Daniel Sarewitz and Richard R. Nelson

Progress in Know-How
Its Origins and Limits

Over the last several centuries, the advance of know-how to meet a wide range of human needs has been phenomenal. In the area of medicine, for example, we now are able to prevent or cure many infectious diseases that used to be scourges. In other fields of human need and activity, however, progress has been quite limited. For example, despite significant resources put into education, many children fail to learn much in the educational environments provided them by modern societies. In contrast to the history of how we have fought infectious diseases, the ways children are taught to read and write today are not much different than educational practices a century ago. This is true even though the percentage of the population that society expects to have strong reading and writing skills has increased greatly.

We are using the term "know-how" to denote the knowledge, some articulated and some tacit, that guides the actions of skilled agents who aim to achieve a particular practical objective. That is, the state of know-how defines best practice in an arena of human activity. In this paper, we develop a theory that we believe helps to explain why know-how has improved markedly in some areas of human endeavor but not in others. Because this uneven evolution of know-how is such a prominent aspect of human affairs, gaining a better understanding of some of the factors that make progress in certain areas particularly difficult is important in its own right. Moreover, improved understanding here may help us distinguish between problems that are likely to be resolved through improved know-how, and thus be responsive to strategies that emphasize more research and creation of knowledge, and problems that are likely to resist this approach.

To introduce and illustrate the problem of the uneven advance of know-how, we describe an apparent paradox directly relevant to the provision of two basic human needs, literacy and health. Consider two distinct policy goals: to progressively improve the literacy levels of the citizens of the United States, and to eradicate smallpox from the world. The paradox we want to highlight is this: while improving literacy in one of the world's most affluent nations might seem a straightforward, if not simple, task, in fact, persistent effort has yielded little if any appreciable improvement in average reading levels among the nation's K-12 students—this despite the strong national consensus on the importance of making progress in this area and the considerable effort directed at improving the practice of teaching reading. In marked contrast, as everyone knows, smallpox has indeed been eradicated from the globe, despite what might reasonably have been seen as

insuperable challenges, ranging from the difficulties of developing the necessary international cooperation among regimes during the Cold War to the complexities of administering vaccines in regions with little medical infrastructure and, at times, strong cultural aversion to Western medicine.

In the United States, the goal of improved literacy has been articulated by policy-makers since the nation's founding.[1] Over the past several decades, considerable research has been aimed at the problem, along with many policy interventions and a rigorous national assessment process. The goal of preventing smallpox through vaccination has also been articulated for several centuries, and increased understanding of both the disease and the vaccination process led, in the late 1950s, to the conviction among some public health officials that global eradication had become both plausible and practicable.

Because ... uneven evolution of know-how is such a prominent aspect of human affairs, gaining a better understanding of some of the factors that make progress in certain areas particularly difficult is important in its own right. Moreover, improved understanding here may help us distinguish between problems that are likely to be resolved through improved know-how, and thus be responsive to strategies that emphasize more research and creation of knowledge, and problems that are likely to resist this approach.

Why has one effort been successful and the other not? In hindsight, one might offer many ad hoc explanations but, as we will show, these explanations actually raise more questions than they answer. We will argue here that the paradox is explained not by the specific details of each case but by a fundamental difference between the practice of vaccinating people against smallpox and the goal of teaching children to read. The former can build upon a foundational core of effective know-how, whereas the latter activity lacks such a core. Moreover, we will argue that the specific details of each case are in part a reflection of this difference in associated know-how, rather than a cause.

THE IMPORTANCE OF A STANDARDIZED CORE

Readers may initially feel that the differences between vaccinating for smallpox and teaching children to read are so fundamental that comparison is not meaningful, so we want to begin by emphasizing that these activities actually do have important similarities. Both are administered by a professional who has been trained to do the job as effectively as possible. Both typically proceed in a controlled and relatively closed environment—the doctor's office, the schoolroom. The vaccine and the equipment used to give the injection are standardized artifacts, as are the textbooks used in teaching children to read. Each of these bodies of practice has been and continues to be the subject of considerable research that aims to illuminate and improve practice. Moreover, society strongly expects and demands that performance in both of these realms improve over time. Vaccines are expected to provide protection against an ever-expanding array of diseases and to do so with ever fewer adverse effects. Aspirations for schooling have similarly increased: we now expect not only that the children of professionally ambitious middle-class parents possess basic reading skills but that all children gain those skills. Indeed, protection against disease by vaccination and preparation for participation in the activities of modern society via effective teaching are both seen as crucial products that society must deliver to all of its citizens.

However, vaccination against smallpox (not to mention a variety of other infectious diseases) has been an extraordinarily successful public health measure, whereas the teaching of reading skills yields highly uneven results that have not improved much over the years. The reading performance of U.S. students has failed on the whole to improve over the past 35 years (the period for which reliable national data are available),[2] despite the significant resources and attention devoted to research, practice, and assessment. Furthermore, there is little reason to believe that much improvement was occurring before this period.[3] While almost every person vaccinated against smallpox is safe from the disease, an alarming number of children who are taught to read in school never really learn to do so at a level necessary to perform well in today's society. If by strong know-how we mean the ability to do something effectively, our know-how bearing on vaccination has advanced much farther in the past century or so than our know-how bearing on the teaching of reading.

In considering the reasons for this unevenly advancing state of know-how, we want to highlight one striking difference between practice in these two arenas that we believe is closely related to differences in effective know-how. Vaccination is a procedure whose effectiveness is largely due to a physical thing, the vaccine. The vaccine must meet certain standards and be delivered appropriately: the standards require strong quality control on production, distribution, and storage, and effective delivery requires certain training of the people doing the vaccinating. If these conditions are met, the procedure of vaccination has a very high probability of success that is largely independent of the person being vaccinated or the professional doing the vaccination. On the other hand, while the teaching of reading

involves skillfully produced artifacts, like the textbooks used and perhaps even a computer, these items do not provide the basic "go" of the teaching and learning, which depends critically on the skills of the teacher and the attributes of both the classroom environment and the individual students themselves.[4]

The vaccine gives the procedure of vaccination a standardized core that is effective in most contexts. Because the vaccine embodies most of the relevant cause-and-effect relations that connect the action (administering the vaccine) to the goal (preventing smallpox), a practitioner does not face the problem of figuring out what will work on a case-by-case basis. This standardization also enables the effective part of the procedure to remain the same with different practitioners and different patients, and in different cultural and social contexts.

Creating a widely effective standardized core for an activity is almost always the result of a learning experience—a process of improvement, of finding out and homing in on what works—that generally involves a long period of time and many participants.

There are two points we want to highlight here. First, creating a widely effective standardized core for an activity is almost always the result of a learning experience—a process of improvement, of finding out and homing in on what works—that generally involves a long period of time and many participants. Second, while there is something of a chicken-and-egg problem here, cumulative learning is almost impossible unless there is something like a standardized core. If there is such a core, it is possible to learn what procedures work better than others. It is also then possible to improve the core—through experiments, natural or deliberate, that involve that core or a contemplated alternative—with some confidence that what is learned can be applied to standard operating practice and be employed by a wide range of practitioners. Absent such a core, there is no fulcrum upon which to leverage cumulative learning and improved practice.

We believe that the aspects of modern vaccination that we have described are necessary aspects of powerful bodies of know-how. We do not deny the need in many cases to fine-tune a practice to the particular contexts of its use, nor do we deny the variety of effective practices we have in some fields. However, we argue that in order to be powerful and to improve over time, the core of a broad practice must be standardized to a certain and usually considerable extent. Practice based on that core must be effective in most contexts, with the key elements under the control of the agents involved. The process must be replicable and usable by most agents trained in the art; that is, by those who have the relevant know-how. The

process also must be amenable to experimentation, deliberate or natural, and what is learned from such experimentation must be able to be built into the standard routine.

This latter aspect—the ability to learn how to improve practice or to create a significantly better mode of practice in a way that lends itself to widespread use—is obviously critical. Today's powerful bodies of know-how and practice almost always are the fruit of a cumulative learning experience that often has occurred over a long period of time. The experience generally involves both learning by doing and deliberate experimentation—in this modern era, organized research and development. We are arguing here that this kind of cumulative learning is not possible unless there is a standardized or standardizable core to a practice upon which efforts to improve it can build, and which can be replaced with something better if such efforts prove successful.

While there is something of a chicken-and-egg problem here, cumulative learning is almost impossible unless there is something like a standardized core.

The vaccine, which provides these attributes to the process of vaccination, is a physical artifact. Most of the powerful bodies of know-how that have been developed over the years have physical artifacts at their core. But there are some exceptions to this rule; for example, our current know-how about assessing the health of newborn babies.[5] Moreover, the fact that artifacts are involved certainly is not sufficient for a body of know-how to continue to advance. Rather, the key characteristics are standardization and controllability, replicability, and ability to learn from experimentation in a way that is widely relevant to practice.

The term "technology" is broadly used to connote effective practices that have these qualities. It generally is assumed, explicitly or implicitly, that the core of a technology is defined by physical artifacts or a physical process. When we employ the term in this paper, we generally will be assuming this, although our usage can encompass any practice with a standardized, controllable, replicable core that successfully achieves what it aims to do.

These characteristics, which we propose are almost always present in effective technologies, are present in education only to a limited degree. There is no single factor or small set of factors that determines the effectiveness of a reading lesson: students with different attributes seem to learn in different ways, individual teachers are good at different things. A teacher can improve the effectiveness of the reading lesson he or she gives by learning from experience what works and what doesn't and experimenting a bit. However, these improvements in educational performance tend to be sensitive to the particular teachers employing them, to the particular student body, and to the general context within which a class or a school

is operating. Hence, what works for one teacher or school or class or student may not work for another. These features also mean that what works in a particular experimental context may provide little guidance as to what will work in a different context.[6] All of this makes progress in reading education as a whole much more difficult to achieve than progress in vaccination. We have not developed a generally effective technology for teaching reading—at least not yet.

This problem scarcely has gone unnoticed in the educational research community. There is a long tradition of trying to find or develop tightly structured teaching routines that are effective, and in recent years a number of attempts have been made to build more of teaching into computers. Invariably, however, routines that work well for some students in some settings do not seem to work well with other kinds of students or in other settings. To date, the effective use of computers in teaching subjects like reading has been limited and disappointing to its proponents. Neurological insights also have failed to lead to the development of a standard, controllable, replicable core of practice that can be reliably applied in the classroom. For a long time educationists have been frustrated by the fact that practices that yield good results in a laboratory setting seem difficult to transfer to widespread use with comparable effectiveness. There just seems to be something different about teaching and learning that makes it impossible, or at least very difficult, to discover or create an artifact or a standardized, easily replicable routine that is highly effective in most contexts.

DISCIPLINARY OBJECTIONS

Conversations with our colleagues suggest that nearly everybody is ready to call on their home discipline to explain why know-how has become strong in some areas, such as vaccination, and why it remains weak in others, such as teaching children to read. Economists tend to say that the answer obviously resides in the structure of incentives, while political scientists maintain that the problem stems from the fact that different groups have different interests, and that some have political power and others don't. Scholars in the social studies of technology argue that technologies are above all socially constructed and that the social or cultural context determines whether know-how will improve and be effective (and even how "improve" and "effective" are defined), not the other way around. Our colleagues in the natural sciences propose that that the principal reason certain problems are more difficult than others to deal with effectively is that the underlying scientific understanding is weaker. We agree that there is something to all these arguments. Indeed, we observe that in areas such as smallpox vaccination, where know-how is strong, the problem is overdetermined: more than one of these disciplinary explanations typically is satisfied. Yet in other cases, none of these explanations seems to hold—for example, in cancer prevention or treatment, where persistent attention from the marketplace and the political arena and a long-term societal commitment to advancing knowledge through research have not translated into significant progress. This suggests to us that standard disciplinary explanations are insuf-

ficient and ad hoc, and that a more general explanation must lie beyond the range of view of any traditional disciplinary lens.

Certainly one reason why know-how is weak in some areas of human need is that only limited resources have been dedicated to advancing know-how there. Some reasons for this are that market incentives may be weak and that those with the greatest need do not have the political power to induce governmental action. This is without question a good part of the reason why our ability to prevent or treat diseases found largely in poor countries generally is much weaker than know-how bearing on diseases that afflict the affluent. Conversely, it is obvious that a medical treatment that costs a lot of money or requires a sophisticated medical infrastructure for effective delivery will be more beneficial in an affluent country than a poor one. In other words, know-how may seem strong if you are rich, but not if you are poor. However, there are many areas of human need, such as improving children's reading abilities, where plentiful resources have been directed to the problem but progress nevertheless has been very limited. Thus we have made some progress in the war on cancer, but overall the scourge remains unabated, in rich countries and poor alike. Lack of market demand, political will, or economic wherewithal is only part of the explanation.

Our uneven progress in dealing with different kinds of cancer might be used by our colleagues in the natural sciences to support their argument that we make progress on solving practical problems where the underlying scientific understanding is strong and tend to be stymied where the underlying science is not strong. Clearly, the state of scientific understanding bearing on cancer is still very uneven. However, a close look at the way effective cancer treatments have been achieved shows that in many cases, probably most, deep scientific understanding had little to do with the development of effective treatments. In any case, uneven development of scientific knowledge is in itself a phenomena that needs to be better explained.[7] More generally, we note that the development of effective technologies has often preceded explanatory scientific understanding, not just in medicine[8] but in many fields of endeavor.[9]

Overall, then, disciplinary explanations, even where they do have some grip, only push the questions deeper. For example, while in the case of smallpox vaccines one did indeed see a supporting package of political will, public support, policy incentives, and scientific understanding, we will argue below that this was partly the result, rather than the cause, of the development of an effective vaccine. In contrast, our theory suggests that the absence of such a strong supporting package regarding the best method of teaching reading might quickly be resolved if a method or a technology that works well for most students could be found or developed.

Thus, whereas particular disciplinary perspectives view the achievement of strong know-how as a product of the conditions within which the know-how is applied (i.e., are the incentives appropriate? is the scientific understanding available?) and privilege their own particular insights about the world in identifying those conditions where know-how is effective, our theory suggests that the capac-

ity to develop strong and improving know-how helps to create the conditions amenable to successful application. We further note that this point, which seems to us rather self-evident, also means that the development of effective know-how helps to create the conditions that allow various disciplinary explanations to hold.

THE COMPLEX CONNECTIONS BETWEEN SCIENTIFIC KNOWLEDGE AND EFFECTIVE KNOW-HOW

Our colleagues in the natural sciences will hasten to point out that, at least in this modern era, the development of powerful, relatively standardized technologies often has been associated with the development of a strong body of scientific understanding bearing on those technologies. In recent years, the biological processes involved in vaccination have been illuminated by an impressive body of scientific understanding of immunology. It is fair to say, however, that scientific understanding provides far less illumination of how to advance educational practice, or even why particular practices seem to be moderately effective and others not.

But the connections here are not simple or straightforward. In the first place, vaccination for smallpox was discovered as an effective treatment by Jenner over 200 years ago. At that time there was no body of scientific understanding that would point to a vaccine made of puss from cowpox as an effective preventative for small pox, much less illuminate how such vaccination worked. The discovery was strictly empirical, guided primarily by folk wisdom. But since the cowpox vaccine provided the basic "go" of the treatment, once it was found to be effective on a small number of patients, the practice was subject to routinization so as to be controllable by the physician and replicable for widespread use. Relevant scientific understanding was not much stronger when, many years later, Pasteur created his effective vaccines. It was known that the diseases in question were caused by microorganisms and that for some reason vaccination with a weakened strain made patients immune to the disease; but the science of immunology came into being later. Even more recently, polio vaccines were created in the absence of fundamental knowledge about how the virus worked—knowledge that emerged only after the disease was largely eradicated.[10]

The rise of immunology as a field of science was strongly motivated by the desire to obtain a practical and deep understanding of how and why procedures like vaccination worked, as well as by a more general interest in the ways bodies fought diseases. The fact that the workings involved physical and chemical substances and mechanisms meant that the conceptions and techniques of the natural sciences would effectively illuminate paths to practical progress. The result has been a great advance in relevant understanding, which in turn has significantly improved the processes involved in vaccine development.

In contrast, scientific understanding bearing on human teaching and learning has advanced much less significantly. Some people concerned with education may propose that the difference in understanding reflects the much higher levels of

expenditure devoted to biomedical research than to research in educational psychology and other fields that might illuminate the processes of human learning and teaching. We argue, however, that scientific understanding of the sort that allows advances related to vaccines is much easier to attain and apply than science that can improve teaching. *This is true in large part because of the existence of a standardized core upon which continued improvement can be leveraged.*

Note that we are proposing that the difficulty of developing a strong scientific understanding that bears on educational practice is as much a result, as it is a cause, of the fact that the stuff of educational practice is, to a far lesser degree than immunization, standardized and controllable, replicable, and amenable to experimentation in a controlled setting with legitimate expectation that what works in that setting can be transferred effectively to actual practice in diverse settings. Absent these attributes, the methods that have been so powerful in fields like immunology, plant nutrition, and the engineering disciplines have little on which to grip. The characteristics of a practice that make it effective also would seem to be necessary if scientific research is to be effective in advancing the know-how behind that practice.

SUPPORT FOR PARTICULAR GOALS AS BOTH FACILITATOR AND CONSEQUENCE OF EFFECTIVE PRACTICE

We would like to propose that a similar pattern of two-way causation often obtains regarding the market or political support for a particular activity, and to the development of the effectiveness of that activity.

It is clear that one cannot judge whether a practice is effective or not without specifying the goals of that practice. If there is general agreement on goals and if performance in terms of these goals can be assessed relatively clearly, it is often not difficult to evaluate whether one way of doing something is better than another. When a new development that is proposed and tried has a significant effect, positive or negative, on goal achievement, there will tend to be agreement on this. In contrast, if goals are not sharp or if there is lack of agreement regarding which goals count more than others, it may be difficult to get general agreement about whether this particular practice or another is better, or whether or not a new way of doing something is an improvement.

But the relationships here are also complex. Thus there may be strong general agreement on high-level goals but significant differences in beliefs about the characteristics of practices that do and don't further the achievement of those goals. This dispute may show up in strong controversy regarding the relevance and importance of various subgoals, and such a dispute is especially likely where no practice is particularly effective. On the other hand, the fact that a practice is very effective in achieving certain objectives may result in an emphasis on what that practice does well and marginalization of its weaknesses or drawbacks.[11] The contrast here between reading education and vaccination is revealing.

There clearly is broad agreement in the United States that it is extremely

important that all children learn to read effectively. However, because teaching methods are often not effective or reliable, people can bring conflicting interests and values about methods and indicators to their discussions about what teaching practices should be implemented and how they should be implemented. We see this in the protracted debate over phonics versus "whole language" approaches to teaching reading. Thus the high-level agreement is mainly conducive to a widespread feeling that we are not doing a good enough job on this front, and that support for more research and other programs has the promise to resolve disputes and enable us to do better. But since the results of research have not been particularly impressive, there is room for continued dispute about which teaching method is best, often with the failure to recognize that, given current know-how, there may be no general answer to that question.

On the other hand, because vaccination does deliver immunity to a disease with high probability of success and this objective is widely shared, disputes about other consequences of vaccination tend to be overcome. The fact that vaccination is a frightening and painful experience for many children, that a few may suffer serious side effects, and that broad effectiveness requires aggressive policies to ensure widespread delivery can be brushed aside as not consequential, given that the vaccine does its central job so well and that the job is deemed by society to be crucial.

Shared operative values are thus in a sense drawn to highly effective bodies of know-how and practice, and they are fragile in areas where know-how is weak. In other words, bodies of know-how and practice that do certain valued things well develop a constituency of supporters who agree on both the means and the ends. Less effective, more variable routines and methods are less able to attract such a strong constituency.

This is not to say that vaccine development and deployment have not been accompanied by controversy or dispute; indeed, since Jenner's earliest efforts, opposition to vaccines has been a continual, if irregular, theme reflecting a variety of religious, philosophical, and practical concerns. At the beginning of the 20th century, proposals for mandatory smallpox vaccination provoked riots in some U.S. cities. In recent years, vaccine-related controversies have centered around alleged (but now disproved) connections between certain childhood vaccines and autism, and around the moral consequences of vaccinating teenage girls against the human papillomavirus.[12] But our key point is that the efficacy of vaccines has continued to improve over time despite such controversies and opposition, and in the face of this improvement and the resulting benefits, opposition to particular vaccines has tended to atrophy. This is a very different situation from the one we see with teaching reading in the United States, where disputes about the best approaches persist while overall efficacy does not seem to improve.

A related point is that the existence of a standardized core of reliable action allows successful adaptation within differing political and cultural contexts—that is, among differing value systems. The story of smallpox eradication in India, for example, is particularly notable for the diversity of approaches that were used to

overcome bureaucratic obstacles and cultural opposition within that large and diverse nation.[13] Teams of vaccinators adapted their approaches to local contexts, even down to the scale of working with village leaders to overcome the religious objections of specific individuals who were resisting being vaccinated. Two aspects of this example are important. First, the flexible strategies used to inoculate against smallpox all gave the same result—immunity to the disease. And second, the social feedback from the vaccination programs was both rapid and positive—smallpox outbreaks were contained. The result was cumulative progress anchored in the deployment of the core technology itself, rather than the method of deployment. Neither of these attributes is present in the case of teaching reading, where, in the absence of an effective core of practice, differences between classrooms in various socioeconomic and cultural settings continue to confound efforts to improve performance at both the local level and the cumulative national level. This lack of progress in turn prevents the convergence of different constituencies around a single approach.

THE HIERARCHY OF PRACTICE AND KNOW-HOW

Giving a vaccination or teaching a class is what one might call a "shop-floor" practice, or one performed directly to achieve a desired result. Most shop-floor practices are conducted under the aegis of an organization that has the responsibility and authority to provide whatever is needed so that those working on the shop floor can do their jobs, and, more generally, to manage the operation. Organizations at this level tend to be part of broader institutional systems. In the cases in question, vaccination and teaching reading, a doctor's office may be part of a clinic or a large private practice or an HMO, while a teacher's classroom is in a school that is part of a school district. The HMO and the school are in turn components of still larger, less definable enterprises—the "health-care system," the "education system." Much of public policy is concerned with these higher level systems.

We note that there are practices for making decisions and taking action at all levels of the hierarchy. The effectiveness of the actions taken is a function of the strength of know-how at each level. While the lines are not sharp, for our purposes here we find it convenient to distinguish between practice on the shop floor, management practice, and practice concerned with the design and overview of broad systems. We propose that as one moves up the hierarchy, it becomes more and more difficult to routinize practice in a model that achieves results broadly deemed satisfactory. As a consequence, effective higher order practice is difficult to replicate or even maintain, or to advance cumulatively. That is, know-how in general is weaker relative to management than to shop-floor practice, and weaker relative to policy-making at the systems level than to management of particular organizations.

There are several reasons for this. First, shop-floor practice is closest to the end results, and the feedback regarding whether practice is achieving the intended

objectives is sharper than for practice higher up in the hierarchy. Therefore, opportunities for productive learning from feedback on shop-floor practice are greater. Second, generally speaking, the farther up one is in the hierarchy, the more the task involves molding, coordinating, and governing activities that other individuals and organizations (at the shop-floor level) will actually do. Managing people is far more difficult to routinize effectively than operating a particular process that one has under direct control, and policy-making involves management on a very broad canvas. Third, partly for this reason but also because most organizations are concerned with a broad range of activities aimed at a variety of objectives, management and policy-making almost always require the weighing of alternative values that can be pursued to different degrees, a process that almost always is at least partially judgmental and inherently conflictual.

In the absence of effective shop-floor know-how that can teach children to read with great reliability, we must depend on better decision-making at the level of complex organizations and policy regimes in order to conquer the contextual diversity that makes it so difficult to improve nationwide reading competency. Yet at these higher levels of organizational complexity, know-how is also weak and unlikely to improve significantly.

Thus, if one considers not a particular medical practice, like vaccination against smallpox, but rather the operation of a hospital or a national health-care system as a whole, the sense of effectiveness and progress that marks medical "technology" fades, and widespread concerns and disagreements about hospital management and the performance of the health-care system more broadly come into view. The situation at these levels of medical care is not much different than that of schools, school systems, and educational policy more broadly, where there is widespread belief that we are not doing very well. The effective know-how that we have gained regarding vaccinations specifically is not matched by effective know-how regarding management or the making and administration of high-level policies.

In the absence of effective shop-floor know-how that can teach children to read with great reliability, we must depend on better decision-making at the level of complex organizations and policy regimes in order to conquer the contextual diversity that makes it so difficult to improve nationwide reading competency. Yet

at these higher levels of organizational complexity, know-how is also weak and unlikely to improve significantly. For this reason, the difficulties policy-makers and administrators have encountered in improving reading performance at the national level are not likely to soon be overcome.

Certainly there are aspects of management where effective practice has been developed, often involving highly standardized procedures. Thus, many businesses and hospitals have made good use of new computation and information technologies to establish effective systems of purchasing and inventory management. We note that these are managerial activities with objectives that are relatively clear and uncontroversial, unlike, say, hospital policies regarding whether nurses are allowed to perform procedures that medical doctors traditionally have done, or the merger and acquisition policies of businesses, or the pedagogical practices of teachers in a large and diverse urban school system. And in some arenas of activity, a large share of management can be more or less routinized. The operation of a modern aircraft carrier or nuclear reactor requires the coordinated actions of many people, and generally appropriate action patterns are achieved.[14] We note that in each of these examples there are clear, agreed-upon goals, strong chains of command, and clear, relatively instantaneous signals of failure that would adversely affect all involved (the carrier fails in its mission, the reactor explodes). But such special cases are far different from the general management problem of multigoal organizations and systems of the sort we have in health care and education.

We want, then, to emphasize that the powerful know-how embodied in the smallpox vaccine is not strongly dependent for its effectiveness on particular types of management and policy regimes. We do not mean to imply that eradicating smallpox was either easy or inevitable, only that the many organizational and policy approaches brought to bear on the eradication effort[15] could be tested, refined, adapted, assessed, and improved because of the underlying effectiveness of the vaccine. Yet this success did not reflect—or lead to—systemic, replicable improvements in know-how about the implementation of global health programs more generally.

So now we can clear up the apparent paradox that we introduced at the beginning of this essay. Certainly the organizational and policy context for global health interventions is more complex and heterogeneous than for teaching children to read in one of the world's most affluent nations. But the level of management and policy practice is not where we would expect to find powerful and improving bodies of know-how that can lead rather directly to desired outcomes. Powerful bodies of know-how are largely a "shop-floor" phenomenon, and while such know-how is clearly present in the practice of vaccinating against many infectious diseases, it remains only weakly developed in the practice of teaching reading. We emphasize that this does not mean that there are not highly skilled educators and effective approaches for teaching reading in a variety of contexts. It does mean, however, that such skill and effectiveness cannot be standardized, replicated, and delivered across the wide variety of contexts where teaching occurs.

CONCLUSIONS

We want to call attention to three consequences of this inability—except in quite special circumstances—to develop powerful standardized procedures at the level of management and policy-making. First, in an area of activity like medical practice or in a problematic policy area like protecting the environment, where we have been able to develop some strong shop-floor level know-how and practices that apparently give us strong capabilities, there tends to be a widespread feeling that we are not employing those capabilities nearly as effectively as we "could." Another way of looking at the matter is that we do not have the managerial and policy-making know-how to effectively use the technological know-how that we have. This obviously is very frustrating. But, we think that this syndrome is a fact of life, a consequence of inherent limitations of our know-how regarding how to advance practice in complex institutional settings characterized by conflicts over goals and values. Social scientists have of course devoted considerable attention to understanding the complexities of human action in such conflictual settings, but we see little evidence that such insights have led to advancing know-how that can actually be applied to yield general progress. In contrast, social science scholarship has compellingly documented the disastrous consequences that can result when limited know-how is inappropriately generalized through the imposition of overly technocratic regimes.[16]

Progress *can* be made, although it may not be dramatic, by coming to better understand the value conflicts that stymie efforts to do anything significant, by reaching some accord on the central things that should be done, and by putting in the perhaps significant resources that are needed because our know-how does not point to any low-cost or highly reliable way of dealing with the problem.

Second, because of the well-recognized limits of purposive action in complex settings, progress in our know-how regarding how to meet human needs is going to continue to be largely driven, as it has been over the last two centuries, by the development and advance of effective technologies that are operative mostly at the shop-floor level. While there certainly will continue to be advances in the technologies used to perform routinizable managerial functions like inventory control and enhancements for policy analysis like more powerful simulation techniques,

what we can expect aside from these functions regarding management and policy-making is that they adjust to allow the new shop-floor technologies to be used to further the objectives for which they are relatively effective. This is what we saw in the example of smallpox eradication in India. The eradication effort succeeded by organizing flexibly around vaccine delivery at the shop-floor level, not by creating newly effective national or global health-system bureaucracies.

Yet the lessons of successes like smallpox eradication are commonly misinterpreted and misapplied. In recognizing that social problems are far easier and less costly to deal with when we have powerful technologies that bear on them relatively directly than when we do not, society pushes hard to develop more effective practice in areas of high human need, like education, where progress has been slow. However, in many important areas, efforts to develop a technological fix[17] will not work because they lack a standardized core upon which to leverage improved know-how, or because the problem is one of purposive action at levels of high organizational complexity, or both. Thus, for similar reasons, we are not going to make major headway on these problems by learning to "manage" them more efficiently.

We do not mean to say that progress cannot be made on such problems. Progress *can* be made, although it may not be dramatic, by coming to better understand the value conflicts that stymie efforts to do anything significant, by reaching some accord on the central things that should be done, and by putting in the perhaps significant resources that are needed because our know-how does not point to any low-cost or highly reliable way of dealing with the problem. In our view, this is the way to look at the challenges of improving education, or protecting the environment, or reducing poverty. The challenge is to the democratic process, rather than to technology or management.

This latter observation may seem uncontroversial, even obvious, but it leads us to our final point. As we have emphasized, when effective shop-floor technologies do exist, and when those technologies do advance widely shared values, the previous two consequences—the disconnect between shop-floor practice and complex institutional settings, and the difficulty of advancing know-how at the management and policy levels—may be overcome. On the one hand, the smallpox eradication program required a significant degree of policy and operational coordination among diverse actors and institutions ranging from international bodies like the World Health Organization to individual nations to local governments to local political leaders and doctors. Such coordination would be unthinkable in the absence of the effective technological core and the shared values that it serves, as is starkly illustrated by the relative chaos in education policy regarding the teaching of reading that persists in the United States. We are suggesting, in other words, that effective know-how can stimulate a reorganization of the political landscape to allow for solutions to previously intractable problems not by creating new types of organizations that can be more effectively "managed," but by facilitating the application of know-how at the shop-floor level.

Thus, our theory may have important implications for choosing strategies and

allocating resources aimed at difficult societal problems. First, as we have sought to explain, even for rather focused problems such as teaching children to read, when lines of research are not leveraged by a standardized core of know-how, R&D programs should not be expected to succeed or be advertised as having any real promise of succeeding, at least in the short and medium run. It is not enough that these programs are "use inspired." They should be understood as aiming only at the creation of fundamental knowledge, and be explicitly designed and described as such. Even so, success is only possible over the long term and the probability of failure is high. Second, as many scholars and practitioners already recognize, the quest for technological fixes to address problems that are actually rooted in complex organizational settings is often not only fruitless but counterproductive. Finally, when a standardized core can be identified, remarkable progress in know-how may occur, and the return on society's investments may be very great indeed.

Acknowledgements

Our work on this topic has been supported by the Sloan Foundation. Our current understanding of the issues, as reflected in this essay, has been strongly influenced by the discussion at a meeting concerned with the uneven evolution of human know-how, held at Columbia University in November 2005. Participating in that discussion were Henry Chesbrough, William Spencer, Richard Murnane, Sidney Winter, Richard Rosenbloom, Roger Pielke Jr., Alan Moscowitz, Annetine Gelijns, Theodore Marmor, and Sherry Glied. Roger Pielke Jr., Helen Ingram, David Guston, and the Arizona State University "sustainability lunch bunch" provided useful comments on earlier drafts.

Endnotes

1. Gordon and Gordon (2003).
2. Trend data are available on the website of the U.S. Department of Education (2004).
3. See, e.g., Donovan, Wigdor, and Snow (2003), Elmore (2004), Foray and Hargreaves (2003), Murnane and Nelson (1984).
4. Murnane and Nelson (2007).
5. Gawande (2006).
6. Murnane and Nelson (2007).
7. See, e.g., Papac (2001).
8. See, e.g., Murnane and Nelson (2007), Baker (1996).
9. See, e.g., Rosenberg (1982), especially chapter 7, "How Exogenous Is Science?" pp. 141-159.
10. Baker and Katz (2004), Stern and Markel (2005), Porter (1997), Garfinkel and Sarewitz (2003), Allen (2007).
11. The complex relations and blurry lines between ends and means is central in Herbert Simon's (1945/1997) analysis of choices and decisions, first presented in Administrative Behavior.
12. Allen (2007).
13. Hopkins (1989), Bhattacharya (2006).
14. Rochlin, La Porte, and Roberts (1987), Roberts (1993).
15. Hopkins (1989).
16. Scott (1998).
17. Weinberg (1966).

References

Allen, A. (2007). *Vaccine: The controversial story of medicine's greatest lifesaver.* New York: W. W. Norton.

Baker, J. P. (1996). *The machine in the nursery: Incubator technology and the origins of newborn intensive care.* Baltimore: Johns Hopkins University Press.

Baker, J. P., & Katz, S. L. (2004). *Childhood vaccine development: An overview.* Pediatric Research, 55, 347-356.

Bhattacharya, S. (2006). *Expunging variola: The control and eradication of smallpox in India 1947-1977*. Hyderabad, India: Orient Longman.

Donovan, M. S., Wigdor, A. K., & Snow, C. E. (2003). *Strategic education research partnership.* Washington, DC: National Academies Press.

Elmore, R. F. (2004). *School reform from the inside out: Policy, practice, and performance.* Cambridge, MA: Harvard Education Press.

Foray, D., & Hargreaves, D. (2003). The production of knowledge in different sectors: A model and some hypotheses. *London Review of Education*, 1(1), 7-19.

Garfinkel, M. S., & Sarewitz, D. (2003). Parallel path: Poliovirus research in the vaccine era. *Science and Engineering Ethics*, 9, 319-338.

Gawande, A. (2006, October 5). The score: How childbirth went industrial. *The New Yorker*, pp. 59-67.

Gordon, E., & Gordon, E. (2003). *Literacy in America*. Westport, CT: Praeger.

Hopkins, J. W. (1989). *The eradication of smallpox: Organizational learning and innovation in international health.* Boulder, CO: Westview Press.

Murnane, R., & Nelson, R. (1984). Production and innovation when technologies are tacit: The case of education. *Journal of Economic Organization and Behavior*, 5, 353-373.

Murnane, R., & Nelson, R. (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.

Papac, R. (2001). The origins of cancer therapy. *Yale Journal of Biology and Medicine*, 74, 391-398.

Porter, R. (1997). *The greatest benefit to mankind: A medical history of humanity.* New York: W. W. Norton.

Roberts, K. H. (Ed.). (1993). *New challenges to understanding organizations.* New York: Macmillan.

Rochlin, G., La Porte, T., & Roberts, K. (1987). The self-designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review,* 40(4), 76-90.

Rosenberg, N. (1982). *Inside the black box*. New York: Cambridge University Press.

Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed.* New Haven, CT: Yale University Press.

Simon, H. (1997). *Administrative behavior* (4th ed.). New York: Free Press. Original work published 1945

Stern, A. M., & Markel, H. (2005). The history of vaccines and immunization: Familiar patterns, new challenges. *Health Affairs*, 24, 611-621.

U.S. Department of Education, National Center for Education Statistics. (2004). *Long-term trend: The nation's report card*. Available online at http://nces.ed.gov/nationsreportcard/ltt/results2004/

Weinberg, A. (1966). Can technology replace social engineering? *Bulletin of the Atomic Scientists*, 22(12), 4-8.