lymphoma and testicular cancers, but not for prostate and breast cancers. These differences are puzzling and disturbing, and will be my focus in the remainder of this essay.

At one level, signaled above, my basic argument is that the key factor is the strength of the understanding bearing on practice in a field. In an earlier paper (Nelson and Wolff, 1997) evidence was provided that the rate of technological advance in an industry is strongly correlated with the strength and vigor of the sciences on which R and D in that industry draws. However, this explanation of course only pushes the question back a level. Why are the sciences that underlie certain technologies so much more powerful than others?

As I have hinted earlier, my tentative explanation of the puzzle at this deeper level involves looking at the causal arrow between strength of understanding and ability to experiment fruitfully with a technology the other way around. I want to propose that the ability to conceive and carry out well defined experimental probes of possible ways to improve technological performance, and to get sharp and reliable feedback on the results, contributes importantly to the human ability to develop an applied science that effectively illuminates that technology.

Of course I recognize that some technologies in effect are born out of prior scientific discovery which was the result of research that was not particularly oriented towards making new technology possible. The rise of radio technology is a good example. But after a new technology emerges, it begins to pose particular scientific problems and puzzles. Rosenberg (1982) has argued that a significant portion of the puzzles that science addresses have been revealed or created by the operation of technologies. In turn, the further advance of a technology depends to a considerable extent on how effectively science is marshaled to illuminate the roadblocks to progress.

I have been proposing that the successful development of an applied science or field of engineering research often is the key to rapid and continuing advance of know-how in a field of activity. Electrical and chemical engineering are fields of research as well as teaching that came into existence as the industries using the technologies on which they are focussed grew in importance. The invention and development of the transistor and integrated circuits provided strong intellectual stimulation (and a reason for financial support) for the new field of material science.

These new technology oriented scientific and engineering fields rapidly enriched and improved their theoretical bases. But from the beginnings they have been very experiment oriented. And much of the experimenting has involved aspects of the technologies that provide the reason for the field's support. In turn, advances in the technologies have provided puzzles and challenges for the sciences. Rosenberg's discussion (2001) of the nature of engineering research and knowledge and its relationship to the advance of practice is particularly apt.

When progress is rapid, there seems to be a strong symbiosis between the particular structure of the technologies and the focus of the sciences underlying them. On the one hand, the technology itself tends to move towards where the understanding is strong. On the other hand, with technology linked to science, the science is able to progress by manipulating aspects of the technology experimentally.

Do I overstate the role of experimentation in the development of science? I do recognize that astronomy, now cosmology, is not strictly an experimental science. However, given its intellectual base in physics, it has been possible to both draw on and focus experimental physical research which probes at the fundamental theoretical conceptions of

astronomy and cosmology. And the ability to make precise empirical observations of the sort needed to rigorously test evolving cosmological theory has enabled that science to proceed almost as if it were experimental. In some cases non experimental data can provide the basis for a strong science. But most of the strong fields of empirical science that have been developed have involved experimentation in an essential way. And I believe that this is especially the case with sciences that illuminate technologies. Those sciences cannot progress effectively, at least not in a way that is useful to advancing the technology, unless the technology itself is suitable for experimentation.

Above I noted that, in fields where technological progress has been rapid, problem solving and inventing is done to a considerable degree off-line, in specialized facilities separated from where the technology actually is being employed. While many of the problems and opportunities are recognized on-line, much of the problem solving is done off-line. For this specialization and separation to work effectively, it must be possible to isolate the technology from much of its operating surroundings, and to work with it in a controlled environment. And performance in that controlled environment must provide reliable information about likely performance “on-line”. For this to be so, it almost always is necessary that the “design” that has been developed and tested in a controlled environment off-line be robust to or protectable from different factors that can vary in actual practice, and which cannot be controlled.

This latter requirement also is important if the variations being explored are to be replicable. Replicability of course is essential if what is learned or created off-line in R and D is to be usable in practice, or at least transferable to an on-line setting so that its efficacy can be evaluated. In many fields of technology one sees progress being achieved through an iterative

process, with the locus of analysis going back and forth between the lab and actual practice. But replicability also is needed so that over the long run many parties can be involved in efforts to advance the technology, building on each others' work, a condition I argued earlier seems to be essential if progress is to be cumulative.

This latter argument would be valid even if experimentation were nearly completely blind, and off line R and D had little power. However, I have been arguing that not only are these characteristics conducive, probably necessary, if a technology is to be advanced cumulatively and rapidly through experimental trial and feedback. They may be necessary, and certainly are conducive, for a body of reliable scientific knowledge, in the sense of Ziman (1978) to grow up that supports efforts to advance know how in an area. For an applied science or engineering discipline to develop a powerful body of knowledge and technique that illuminates a body of practice and aids in its improvement, that body of practice must lend itself to rigorous study and experimentation, with a capability to evaluate reliably the results of variation. Vincenti's study (1990) of aeronautical engineering knowledge and its development provides strong evidence for this argument.

Am I underplaying the role of the basic sciences, like physics, mathematics, various areas of biology, whose orientation is not defined in terms of a particular technology or solving a set of practical problems? I do not want to underplay their role. I would propose, however, that advances in basic science mostly have their impact on technological advance by informing and strengthening the applied sciences and engineering disciplines that do have a practical focus. Thus I am proposing that fields of technology that advance rapidly and cumulatively have under them strong applied sciences that in turn are able to draw from strong basic science.

This is not an endorsement of the “linear model”. Rather it is a proposition about the structure of a knowledge systems that exist in areas where the advance of know-how is strong. I am calling attention to the critical role of what has been called the “bridge” sciences, and proposing that to be effective they need to be, at once, closely oriented to the technologies they are designed to illuminate, and close enough to the basic sciences so that they can draw power from them. A large gap, on either side, limits their effectiveness.

Consider some of the implications, if this argument is broadly correct. First, as advocates of support of science long have argued, it is a poor bet, and a likely waste of money, to pour resources into advancing practice in a field, if understanding there is weak. There is little then to guide efforts to develop technology that will perform significantly better than prevailing practice. And information as to whether or not the new departures are effective may be slow in coming and inconclusive. For this reason, a necessary first step to solving the practical problem or meeting the pressing need is to support the scientific research that enables the problem to be understood. This argument of course is an old one, and often made in a self serving way by scientists.

But second, my argument points to the major difficulties that may need to be overcome, and the long time period that may be required, for a strategy of trying to develop a useful underlying science to be successful. The scientific understanding, to be useful, must link up with the available technologies for operating in the area, or point relatively clearly to practical new ones. Understanding far removed from possible practice does not provide sharp guidance as to how practice can be improved. On the other hand, an attempt to build an applied science that is far removed from strong fields of basic science may yield knowledge of

limited power. One implication of this is that the achievement of a science that illuminates a technology may depend on transforming the technology so that it becomes more amenable to scientific inquiry. As I shall argue shortly, there may be strong constraints that make this difficult.

V. THE CASE OF EDUCATION

Consider a highly relevant case that illustrates, I believe, several of the points I have just made: the efforts to develop more effective school educational practice. (See Murnane and Nelson, 1984, Hagarty, 2000). I think it apparent that neither of the two attributes that I argued earlier made R and D in a field powerful are strong in the case of education. It is very difficult in education to predict with any precision just how a proposed change in teaching method actually will work out in practice. General understanding of the education process and schooling may provide a broad prediction, but the devil is in the details. And it is difficult, perhaps impossible, to get reliable information on this from simple inexpensive pilot experiments.

These limitations are closely related, I would argue, to the following problem. The fact that a particular practice seems to work well in a particular context does not mean that it can easily be transferred to another context, or if this is tried, that it will work well there. Partly the problem is that it is difficult to specify in any detail, or to know, the essential aspects that determine its performance; thus replication is chancy. Another problem is that what works well in one context may not work so well in another, and it is hard to control for the relevant variables. Still a third problem, related to the above but of central importance in its own right, is that evaluation is extremely difficult. It may take many years before the lasting effects of a new

mode of instruction can be learned. And there may be many different kinds of impact to be considered.

The difficulties here clearly reside in the education process itself. Education as currently practiced largely involves a set of strategies and practices that are generally understood as appropriate in particular contexts, but with a lot of variation across individual classrooms and teachers. There are indeed canons of good practice. But not many educationists are ready to propose that there are a set of foolproof “cake recipes” that define best practice in teaching. And while novice teachers may learn a lot from observing able experienced teachers, every teacher has their own particular strengths and weaknesses, and style of operating.

A certain amount of classroom equipment is used: textbooks, perhaps film, recently at least computers. But, while some students of education hold out hope for the Internet and the computer, at present there are no powerful devices used in education, comparable to the apparatus that dispenses and monitors anesthesia, that are used in heart surgery. Some years ago Cuban (1986) reviewed experience with using computers and other forms of teaching equipment in classrooms, and concluded that their impact had not been dramatic. In a recent paper, Murnane, Sharkey, and Levy (2001) review a particular educational program that centrally involves use of the Internet, and also a considerable amount of programmed instruction, and which has been implemented in a number of schools. They highlight the apparent broad effectiveness of the program, but also argue that the standardized instruction package and the use of the Internet should be understood as a complement not a substitute for an effective teacher working with students.

It is well known that how an individual child learns in a classroom is strongly affected by the behavior and attitudes of other children in that classroom, and is not independent of what is going on in a child's life outside of school. A major portion of the challenge for a teacher is to organize and manage classroom interaction, as well as to deal with the particular problems or challenges of individual students.

In section II I proposed that all bodies of human know-how bearing on complex activities, like a heart operation, or designing and building an aircraft, or education, involve a mix of articulated and tacit knowledge, and physical and social technologies. It is apparent that the mix in education is heavily weighted towards the tacit and social.

These characteristics are reflected in the limited ability to conduct educational experiments, the results of which provide reliable guides to how to improve educational practice in real world settings. For many years such experimentation has been high on the agenda of scientifically oriented Schools of Education. But consistently the record has been that what is reported to work in a lab school or in another chosen testing locus has been hard to duplicate outside of the locus of the original research. As noted, part of the problem clearly has been that it is impossible to describe what the experimental treatment was with sufficient precision and detail so that one could know whether one was replicating the key elements of it or not. Part is that the context conditions that enabled a particular treatment to work were not fully known, and not necessarily in existence in other places. And part surely is that evaluation takes time and in many cases does not yield unambiguous results.

These basic characteristics of education also limit what can be learned from large scale statistical studies that collect and analyze data from a number of different schools or classes or

modes of teaching. It is not that statistical studies do not identify important correlates of good educational performance. One important correlate is the education and income of a student's parents. Another is the training and experience of a student's teacher. But the former provides no information as to how to improve the performance of schools, given the backgrounds of the students. And while the latter does provide guidance to schools regarding the kind of teachers they ought to hire and about the importance of encouraging promising teachers to stay in the system, it tells very little directly about the educational practices that work best.

The fields of research that one would hope would illuminate the educational process and guide efforts at improvement in fact provide only a dim light. On the one hand, research that is focussed on subject matter that arguably is closely related to the education process at best seems to yield course grained and often unreliable conclusions. On the other hand, scientific research that limits itself to subject matter where relatively fine grained and reliable knowledge can be attained, tends to generate findings that are a far distance from anything useful in the education process.

Thus a recent (U.S.) National Research Council report, How People Learn: Bridging Research and Practice (1999), gives the following as an example of the former kind of research finding, and how such knowledge is useful in education:

“Students come to the classroom with preconceptions about how the world works. If their initial understanding is not engaged, they may fail to grasp the new concepts and information that are taught, or they may learn them for purposes of a test but revert to their preconceptions outside the classroom” (p.10)

The contrast of this bit of knowledge, useful as it is, with say the discovery that scurvy among seamen was caused by the absence of a class of foodstuffs in their diet, is striking. The latter led relatively directly to dealing with the problem by assuring the availability of certain of those foodstuffs. The understanding highlighted in the NRC report points, but only very broadly, to good teaching practice.

I note that the original discovery about scurvy was not associated with a theory to explain it. Theoretical understanding was achieved only much later, with the discovery of vitamins and their association with body function. But the initial finding regarding the causes of scurvy was sharp and precise enough to identify a treatment that worked. The NRC publication, from which the above quote is drawn, seems blind to the difference between the nature of the findings of educational research, and the knowledge base under medical practice, or at least is mute about it.

The same NRC reports mentions that more fundamental research has been going on in brain science and cognitive science. Considerable progress has been made towards understanding areas and mechanisms in the brain associated with various kinds of perception, and thinking. However, the detailed hard findings at this level are many layers away from providing useful input to guiding teaching. The NRC report acknowledges this, and points to the intellectual gap as a real problem:

“The concern of researchers for the validity and robustness of their work...often differ from the focus of educators on the applicability of these constructs in real classroom settings..”
(P.6).

But the report does not draw the obvious conclusion that the fine grained and reliable knowledge coming out of fields like brain science are that way because the subject of research is carefully controlled and far removed from the hurly burly of the educational process. Again, the contrast with medical care, where biological understanding often is very close to what one needs to know to cope with a disease, is striking.

Since both education and medical care are activities focussed on helping individuals, and the recipient of the treatment is a vital element of the process of teaching or healing, I believe the contrast here is well worth exploring further. Most of the significant advances in medical care have occurred over the past one hundred and fifty years, and have been associated with a tremendous increase in scientific understanding of human illness of various kinds, and of the effects of various treatments. The basic mechanisms in question are biological, and often the biological mechanisms can be understood in terms of the chemistry and (occasionally) physics involved, all strong fields of science. Animals in many cases provide convenient models of humans, in circumstances where in vitro chemistry does not illuminate what is going on.

In general the improvements in performance of medical care have occurred in areas where understanding has become strong, but this is not always the case. In many cases we have learned that certain treatments work (like limes for scurvy, and aspirin for headaches) but initially at least have had little understanding of just why. But we were able to learn that lime juice prevents scurvy, and aspirin relieves headaches and seems to reduce the risk of certain heart ailments, and make use of that knowledge in the practice of medicine, because limes and aspirin are well defined substances. Thus "swallowing lime juice" or "taking aspirin" are

routines that can be well enough described so that people instructed to do it can, with only a small chance of getting it badly wrong.

As these examples indicate, the medical treatments that we have learned work well have tended to be well specified; indeed most of them are substances or other artifacts (glasses) that we have learned (often scientifically) to characterize precisely. And by and large their effects are not greatly influenced by factors from which they cannot be shielded (but consider the warnings on medicines regarding what not to take at the same time). Thus we are able to control and calibrate the treatment, and are able to learn from variation, either accidental or deliberate.

And of particular importance for the current discussion, these characteristics, where they exist, permit both controlled experimentation regarding new medical practice--new drug regimes, surgical procedures, etc--and the development of a relatively strong body of biomedical scientific knowledge. While biomedical scientists have a tendency to underplay the importance of what is learned in "on line" actual practice, "off line" R and D, and controlled tests, play a very powerful role in facilitating the evolution of medical know-how. (For a careful balanced discussion rich with empirical examples, see Gelijns, 1991).

Some scholars deeply committed to research to advance educational practice have taken as an insult my argument that the findings of research in these fields simply do not have the power of the findings of biomedical research to illuminate and facilitate the improvement of practice. My argument has nothing to do with the quality of the researchers in the field of education, but rather with the innate limitations on the ability of research to contribute to the advancement of technologies that are largely tacit and social.

Earlier I put forth business management as another field where, like education, advance scarcely has been dramatic. I propose that the reasons are very similar.

There probably has been less “off-line” research aimed to develop better management practice than there has been off-line experimental research in the field of education. Most of the research in this area has proceeded by trying to identify firms or cases where a particular practice is or has been employed, and to compare performance in these instances with cases where the practice has not been employed. But as with the case of cross sectional studies of the efficacy of education practice, such efforts have been bedeviled by, on the one hand, great difficulty in pinning down the essentials of the practice being studied and hence being able to determine when it was actually employed, and second, being able to distinguish the effects on firm performance of use of the practice from the effects of other variables. These two basic problems are, of course, not unrelated. The various studies of the value of employing “quality circles” is a good illustration of these problems. (See Cole and Scott, 2000, and Nelson, Peterhansl, and Sampat, forthcoming.)

My mother discipline is economics. The science of economics has much the same weaknesses as the science of education, and of business management practice. and I would argue for the same basic reasons. The limitations of all three fields largely reflect, under my argument, that the basic human activities in these arenas are highly tacit and social, and difficult to specify with precision. In each of these fields the motivation for study is largely to enable policy to be more effective, and in these fields there is strong awareness that the prevailing science provides at best only general and hedged guidance to policy. In economics, as in education, there is strong faith that “if we only had better scientific understanding” we could

develop more effective and reliable policies. But if I am right, the fact that economics as a science provides only broad and uncertain guidance to policy is in good part the result of the fact that the objects of interest are impossible to define and measure with precision. The science of economics can be made precise only by shifting the study to an arena far simpler than that in which we really are interested. And this, many would argue, is exactly what has happened in much of economics. While the results may make for some nice economic theoretical arguments, they do little to illuminate real policy issues.

But to return to the medicine-education comparison, it is interesting to note that, where medical treatment can not be specified in terms of pills or other physical substances, or a clear cut procedure like splinting a bone break, or where the effects of treatment cannot be isolated from those of other variables and actions (as in treatment of obesity), or where understanding is weak and animal tests do not provide much information (as in study of the effects of environmental factors on the incidence of cancer) medical R and D does not demonstrate much power. Here the situation is not very different, it seems to me, than in education and business management.

V. SOCIAL TECHNOLOGIES AND THE EVOLUTION ON KNOW-HOW

Are technologies that are strongly social and tacit important exceptions that fall outside of the remarkable abilities human societies have developed to advance their practical know-how? The discussion above has been concerned only with education in any detail; the discussion of management and economics was at best cryptic. But the elements that seem to make progress difficult in these areas seem quite similar, and to hold as well for areas like the

prevention of crime, or teen age pregnancies, or managing the medical care system, or the Internet. Interestingly, the two last examples are of cases where the underlying physical technologies have become very powerful, but the social technologies needed to manage them are not very effective.

In a recent book, Kline (1995) argued that human behavior in a social context was intrinsically more complex than the operation of a physical machine or other artifact, according to the particular measure of complexity that he lays out. He proposed, persuasively in my view, that fields of science that deal with very complex subjects cannot be expected to come up with the precise laws and relationships that have come out of physics. Is the reason why the sciences underlying social technologies are relatively weak simply that these kinds of technologies are very complex?

This is one way of looking at it. However, I have put forth a particular set of arguments regarding just why these kinds of technologies are difficult to advance, that involves their tacit and social nature in an essential way. I want to stand by my argument that the heart of the problem is the difficulty in these technologies of doing precise and replicable experimentation, and gaining reliable and generalizable knowledge from variation.

This formulation among other things has the advantage of leading to the question of whether these characteristics are innate, or whether they can be modified. I am not alone in pointing to these characteristics as an important part of the problem in advancing education.

Indeed there has been a long standing argument between educators who have advocated bringing more tightly controlled and explicit routine to the education process, and those who have resisted this strongly saying that this hinders tailoring education to the particular

needs and capabilities of individual students and the characteristics of particular groups of students assembled in a class (see e.g. Murnane and Nelson, 1984). This debate has ranged from argument about whether or not there is one particular way that reading is best taught, to the appropriate use of computers in education. A common strand, however, is the pluses and minuses of developing and using standardized methods.

Recently several economists (see e.g. Arora and Gambardella, 1994, Dasgupta and David, 1994, Cowan and Foray, 1997) have argued that the extent to which a technique is tacit or articulated and codified depends to a good extent on the magnitude and skill of the efforts to codify it. While it is not plausible that even a major effort could fully codify the skills of an expert surgeon, or an effective teacher, surely there is something of a common core of good practice that, to some extent can be codified. There certainly are relatively programmed teaching methods, including those built into computers, that have had a certain amount of effectiveness. The question is how far this can be pushed without running into the problem raised by those skeptical of routinization. One size of shoe does not fit all feet. But are there a reasonable number of well defined shoe sizes that mostly will do the job? That turns out to be the case, mostly, with shoes. How about education?

Much of the tacitness of educational practice is bound up, I would argue, with the innately social aspects of teaching and learning. There needs to be effective interaction between teacher and student, and to a considerable extent that interaction is influenced by the larger group in a classroom. The problem with advancing social technologies is that there are strong constraints associated with the capabilities and wills and beliefs of the people who's actions somehow must be enlisted, coordinated, or managed. In turn, these individual and idiosyncratic

constraints make it difficult or impossible to standardize a technique, or even to describe what is being done with precision, and make reliable experimentation, or generalizable feedback from operating experience, very difficult as well. Perhaps the course to greater effectiveness is to get rid of these constraints, by substituting physical for social technologies.

Indeed in many arenas exactly this has been done. Taylorism routinized and made explicit the jobs that workers did in manufacturing technology, and machinery and later more general automation transformed much of what had been a social technology of management and control into physical technology. Once this was done, it was possible to experiment with new designs for machines and automated coordination mechanisms, and make real progress on the management and coordination problem. In turn, routinization and mechanization greatly facilitated the development of strong engineering knowledge.

To some extent computer programmed instruction does this in education. But it is highly uncertain how far mechanized instruction can be pushed. And there remains the nagging problem that in this society at least individual differences are valued, not seen as something to be strongly repressed.

Improving the way we educate children surely is an extraordinarily important goal. Research that will help to guide experimentation and evaluation is of top priority. But perhaps we need to recognize that advancing knowledge and practice here is innately more difficult than advancing know-how in many areas of medicine, or agriculture, or telecommunications. And it is not at all clear that the strategies and organizational structures that have worked well to advance know-how in areas where it has been possible to routinize practice, to make

knowledge of best practice well articulated to a considerable degree, and to control or mechanize the processes closely, are the ones that will work well in education.

To some extent the constraints here are of our own making, and we can relax them if we choose. We now use drugs to help control certain individual behaviors that are judged likely to be destructive to self and others, but thus far society has shown reluctance to heavily drug all individuals who are judged likely to commit crimes. Can we require that children deemed likely to be disruptive in class go on drugs? Are we willing to jail parents whose children skip school? We can if we wish control at least some of the variables that make it so difficult to routinize and standardize education. And that probably would make it easier to learn from educational experimentation. However, most of us don't want to go very far down this road.

In education, and in other areas, there clearly are limits on our willingness to routinize and mechanize for the sake of better control, and the ability to make faster progress. A Brave New World is not all that attractive.

My exploration of the factors behind the uneven evolution of human know-how is just starting. I bet that my conjecture about the central importance of ability to recognize, generate, evaluate, and duplicate on-line variation will hold up after wider and deeper study. I bet that my proposition about the importance of a strong underlying applied science or engineering discipline, for which the former condition is necessary (but not sufficient) will hold up. According to this theory, areas where a major portion of the know-how is social and tacit, and there are constraints on changing this, are innately difficult to advance.

But they are only a portion of such areas. Lerner's recent study (2001) of frustration in the attempts over many years to deal effectively with breast cancers shows continuing dispute about the efficacy of physical technologies, associated with inability to get sharp persuasive evidence regarding the efficacy of different treatments. I view the case here as strongly supporting my proposition that a necessary condition for making progress is the ability to learn by doing and get reliable feedback from on-line variation. I see the inability of the "science" to advance here sufficiently to enable the development of significantly better practice as being as much the result of inability of trial and error learning to generate knowledge of what works, as a reason for the continuing blindness. That a technology is largely tacit and social seems to doom it to slow progress, but these conditions certainly don't seem to be necessary for such frustration. I hope the reader will agree with me that there are a range of fascinating and important puzzles here. Come join in the exploration.

REFERENCES

- Arora, A., and Gambardella, A., "The Changing Technology of Technological Change: General and Abstract Knowledge and the Division of Innovative labor," Research Policy, 1994, pp. 523-532.
- Basalla, G., The Evolution of Technology, Cambridge University Press, Cambridge, 1988.
- Bucciarelli, L., Designing Engineers, MIT Press, Cambridge, 1994.
- Cole, R., and Scott, R., The Quality Movement and Organization Theory, Sage Publications, Thousand Oaks, 2000
- Constant. E. The Origins of the Turbojet Revolution, Johns Hopkins Press, Baltimore, 1980
- Cowan, R., and Foray, D., "The Economics of Codification and the Diffusion of Knowledge," Research Policy, 1997, pp. 595-622.
- Cuban, L., Teachers and Machines: The Classroom Use of Technology Since 1920, Teachers College Press, N.Y., 1986
- Dasgupta, P., and David, P., "Towards a New Economics of Science," Research Policy, 1994, pp. 487-522.
- David, P., "Computer and Dynamo: The Modern Productivity Paradox in a Not-Too-Distant Mirror," in Technology and Productivity: The Challenge for Economic Policies, OECD, Paris, 1991.
- Dosi, G. "Sources, Procedures, and Microeconomic Effects of Innovation", Journal of Economic Literature, 1988, pp.1120-1171
- Edmondson, A., Bohmer, R., and Pisano, G., "Disrupted Routines: Team Learning and Technology Implementation in Hospitals", Working paper 00-003, Harvard Business School, 2001

Gelijns, A., Innovation in Clinical Practice: The Dynamics of Medical Technology Development, National Academy Press, Washington D.C., 1991

Hagarty, S., "Characterizing the Knowledge Base in Education," in Knowledge Management in the Learning Society, OECD, Paris, 2000

Hutchins, E., Cognition in the Wild, MIT Press, Cambridge, 1996

Klevorick, A., Levin, R., Nelson, R., and Winter, S., "On the Sources and Significance of Interindustry Differences in Technological Opportunities," Research Policy, Spring 1995.

Kline, S., Conceptual Foundations for Multidisciplinary Thinking, Stanford Un. Press, Stanford, 1995

Lerner, B., The Breast Cancer Wars, Oxford University Press, New York, 2001

Metcalf, S. Evolutionary Economics and Creative Destruction, Routledge, N.Y., 1998

Mowery, D., and Rosenberg, N. "The Influence of Market Demand on Innovation: A Critical Review of Some Recent Studies", Research Policy, 1979

Murnane, R., and Nelson, R., "Production and Innovation When Techniques Are Tacit: The Case of Education," Journal of Economic Behavior and Organization, 1984, pp. 353-373.

Murnane R., Sharkey N., and Levy F., "A Role for the Internet in American Education? Lessons from the CISCO Networking Academies", Manuscript, July 2001

Mokyr, Joel, The Lever of Riches, Oxford University Press, Oxford, 1990

National Research Council, How People Learn: Bridging Research and Practice, National Academies Press, Washington, 1999

Nelson, R., Peterhansl A., and Sampat, B., "Four Models", Manuscript, 2002

Nelson, K., and Nelson, R., "On the Nature and Evolution of Human Know-How," Research Policy, forthcoming

- Nelson, R., "Knowledge and Innovation Systems," in Knowledge Management in the Learning Society, OECD, Paris, 2000
- Nelson, R., and Sampat, B., "Making Sense of Institutions as a Factor in Economic Growth," Journal of Economic Organization and Behavior, 2001
- Nelson, R., and Winter, S., An Evolutionary Theory of Economic Growth, Harvard University Press, Cambridge, 1982.
- Nelson, R., and Wolff, E. "Factors Behind Cross-industry Differences in Technical Progress", Structural Change and Economic Dynamics, 1997, pp 205-220
- Petroski, H. The Evolution of Useful Things, Alfred Knopf, N.Y., 1992
- Polanyi, M., Personal Knowledge: Towards a Post-Critical Philosophy, Routledge and Kegan Paul, London, 1958.
- Rosenberg, N., "Science, Innovation, and Economic Growth" The Economic Journal, 1974
- Rosenberg, N., "How Exogenous is Science?", in Rosenberg, N., Inside the Black Box, Cambridge Un. Press, Cambridge, 1982
- Rosenberg, N., "Uncertainty and Technological Change", in Landau, R, et al, The Mosaic of Economic Growth, Stanford Un. Press, Stanford, 1996
- Rosenberg, N. "Engineering Knowledge", manuscript, Sept. 2001
- Rosenberg, N., and Nelson, R. "American Universities and Technological Advance in Industry", Research Policy, 1994
- Saviotti, P., Technological Evolution, Variety, and the Economy, Elgar, 1991
- Smith, A., The Wealth of Nations, Modern Library, N.Y., 1937, first published 1776

Solow, R., "Technical Change and the Aggregate Production Function", Review of Economics and Statistics, 1957

Vincenti, W., What Engineers Know and How They Know It, Johns Hopkins Press, Baltimore, 1990

Womack, J., Jones, D., and Roos, D., The Machine That Changed the World, MIT Press, Cambridge, 1990

Ziman, J., Reliable Knowledge, Cambridge Un. Press, Cambridge, 1978

Ziman, J., Technological Innovation as an Evolutionary Process, Cambridge Un. Press, Cambridge, 2000