

On the Uneven Evolution of Human Know-how

June 7, 2006
Richard R. Nelson
Columbia University

0I. Introduction

The last two centuries have seen enormous improvements in the standard of living of much of the world's population. Scholars from many different disciplines, not just economists, are in accord that the central driving force behind this progress has been the advance of technology, or more broadly the know-how that human society is able to marshal as it goes about various activities involved in the production of goods and services.. The advances in know-how that have been achieved are partly measured in the statistics on national product and income per capita, but these numbers only partially reflect the capabilities we now have to do things that were impossible in an earlier era, like to fly across the Atlantic ocean in seven hours, or to prevent or cure many infectious diseases.

Less attention has been given to the fact the advances in know-how that have been achieved have been far more impressive in certain areas of human activity than in others. Thus, in comparison with the vast advances in know-how that have been achieved in various areas of medicine, there has been very little progress made in the productivity and effectiveness of educational practice. Within medicine, while a number of diseases that used to be terrors now can be prevented or treated effectively, little progress has been made on curing, much less preventing, various kinds of cancers. We now have the know-how to palliate AIDS, but not cure it, and the cost of treatment is extremely high.

If attention is shifted from operating technologies and practice to modes of coordinating and managing work in the organizations that operate those technologies, one sees a comparable unevenness of progress. Within business organizations there have been great advances in ability to manage inventories, but it is not clear that strategic decision making regarding mergers and acquisitions is any better now than it was fifty years ago. And while we now have medicines and practices for dealing effectively with a variety of diseases, the management of hospitals, and the governing of the healthcare system more broadly, remain serious problems in virtually all societies.

Clearly I am defining know-how here relatively broadly, including not only know-how bearing on “technology” as that term conventionally is conceived, but also know-how bearing on operating practice more generally, as the know-how involved in teaching children how to read. I also am including as know-how the learned skills involved in managing a complex of many people and activities that need to be coordinated if the desired results are to be achieved. These latter kinds of “know-how” seem to be significantly harder to advance than know-how that is associated with technologies in the more narrow sense of that term. But if that is so, that is an important insight, and a phenomenon that needs to be explained.

For some years now I have been trying to develop a better understanding of what lies behind the very uneven evolution of human know-how, in the sense above, and of what if anything can be done to speed up the advance of know-how in areas where to-date progress has been very limited. One obvious potential explanation is that significant resources have been put into efforts to advance know-how of certain kinds, while only limited resources have been put into trying to advance other kinds of know-how. This certainly is part of the reason why, for example, there has been much more progress in preventing or dealing with diseases that afflict

residents of rich countries, than diseases which are mainly prominent in countries where the people and the government are very poor.

However, aside from cases like this, I would argue that differences in “effective demand” account for only a small portion of the uneven evolution of human know-how that we have experienced. Thus, progress in dealing with breast cancers, while nontrivial, is nowhere near as great as progress in dealing with small pox, despite the very large amounts of money, both private (drug companies) and public (NIH), that has been put into finding a cure for breast cancer. While it is true that far fewer resources have been put into trying to advance educational practice than trying to improve medical practice, an important reason for this is that little has been achieved from the efforts to advance educational practice that have been mounted, and there is widespread skepticism that simply putting more money into the effort would accomplish much.

I think it evident that a principal reason that certain areas and kinds of know-how, whose advance is widely agreed would yield large payoffs, have progressed hardly at all is that progress in these areas has proved very difficult to achieve.

Particularly among scientists it has been proposed that the reason why certain fields of practice, certain bodies of know-how, are able to advance rapidly, while others are very difficult to advance, lies in differences in the strength in the underlying sciences. I believe there is a lot to this argument. However, I will argue that the relationships between the ability to advance practical know-how, and the strength of the sciences underlying that know-how, are complex, not simple. And in any case, a theory that tries to explain the uneven evolution of know-how on the uneven strength of scientific understanding would seem to have to come to grips with the

question of why certain bodies of scientific knowledge seem to have advanced much more rapidly than others. These issues will be a central focus of this essay.

II. Technological Advance as an Evolutionary Process: What Makes Progress Easy or Difficult

Over the past half-century a large interdisciplinary body of scholarship has developed concerned with understanding the processes through which technologies, of the sort generally given that name, evolve over time. In contrast, there has been much less research and writing on the evolution of teaching practices, or modes of business organization and management. In the first part of this discussion, where I will be reviewing what is known about how “technologies” advance, the referent will be mostly technologies of the sort conventionally so designated. Later, where I explore what makes progress hard or easy, I open up the term to include “technologies” more widely construed.

First of all, scholars of technological advance from a wide variety of disciplines have converged on the proposition that technological advance needs to be understood as proceeding through an evolutionary process. The process is evolutionary in the sense that at any time there is generally a wide variety of efforts going on to advance the 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 largely through an ex-post selection process.

However, the proposition that technology advance is evolutionary in this sense in no way denies, or plays down, the role of human purpose in the process, or the often extremely powerful body of understanding and technique used to guide the efforts of those who seek to advance technology. Efforts at invention or innovation are by no means totally blind, or strictly random, as often is assumed to be the case regarding biological “mutation.” A strong base of

underlying knowledge helps those that attempt to improve technology to focus their efforts on pathways that have considerable promise, and to avoid those that are likely dead ends. A strong knowledge base also often permits a variety of tests to be made in the course of an inventive effort that can confirm earlier expectations or indicate changes that need to be made. However, hard proof of what works and what does not, and what works better than what, must be learned through actual experience and actual competition.

As the above discussion suggests, technology needs to be understood as comprising both a body of practice, and a body of understanding. Professionals in the field possess not only a considerable amount of “know-how” regarding effective practice, but also a body of more analytic knowledge that explains or rationalizes why practices work as they do, illuminates current limitations of the technology, and suggests promising ways through which improvements may be achieved. And in the process of technological advance, practice and understanding co-evolve. 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 artifact or practice. New understandings learned through this route in turn provide clues and opportunities to the further advance of technique.

But, of course, there is much more to the story about how understanding advances than simply learning by doing and using. In the contemporary world, most important technologies are associated with specific institutionalized fields of engineering or applications oriented science. The codified part of the relevant body of understanding is largely contained in these fields, and serves as the basis for the training of new technologists and applied scientists. And these also are fields of research. I will say more about these fields of science shortly.

From this perspective -- that the advance of know-how is an evolutionary process, but with efforts to improve practice guided by the prevailing body of understanding -- what are the characteristics of a field that can enable progress to be rapid and cumulatively great, or on the other hand, limit the possibilities for sustained progress? I propose the following

First, those that employ a practice must be able to identify and control relatively tightly what they are doing, and to return it to a standard if it inadvertently wanders away. They must be able to replicate it, if that is desired. The evolutionary perspective points to ability to learn from experiments, natural or deliberate, as the key to being able to make progress. But one can not learn from experiments if one does not have ability to identify, specify, control, and replicate. These characteristics also are necessary for a successful practice used by one actor is to be imitable by others. In the language developed some time ago by Sidney Winter and myself, for progress to be made the practices involved must have a certain amount of the “routine” about them.

Second, the criteria for better performance must be clear and relatively stable, and competing practices must differ non-trivially in efficacy under those criteria. Further, the evidence of efficacy must be relatively sharp, and available in a timely fashion.

Third, it must be possible to focus experimentation on pathways likely to be promising, and gain reliable information relatively quickly as to whether an approach being explored is promising or not. As proposed above, these abilities depend to a considerable degree on the strength of the understanding underlying the technology in question. Here I note further that, if the underlying understanding is strong enough, it may be possible to learn from experiments done off-line, in R and D as it were. There are obvious advantages in being able to do this. However, it is important to recognize that, for off-line R and D to be effective, understanding

must be strong enough so that what one learns from theoretical calculations and through experimenting with simplified models can be applied to actual practice.

I note that ability to learn from R and D obviously presupposes the first two conditions I laid out. But these themselves are not independent of the understandings that enable R and D to be effective. On the one hand, a strong body of underlying understanding can enable identification of the key elements of a practice that determine performance, and hence facilitate effective control and replication. Similarly, strong understanding enhances ability to identify key indicators of strong or weak performance, and hence to facilitate evaluation.

The discussion above supports, and fleshes out, the common proposition that progress in operating know-how in a field is greatly facilitated if there is a strong body of underlying scientific understanding. However, I would like to argue that the conventional wisdom is blind to the fact that causation runs here from the nature of the technology to the ability to develop a strong underlying science, as well as from a strong science to ability to advance the technology.

Above I suggested that the body of scientific understanding underlying a technology tends to be contained in the applications oriented sciences, including the engineering fields. These kinds of sciences are explicitly oriented to understanding of the technology, and the subject matters that the technologies are intended to deal with. It is a mistake to see the applications oriented sciences and engineering disciplines as simply applications of the fundamental knowledge garnered by more “basic” fields of study. A large part of the research that goes on in these fields involves analysis of and experimentation with aspects of practice, or the subject matter that practice addresses. If one can not closely identify, manipulate, and reproduce if wanted, what one is experimenting with, one can not learn much from this kind of research. I note, highlight, that an essential part of the understandings contained in successful

engineering disciplines and applications oriented sciences relates to principles that have been found relevant to the operation of the relevant field of practice, and knowledge about how to advance it. In Thomas Kuhn's terms, these scientific fields provide a paradigm for work to advance practice, as well as work to advance the science. The paradigms they provide may, or may not, have a solid basis in more fundamental science, although for reasons I will discuss later the most powerful of such fields do have such a basis.

Put more generally, while the rapid sustained advance of a field of practice is greatly facilitated by the existence of a strong relevant body of scientific understanding, the principle that practice and understanding co-evolve holds for the scientific part of that understanding, as well as the part gained through learning by doing and using. When the advance of practice in a field is difficult for the reasons I laid out above, it may be difficult to advance the science that underlies that field, and for many of the same reasons.

III. Medical Practice

The conquest of many infectious diseases, at least in countries where the people and the government are rich enough to afford the treatments, surely ranks among the advances in know-how achieved over the last century and a half that have brought greatest benefit to humankind. And in view of Pasteur's role in these developments, it seems highly appropriate at this occasion to discuss some aspects of these developments.

I want to observe that the advances here were very much the result of "research" and the advance of "scientific understanding". Pasteur's research provided convincing evidence for the theory (which some had held for a long time) that at least certain diseases were caused by microorganisms. This enabled researchers concerned with particular diseases to try to identify a particular microorganism as a cause. And the treatment Pasteur devised for a couple of the

diseases he was studying -- development of a vaccine or serum and its application – provided a general model that was followed by others in the search for treatments of other diseases, However, Pasteur and other researchers who followed his lead did not provide any light on just how microorganisms caused disease. That understanding came many years later. And understanding of why vaccines worked (on some diseases) had to await the development many years later of understanding of the immune system. Nor was it clear for many years just what human diseases were the result of infections and which ones, like vitamin deficiency diseases, and most cancers, were not.

The point I want to make here is that the kind of understanding made by Pasteur and the other great scientists who contributed to the field was very much “practice” and “problem” oriented. And that understanding provided a paradigm both for efforts to deal with other diseases, and to advance the relevant science.

In particular, the scientific understandings won in the late nineteenth century provided light for the search for causes and treatments. While the light was quite diffuse, it pointed efforts to deal with a disease towards trying to find a microbe that was the culprit. Once a microbe could be firmly associated with a particular disease, biomedical researchers could search for a cure. And if an effective vaccine or serum was developed, the techniques of vaccine production and application tended to be reliable enough, at least after a period of learning by doing and using, so that vaccines and a vaccinations became part of a well specified routine to deal with that disease.

The identification and development of pharmaceuticals that killed microorganisms in the body came later. As with treatment by vaccines and serums, use of antibiotics (to use the modern

term) was and is effective on certain diseases because the substances and procedures are highly routinized, and they work with almost all patients that have the disease in question.

My reading of the history of the evolution of effective know-how for disease prevention and treatment tells me that in the successful cases I know about the effective procedure involved a well specified routine, generally built around a chemical substance or other artifact of some kind. The research which led to the successful development involved a search for such an artifact, and the testing it out first in “models” of the human system (both in vitro and animal models), and later on humans. In more modern time, this search has been guided by scientific understanding of the human body and of disease processes to a much greater extent than in Pasteur’s day. But as in that earlier era, the light shed by basic scientific knowledge has been diffuse and seldom sharply focused, except in the last stages of a successful development.

Many human illnesses continue to defy this mode of attack. While I have a number of propositions I could put forward about the characteristics of these diseases, space and time constraints prevent me from doing that here. Rather, I must turn now to the case of education.

IV. Education

Today, virtually all modern societies are frustrated with the performance of their systems of primary and secondary education. Costs per student have been rising briskly, a symptom of the fact that productivity in education has been almost static, in contrast with the rapid rates of productivity growth in many other sectors of the economy. And there is growing dissatisfaction with the effectiveness of the system in providing children with a strong education, with particular concerns that a large fraction of children in disadvantaged economic conditions are not taking to what is taught.

In the United States over the past thirty years there have been a number of attempts to raise educational productivity and effectiveness through expanding and reforming educational research and development. The example of the success of research in various areas of medicine has been used as an example which gives hope for similar success in education. But very little of value has come out of the efforts. It seems apparent that the conditions that make research powerful in some other areas are not strong in education.. There do not seem to be any ways of learning what will work in practice through theoretical analysis, or controlled experiment short of a full scale trial that is effectively an on-line operation. And what works in a controlled trial setting is often very hard to transfer to another setting without experiencing significant modification.

The situation in education thus contrasts sharply with that in medicine, where a lot can be learned from in vitro research and research on animal models, and where, after a period of testing and adjustment, a new pharmaceutical or practice usually can be transferred reliably from operation in a controlled setting, to more general use.

There would appear to be several reasons. One is that, while the state of understanding regarding education provides a broad set of principles regarding practice, the devil is in the details, and the details can only be established in practice. Second, even with experience in actual practice, it is difficult to identify the key elements that are influencing performance, and therefore to be able to specify what is essential and what is not if there is an effort to reproduce or transfer a practice that seems to be performing well. Third, what works in one context may not work in another, and it is generally impossible to fully control context variables.

Obviously some of these same problems show up in efforts to advance practice in medicine. But generally they are far less severe there.

These characteristics reflect, as both effect and cause, that the scientific understanding bearing on educational practice is not strong. The fields of research that one would hope would illuminate the educational process and guide efforts at improvement in fact shed only a dim light. On the one hand, research that is focused on subject matter that arguably is closely related to the educational process at best seems to yield course grained and unreliable conclusions. On the other hand, scientific research that limits itself to subject matter where relatively fine grained and reliable results can be attained, as modern study of the physical workings of the brain, tends to generate results that are a far distance from anything useful in the educational process. There is no field like electrical engineering, or bacteriology or oncology, that relates to and guides analysis and experiment bearing on educational practice.

I believe a major reason why this is so is that little success has been achieved to date in creating or discovering in education tight routines, or aspects of them, that are powerfully effective and have wide applicability. Physical equipment used in education, like textbooks, film, recently computers and computer programs, build in elements of the routine. However, to date no artifact has been developed for educational purposes that has the power of an antibiotic for dealing with an infection, or a computer used for data processing. Artifacts in education play a role as tools an effective teacher can use as elements of the practices he or she employs, but do not provide the core of those practices.

V. Modes of Organizing and Governing Economic Activity

One very important feature of the practice of education, at least in the standard school setting, is that issues of organization and governance are central. There are many children in the classroom. The teacher has to maintain order in the classroom, and establish and hold an

environment where a student's peers support learning rather than denigrating it. Teaching would be a hard enough task absent these complications, and it would be difficult to find or develop a well defined set of routines that would work effectively for most teachers and individual students, without a lot of individualized tuning, even if they were not there. But the fact that they are there certainly is an important reason why it has proved so hard to develop simple effective routines of teaching children to read, analogous to what doctors do to deal with a child's throat infection.

Bhaven Sampat and I have used the term "social technologies" to refer to ways of doing things where interaction among people is a much larger part of what is going on than in activities where people largely interact with artifacts of some kind to get things done; the latter we called "physical technologies". We have argued that it is much more difficult to achieve significant cumulative progress of social technologies than physical technologies. Education clearly involves social technologies to a significant extent. There is only limited employment of physical technologies. The governance of large organizations is largely through social technologies. This, I would argue, is why our know how regarding governance and organization has progressed so slowly.

There certainly are examples of aspects of organizational management where progress has been dramatic. As I suggested earlier, inventory management is one such area. But the reason for such progress here clearly is that physical equipment, computers and a linked electronic network recording sales and generating new orders, that is a physical technology, has come to replace what had been largely a social technology. An interesting question is why efficient low cost computer related technology was able to be employed so effectively in inventory management, when it has proved so difficult to develop effective teaching practices in

which computers play the central role. The answer, I would argue, is that the criteria for effective inventory management involve only a few dimensions, and the particular routines used to manage inventories can be kept in relative isolation from other activities that are involved in business organization and practice.

The contrast with the apparent inability of business management to do significantly better than they have in the past with decisions and processes of mergers and acquisitions is striking. Here, the criteria for “success” are complex in the short run, and if long run profitability is the ultimate criterion, this cannot be assessed except in the long run. And in any case, what a merger brings to a firm cannot be isolated from a wide variety of firm operations.

The same sort of considerations obtain regarding activities at a hospital. Decision making and administration of some activities can be routinized, because the objectives are relatively simple and performance easy to evaluate, and because how the operations in question are performed has little impact on and is not much affected by operations in other arenas. This is not so of other kinds of activities. It is possible to advance the effectiveness of the former through operations research and in some cases by mechanization, while the latter may be very difficult to advance over time.

VI. The Science-Technology Connection

I want to conclude this talk by pulling together some of the comments I have made on the connections between technological advance and the strength of the underlying sciences, and extending them a bit. I have stressed the key role of the applications oriented sciences in the advance of practical know-how. I note here that in several recent survey studies, industry respondents almost always have said that the applications oriented sciences are the ones from

which they draw most in their efforts to develop new products and processes. This finding should not be interpreted as indicating that the fundamental sciences, like theoretical physics, or mathematics, are irrelevant to technological progress. The strong applications oriented sciences and engineering disciplines draw extensively and productively on sciences “upstream” from them, as it were. Since the fields of technology that are advancing rapidly almost always are supported by a strong underlying applications oriented science, and the latter tend to be marked by good ability to draw from the more fundamental sciences, in effect one sees rapid advance of practical know-how where the practice, the applications oriented science, and basic scientific understanding, are relatively closely connected.

My reading of technological histories in such fields suggests that the key has been the designing of practice around what is known scientifically. This has been far easier in fields of practice where artifacts – chemical substances, mechanical or electrical devices -- can be made to play key roles. Where this is the case, it has proved possible to develop strong fields of applications oriented sciences concerned with the nature of these artifacts and the context in which they operate. In turn, these applications oriented sciences can draw on the understandings of the natural, and particularly the physical sciences. Note that the strength here is the result of a two-way connection. On the one hand, the prominent use of artifacts facilitates the development of a strong applications oriented science. On the other hand, an applications oriented science focused on artifacts and their context of use can draw on the strong natural sciences. I can identify a few bodies of strongly effective know-how where artifacts do not play a central role, but not many.

However, in many fields of practice it has not been possible to find or develop artifacts that are effective in enabling productive routines. As noted, many human diseases still lack an

effective prevention or cure. And in fields like education, or stopping crime, while it is possible that new artifacts will help performance to increase somewhat, it is extremely unlikely that they will be effective, like an antibiotic is for an infection. Thus practice in these fields, and the applications oriented sciences underlying practice, have little contact with the strong physical sciences.

I suggested above that the areas of know-how where we have had the greatest difficulty achieving significant progress have been those where practice mostly involves social technologies, and there has been little success in discovering or inventing powerful physical technologies to help with the work. These social technologies in turn are illuminated, but only dimly, by the behavioral and social sciences, most of which support both basic and applied research. But the knowledge that has been won by these fields is far less powerful than the knowledge that had been won through the natural sciences. I would argue that the reason lies in the nature of the ontologies. By now what I think are the key differences should be obvious.

I find it interesting that in recent years a number of philosophers of science have given up on the idea that the behavioral and social sciences ultimately can be as strong as physics, or that they productively can use the methods that have worked so well in physics. I think they are quite correct on this judgment. And I think this means that the bodies of practice that are supported by the behavioral and social sciences will continue to be areas where the progress of know-how is slow.

I am an economist. These last remarks would seem naturally to lead me into some comments on the nature of my “science” and its relationship to economic policy making. But I have run out of time.