Industrial and Corporate Change, Volume 17, Number 3, pp. 485–497
doi:10.1093/icc/dtn010
Advance Access published May 2, 2008

Factors affecting the power of technological paradigms

Richard R. Nelson

It is clear that the power of “technological paradigms” proposed by Dosi (1982) varies greatly across fields of practice, in the sense that in certain field’s progress has been much more rapid than in others where comparable resources have been applied to the effort. This essay explores the factors behind these differences. It proposes that one important factor is the extent to which the technology in a field is controllable and replicable. Another factor is the strength of the supporting sciences. It is argued that these factors are strongly intertwined with the causal arrows going both ways.

1. Introduction

It now is more than 25 years since Giovanni’s article appeared in print (Dosi, 1982). It is still a pleasure to read: rich in insights regarding important empirical phenomena, and pregnant with ideas that beg to be developed farther, the concept of technological paradigms clearly has provided a major part of the evolving scientific paradigm guiding research on technological change that has held sway over the last quarter century.

Giovanni’s technological paradigm concept (the concept of a technological regime developed in Nelson and Winter, 1977, had many things in common) accomplished a number of important things. It brought together and clarified the relationships between the literatures on the role of demand and the role of technological opportunities in influencing the direction of technological advance by proposing that both beliefs about what kinds of advances would be useful and what kinds were feasible were involved in a paradigm. It argued that both understandings closely tied to practice and how to advance practice, and deeper and perhaps more abstract understandings of principles bearing on practice, were part of a paradigm, thus recognizing a rough distinction between technological and scientific knowledge and pointing to how they were connected. It proposed that there might be several quite different paradigms recognized as possible ones for advancing the effectiveness of technology aimed to meet a particular range of needs, but that at any time most of the actual work tended be concentrated on one of these, and focused in particular ways; the result of this was that, as long as a particular paradigm held sway, technological advance tended to proceed along a particular trajectory.

Giovanni's technological paradigm concept was very much part of his broader commitment to an evolutionary theory of technological advance. But the evolutionary process in the theory that Giovanni was helping to articulate differed in essential ways from biological evolution. In particular, central in the paradigm concept is both a notion of conscious direction of efforts to advance practice, and recognition that efforts in that broad direction are far from blind, but are strongly oriented by the body of human know-how about how to advance practice. But the process also involves strong evolutionary elements in the sense that, while paradigmatic knowledge helps those aiming to advance a field to avoid many dead ends and to point themselves in directions that are likely to be fruitful, it is in general impossible to judge accurately in advance what will work better than what. Therefore progress in a field generally involves the continuing creation of a variety of alternatives in competition with each other and with prevailing practice, with ex-post selection determining winners and losers.

I note that many historians of technological advance have developed this point of view independently. Walter Vincenti (1990) is a good example. More generally, it is fair to say, as I noted above, that this perspective now is a major part of the broad paradigm employed by most empirically oriented scholars of technological advance, regardless of the academic discipline within which they got their original training. This certainly has been so of my own work.

Much of that recently has been focused on a problem that, using Giovanni's paradigm language, can be posed this way (Nelson, 2003). It is apparent that the paradigms that have molded the advance of technology (or know-how: I will use the terms interchangeably), concerned with the meeting of some human needs, have been far more effective in achieving progress than have those concerned with other areas of need. Thus in comparison with the great advances in the technologies determining the effectiveness and cost of long distance communication, and computation, and many areas of medicine, very little progress has been made in the productivity and effectiveness of educational practice. Or, within medicine, while a number of diseases that used to be terrors now can be prevented or treated effectively, little progress has been made in curing cystic fibrosis, or many kinds of cancers. What are the reasons why paradigms in some areas have been much more effective in advancing know-how than paradigms in others. What explains the power, or lack of it, of a technological paradigm?

One obvious potential explanation is that significant resources have been put into the paradigms guiding 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 a good 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 know-how that we have experienced. Thus, progress in dealing with breast cancers, while non-trivial, 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 (government agencies), 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 is 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. Effective paradigms to guide and power the advance of know-how in these areas have not been developed, at least not yet.

Particularly, scientists are prone to argue that a key difference between fields where technological paradigms have supported sustained and rapid progress in practice, and those where practical progress has been slow, is that the former are supported by strong scientific knowledge and the latter are not. I believe that there is merit to this argument. However, I will argue that the relationships between the ability to advance practical know-how and the strength of scientific knowledge underlying that know-how are complex not simple. And in any case, differences in the strength of various fields of science are an important matter to be explained in its own right.

2. What makes progress easy or difficult?

Professionals in a field operating within a technological paradigm work with a wide variety of knowledge; (for a splendid discussion see Pavitt, 1987). An important part of their expertise is empirical knowledge about the nature of current practice and its performance, including its strengths and weaknesses. Empirical knowledge about practice is complemented by a body of more analytic knowledge that explains or rationalizes why practices work as they do, illuminates the factors behind current limitations of the technology, and suggests promising pathways for improving it. A portion of the practical analytic knowledge may be supported by knowledge of more fundamental scientific principles.

In virtually all fields of technology, much of professional knowledge is acquired through learning by doing and using, more generally through detailed acquaintance with the practice. But 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. Vincenti's discussion (1990) of engineering knowledge bearing on aircraft design and its evolution, and Constant's history of the turbojet (1980) are particularly apt.

From this perspective—that the advance of technology is an evolutionary process, but with efforts to improve practice guided by the prevailing body of understanding, some of the product of experience and some coming from an institutionalized science—what are the characteristics of a field that can enable progress to be rapid and sustained, or on the other hand, limit the possibilities for sustained advance. I suggest the following.

First, those that employ a practice, and the professionals responsible for maintaining, reproducing, and advancing it, must be able to identify and control relatively tightly the practice they are using or analyzing, and be able to return it to a standard if it inadvertently wanders away. They must be able to replicate it when they so desire. The evolutionary perspective points to being able to learn from experience and variation, from experiments natural or deliberate, as the key to being able to make progress. But one cannot learn from experience or experiment if one does not have ability to identify, control, and replicate. In the language developed some time ago by Sidney Winter and myself (Nelson and Winter, 1982), for progress to be made the practices involved must have a certain amount of the "routine" about them.

Second, the criteria for assessing efficacy of performance must be clear and relatively stable, and it must be possible to compare competing practices under those criteria. Further, the evidence of efficacy must be relatively sharp, and available in a timely fashion. I note that for this condition to be met, the first condition is of central importance. If one cannot identify and control a practice, one cannot know what practice it is that is associated with the outcomes one is observing. But of course while condition one is necessary for evaluation of alternative practices, it is not sufficient.

Third, while if conditions one and two hold, under favorable circumstances a lot can be learned from unplanned natural experiments and the variety of practice at any time, it is clearly of great potential advantage to be able to engage in and learn from deliberate controlled experimentation. And it is specially valuable if one can gain reliable information bearing on practice from experimenting with and analyzing simplified versions or models, off line from actual practice. The costs, including the opportunity costs, of exploring different designs, ways of doing things more generally, then can be greatly reduced. And so too the riskyness of trying out things that might be very effective, but also might prove worthless or worse. Indeed, such off-line exploration using simplified models is much of what is involved in what is now called R&D.

All of these abilities obviously depend to a considerable degree on the strength of the understanding underlying the technology in question. Thus a strong body of understanding can enable identification of the key elements of a practice that determine performance, and hence facilitate effective control and replication. Strong understanding enhances ability to identify key indicators of strong or weak performance, and hence to facilitate evaluation of prevailing practice, and guide attempts at improvements. And to be able to learn from deliberate experiments, and from off-line R&D more generally, understanding must be strong enough so that what one learns from experimenting in a controlled setting and with simplified models can be applied to actual practice.

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. Empirical studies of the knowledge bases of technological advance indicate that, as I suggested above, much of the body of scientific understanding underlying a technology tends to be contained in the applications oriented sciences (See for example, Klevorick *et al.*, 1995). These fields are explicitly oriented to understanding of the technology, and the subject matters that the technologies are intended to deal with. A large part of the research that goes on in these kinds of fields involves analysis of and experimentation with aspects of practice, or the subject matter that practice addresses.

If that research is successful, these fields can support powerful technological paradigms for advancing practice. The paradigms they provide may, or may not, have a solid basis in more fundamental science. For reasons I will discuss later the most powerful of such fields do have such a basis. But, as Giovanni has argued persuasively, they are bodies of technique and knowledge in their own right.

While it is broadly understood that a strong body of scientific understanding enables technological progress to be rapid and sustained, conventional wisdom tends to be blind to the fact that, for the engineering fields and applications oriented sciences, the causation here runs 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. I have noted that much of the research in the engineering disciplines and applications oriented sciences aims to develop understanding of what is going on in the operation of the relevant field of practice, so as to illuminate how to advance it. However, if one cannot closely identify, control, manipulate, reproduce, and evaluate what one is experimenting with, one cannot learn much from this kind of research.

I do not mean here to deny the role that the fundamental sciences—fields generally thought of as not closely tied to areas of application—often play in supporting technological advance, but rather to put that role in context. In the first place, major parts of what generally are thought of as fundamental sciences often have been developed in efforts to understand a technology or an area of application; consider for example thermodynamics, or many sub-fields of chemistry.

Second, while scientific knowledge won in efforts where advancing practice was not an important goal often serves to illuminate practice, it tends to do so as supports for analysis in an applications oriented science; thus much of physics useful in analysis of electronics technology does so through incorporation in the analytic structures of fields like electrical engineering and materials science. Third, many aspects of quite powerful fields of applications oriented science have only limited connection with fundamental scientific understandings. This is the case, as I will argue below, with much of biomedical science.

I now turn to some examples to make these arguments more concrete.

3. Dealing with infectious diseases and with education

The conquest of many infectious diseases surely ranks among the advances in know-how achieved over the last century and a half that have brought greatest benefit to humankind. I now want to argue that the kinds of understanding developed by the evolving bio-medical sciences, that underlay and enabled these developments, were and continues to be largely oriented towards illuminating paths to better treatments of disease. Where, successful one tends to find both the practice, and efforts to improve practice, proceeding under relatively tight technological paradigms that operate under the favorable conditions laid out in the preceding section. The more basic understanding under the applications oriented paradigms sometimes may be strong, but often is relatively shallow. (Porter, 1997, provides detailed accounts of the cases discussed below).

The latter certainly was the case when Pasteur first provided convincing evidence for the theory (which some had held for a long time) that at least certain diseases were caused by micro-organisms. This oriented researchers concerned with particular diseases to try to identify a particular micro-organism as a cause. And the treatment Pasteur developed for a couple of the diseases he was studying—a vaccine employing dead or weakened micro-organisms implicated in the disease—provided a general model that was followed by others in the search for treatments of other diseases. And an effective vaccine, once developed, could be applied widely and routinely across the relevant patient population.

That is, after the great discoveries of Pasteur and Koch, a broadly effective "technological paradigm" was developed for dealing with an expanding range of infectious diseases.

However, Pasteur and other researchers who followed his lead did not provide any light on just how micro-organisms caused disease. That knowledge came later. Understanding of why and how vaccines worked (on certain diseases) had to await the development many years later of understanding of the immune system. Nor it was clear for many years just what human diseases were the results of infections and which ones, like vitamin deficiency diseases, cystic fibrosis, and most cancers, were not.

Similarly, while a micro-organism theory of disease clearly pointed attention in this direction, the identification and development of chemical compounds, natural or man made, that killed micro-organisms 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. However, until recently at least the search for effective antibiotics has tended to proceed without deep understanding of why a particular substance might be effective. And, even today, the reasons why many antibiotics work as they do is not well understood.

More generally, my reading of the history of the evolution of effective know-how for disease prevention and treatment tells me that in the successful cases 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 study of the nature of the disease which identified at least broadly the kind of artifact that might be effective, a search for or attempt to design 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 many of the successful cases I know about, there was a broad paradigm available which had proved helpful in the search for cures in related illnesses. In turn that broad paradigm was supported, but in most cases only loosely, by deeper scientific understanding.

In stressing the importance of practice oriented paradigms as guiding successful efforts to deal with diseases, I am not trying to deny the important role that basic scientific understanding of the human body and of disease processes often have played in sharpen these paradigms, or leading to new and better ones. Thus understanding of body chemistry and the working of cells in regulating those processes has been important in the discovery and design of drugs that have been effective in cardiovascular disease. Understanding of the genetic changes in cells that make the cancerous has led to a new orientation of searching for effective ways to stop cancers. But advances in fundamental understandings of this sort almost never point sharply to cures, but rather help in the discovery of cures by becoming part of the operating technological paradigm.

I also note that there are a number of cases where a disease is well understood scientifically at a relatively fundamental level, but no effective prevention or cure has yet been developed. Thus biomedical scientists now understand what goes wrong in body chemistry under cystic fibrosis, and the genetic sources of the problem. That understanding may ultimately lead to an effective paradigm for finding or developing better treatments. But to date progress has been halting at best.

I want to turn now to an area of human activity where there has been very little progress. 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 (Murnane and Nelson, 2007).

In the United States over the past 30 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 experiment short of a full-scale trial that is in effect a controlled on-line operation. And what works in a controlled full-scale trial setting is often very hard to transfer effectively to another setting.

The situation in education thus contrasts sharply with that in many areas of 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.

The principal reason, I would argue, is that the first two basic conditions for a field of practice to progress that I laid out earlier are not there in education. There is only limited ability to control tightly, specify accurately, and replicate educational practices, that seems to be effective. Related to that, it has proved nearly impossible to identify the key elements that lie behind effective practice when that occurs, other than very broadly. As a result, experience with trying to transfer to another setting a practice that seems to be working well has not been particularly good.

Also, there are significant problems in assessing the performance of an educational practice, even when a practice can be tightly identified. Test scores measure only a part of what education is trying to achieve. And the long run effects of education on ability to get a good job, be productive generally, be a good citizen, take many years to see, and even then are almost impossible to disentangle from other factors influencing a persons life and success.

These same problems obviously limit ability to design experiments that have a good chance of being productive, as well as limiting the ability to spread widely practices that look good in an experimental setting

These characteristics reflect, again 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 directly on educational practice 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 that an effective teacher can use as elements of the practices he or she employs, but do not provide the core for those practices. Effective practice seems to be closely tuned to the skills and personality of the individual teacher, and the backgrounds, knowledge and motivations of the student body. This is a context in which good educational practice can emerge. But it is not one within which it is easy to make significant and continuing progress in teaching methods that are broadly applicable to the student population.

4. Reflections on the science-technology connections

I now want to pull together some of the comments I have made on the connections between technological advance and the strength of the underlying sciences, and extend them a bit.

In a recent article Paul Nightingale (2004) has put forth an argument, very similar to mine, that it is basically through design that builds in the factors known to be associated with good performance, and shields an artifact or process from influences that can muddle how and how well it works, that the performance of a technology becomes reasonably reliable and predictable. Nightingale takes a position, similar to mine, that it is only through long experience, and research of the sort we now associate with the applications oriented sciences, that strong knowledge that enables good design comes into existence.

I have argued that this perspective should not be interpreted as indicating that the fundamental sciences, like molecular biology, or theoretical physics, or mathematics, are irrelevant to technological progress. The stronger of the applications oriented sciences and engineering disciplines draw extensively and productively from sciences "upstream" from them, as it were.

However, the language "draw from" is relevant in this context. As I have argued earlier, advances in fundamental understanding almost never lead directly to advances in technologies. Rather, the route is more indirect. As historians and philosophers of science (see for example Cartwright, 1983) increasingly have argued, the power of a science, its ability to develop reliable knowledge, generally is dependent upon its ability to isolate and control the phenomena it is investigating, and the "laws" the science finds relate to what goes on in those controlled contexts. As a consequence those laws may not provide much understanding of what is going on in messy contexts where controls and modes of isolation are not effective.

But on the other hand, scientific knowledge that points to something that can be productively used in a practice, if the circumstances could be controlled, can provide very useful information to designers, whose province of inquiry very much does include modes of control and isolation. Nightingale (2004) poses the issue here as whether those seeking to advance the technology are able to control the way it operates so as to fit the conditions under which scientific research has shown a particular desired thing happens. In this way, advances in basic scientific understanding can sharpen up and provide new insights to the applications oriented sciences that draw from them.

My reading of what has been going on in many of the fields where technological progress has been rapid and sustained suggests that, the key often has been the ability to design a controllable replicable practice that is broadly effective around what is understood scientifically. In some fields the identification of a controllable broadly effective practice preceded scientific understanding of what makes the practice work and what influences it must be shielded from, but once that practice was established a body of understanding bearing on it came into being and facilitated its advance. In other cases, understanding led to the development of controllable broadly effective practice. But whatever the genesis, where progress in practice has been rapid and sustained, practice and understanding have tended to co-evolve.

This has been far easier in fields of practice where physical 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 good part because the artifacts are amenable to control, and replication, and therefore are good subjects for experimentation. This is not a new argument. Rosenberg (1974) shows that Marx made it a long time ago, in his analysis of why capitalist modes of production were amenable to continuing improvement through the application of science.

In turn, these applications oriented sciences often can draw on the understandings of the natural sciences, whose domain of study is amenable to the control and isolation that permits the development of reliable knowledge. While that knowledge is of what is going on in stylized contexts, it provides insights into how artifacts might be made to work if they could be similarly controlled.

On the other hand, a strong underlying science does not necessary engender a strong broad technological paradigm. In some fields of practice where the underlying basic sciences are reasonably strong, it often has proved difficult to find or develop controllable replicable ways of doing things, based on physical artifacts, that are broadly effective. As noted, many human diseases still lack an effective prevention or cure, despite the existence of a strong body of both applications oriented and basic science in the biomedical field.

Ohid Yaqub (2008) has proposed an interesting argument regarding why the development of vaccines for certain life threatening diseases has proved so difficult to achieve; HIV AIDS is an important example. The use in humans of vaccines that are not well tested for safety can be extremely dangerous, since the vaccine itself potentially can cause the disease it is designed to prevent. Thus development of an effective vaccine in such contexts requires the ability to do effective testing for safety without trials on humans, prior to any significant testing for efficacy on humans. While for some diseases and proposed vaccines there exist well known reliable in vitro tests that do at least a rough screening, and good reason to believe that follow-on trial on non-human animals will test reliably safety on humans, in other cases this is not so. Where this is not the case, as apparently is the situation with HIV AIDS, learning about the efficacy of a possible new vaccine, without jeopardizing safety, can be very difficult, expensive, and time consuming.

In another essay I have suggested the areas of know-how where, we have had the greatest difficulty achieving significant progress have been those where practice mostly involves what Bhaven Sampat and I (Nelson and Sampat, 2001) have called social technologies, practices where interaction among people is a much larger part of what is going on than in practices where people largely interact with artifacts of some kind to get things done. Ways of doing things that are highly sensitive to what people do, and how they interact, are much more difficult to control and replicate than ways of doing things that are largely dominated by physical artifacts. And they are therefore more difficult to advance through experience and research.

I note that, many of the areas of human action where the technologies are largely social are so by default. Education is a canonical example. Despite lots of trying, there has been little success in discovering or inventing powerful physical technologies to help with the work.

Social technologies in turn are illuminated by the behavioral and social sciences, most of which support both basic and applied research. But the illumination is dim. 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 their subject matter. In particular, ability to learn through scientific research depends on the extent to which context can be controlled, and whether or not what is learned in closely controlled contexts provides useful insights into what is going on, and how to operate effectively in, a messier world.

For a number of reasons, many of which by now should be obvious to the reader, these conditions have proved very difficult to achieve in the social sciences.

5. A summing up

The proposal advanced by Giovanni Dosi that the advance of technology, or practical know-how, generally tends to proceed guided and constrained by "technological paradigms" has been become a basic part of the scientific paradigm of scholars studying technological change. In this essay I have highlighted that there have been vast differences in the power of the technological paradigms that have been developed in various arenas of human activity, and put forth what I think are the key factors behind those differences.

My argument supports the widely held view that technological advance tends to proceed rapidly where scientific understanding is strong, and slowly where scientific understanding is weak, but have proposed that the critical fields are the engineering disciplines and the applications oriented sciences. I also have argued that the causal arrows go in both directions. A principal requirement for a technological paradigm to be capable of generating rapid sustained progress is that the practices it is seeking to advance be controllable and replicable. A strong underlying science can help this to be achieved. But on the other hand, unless this is achieved, it is virtually impossible for a strong underlying science to develop.

Acknowledgements

I very much appreciate the helpful comments and suggestions on an earlier version of this paper given to me by Paul Nightingale. As I note in the text, Paul and I are developing very similar points of view on the phenomena addressed here. The author is grateful to the Sloan Foundation for support for the research behind this article.

Address for correspondence

The Earth Institute at Columbia University, and the School of Business at the University of Manchester.

References

Cartwright, N. (1983), *How the Laws of Physics Lie*. Clarendon Press: Oxford.

Constant, E. (1980), *The Origins of the Turbojet Revolution*. Johns Hopkins, University Press: Baltimore.

Dosi, G. (1982), ‘Technological paradigms and technological trajectories,’ *Research Policy*, **11**, 147–162.

Klevorick, A., R. Levin, R. Nelson and S. Winter (1995), ‘On the sources and significance of inter-industry differences in technological opportunities,’ *Research Policy*, **24**, 185–205.

Murnane, R. and R. Nelson (2007), ‘Improving the performance of the education sector: the valuable, challenging, and limited role of random assignment evaluations,’ *Economics of Innovation and New Technology*, **16**, 307–327.

Nelson, R. and S. Winter (1977), ‘In search of useful theory of innovation,’ *Research Policy*, **5**, 36–76.

Nelson, R. and S. Winter (1982), *An Evolutionary Theory of Economic Change*. Harvard University Press: Cambridge.

Nelson, R. and B. Sampat (2001), ‘Making sense of institutions as a factor shaping economic performance,’ *Journal of Economic Behavior and Organization*, **44**, 31–54.

Nelson, R. (2003), ‘On the uneven evolution of human know-how,’ *Research Policy*, **32**, 909–922.

Nightingale, P. (2004), ‘Technological capabilities, invisible infrastructure, and the un-social construction of predictability: the overlooked fixed costs of useful research,’ *Research Policy*, **33**, 1259–1284.

Pavitt, K. (1987), *On the nature of technology*. Lecture given at the University of Sussex, UK, June 23.

Porter, R. (1997), *The Greatest Benefit to Mankind*. W.W. Norton and Company: New York.

Rosenberg, N. (1974), ‘Karl marx on the economic role of science,’ *The Journal of Political Economy*, 713–728.

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

Yaqub, O. (2008), ‘Knowledge Infrastructure and Vaccine Innovation,’ Unpublished D. Phil., SPRU, University of Sussex, UK.