49. B. J. Meyerson, *Acta Physiol. Scand.* **63** (Suppl. 241), 1 (1964); F. P. Zemlan, I. L. Ward, W. R. Crowley, D. L. Margules, *Science* **179**, 1010 (1973); M. M. Foreman and R. L. Moss, *Horm. Behav.* **10**, 97 (1978); V. N. Luine and C. M. Paden, *Neuroendocrinology* **34**, 245 (1982); V. N. Luine and C. T. Fischette, *ibid.*, p. 237.
50. M. Warembourg, *Mol. Cell. Endocrinol.* **12**, 67 (1978); *Neurosci. Lett.* **9**, 329 (1978).
51. W. C. Young, *Sex and Internal Secretions* (Williams & Wilkins, Baltimore, 1961); H. H. Feder and B. L. Marrone, *Ann. N.Y. Acad. Sci.* **186**, 331 (1977).
52. B. S. McEwen *et al.*, in *Progesterone and Progestins*, C. W. Bardin, P. Mauvais-Jarvis, E. Milgrom, Eds. (Raven, New York, in press).
53. B. Parsons, N. J. MacLusky, L. Krey, D. W. Pfaff, B. S. McEwen, *Endocrinology* **107**, 774 (1980).
54. B. Parsons, T. C. Rainbow, D. W. Pfaff, B. S. McEwen, *Nature (London)* **292**, 58 (1981).
55. B. Parsons, B. S. McEwen, D. W. Pfaff, *Endocrinology* **110**, 613 (1982).
56. B. Parsons, T. Rainbow, D. W. Pfaff, B. S. McEwen, *ibid.*, p. 620.
57. U. Clever, *Science* **146**, 794 (1964); M. Ashburner, *Dev. Biol.* **39**, 141 (1974); ——— and H. D. Berendes, in *The Genetics and Biology of Drosophila*, M. Ashburner and T. R. F. Wright, Eds. (Academic Press, New York, 1978), pp. 315–395.
58. B. J. Meyerson, *Horm. Behav.* **3**, 1 (1972); C. Kubli-Grafias and R. E. Whalen, *ibid.* **9**, 380 (1977).
59. M. Y. McGinnis, B. Parsons, T. C. Rainbow, L. C. Krey, B. S. McEwen, *Brain Res.* **218**, 365 (1981).
60. T. C. Rainbow, P. G. Davis, B. S. McEwen, *ibid.* **194**, 548 (1980).
61. B. J. Tennent, E. R. Smith, J. M. Davidson, *Horm. Behav.* **14**, 65 (1980).
62. B. S. Rubin and R. J. Barfield, *Abstracts of the Eastern Conference on Reproductive Behavior*, New Orleans, 1979; *Endocrinology*, in press; B. S. Rubin, thesis, Rutgers University (1981).
63. J. M. Davidson and S. Levine, *J. Endocrinol.* **44**, 129 (1969); L. G. Clemens, J. Shryne, R. A. Gorski, *Physiol. Behav.* **5**, 673 (1970); R. E. Whalen, W. G. Luttge, B. B. Gorzalka, *Horm. Behav.* **2**, 83 (1971).
64. T. C. Rainbow, B. Parsons, B. S. McEwen, *Nature (London)* **300**, 648 (1982).
65. J. D. Blaustein and H. H. Feder, *Brain Res.* **177**, 489 (1979); also see J. D. Blaustein and H. H. Feder in (*14*).
66. B. Parsons and B. S. McEwen, *J. Neurosci.* **1**, 527 (1981).
67. M. Palkovits, *Brain Res.* **59**, 449 (1973).
68. T. C. Rainbow, A. Biegon, W. V. Bleisch, *Life Sci.*, **30**, 1769 (1982).
69. J. L. Gerlach and B. S. McEwen, *Science* **175**, 1133 (1972).
70. Supported by USPHS grants NS07080 (B.S.M.) and HD05751 (D.W.P.) and by an institutional grant from the Rockefeller Foundation for research in reproductive biology. We acknowledge the many contributions of laboratory colleagues: Drs. Anat Biegon, Jose Bueno, Rochelle Cohen, Paula Davis, Philip Femano Christine Fischette, Richard Harlan, Lee-Ming Kow, Lewis Krey, Ivan Lieberburg, Vicky Luine, Neil MacLusky, Marilyn McGinnis, Joan Morrell, Charles Paden, Bruce Parsons, Tom Rainbow, Edward Roy, Yasuo Sakuma, Susan Schwartz-Giblin and Brenda Shivers.

Industrial Innovation Policy: Lessons from American History

Richard R. Nelson and Richard N. Langlois

Government involvement in the research and development (R & D) process has a long history in this country. As is too often the case, this rich experience has seldom been consulted in policy debates over government programs to stimulate industrial innovation.

Summary. The historical interrelations of government support of R & D and technical change in seven major American industries point to three types of policy that have been successful in the past: (i) government R & D support for technologies in which the government has a strong and direct procurement interest; (ii) decentralized systems of government-supported research in the "generic" area between the basic and the applied; and (iii) a decentralized system of clientele-oriented support for applied R & D. A fourth type of policy, under which the government attempts to "pick winners" in commercial applied R & D, has been a clear-cut failure.

This article is an attempt to identify some of the lessons of past federal R & D policy. It summarizes the conclusions of a study, recently completed at the Center for Science and Technology Policy at New York University, of how such policies have shaped technological change in seven major American industries—semiconductors, computers, aircraft, pharmaceuticals, agriculture, residential construction, and automobiles (*1*). What makes this study unique is its explicitly comparative, cross-industry focus: it was predicated on the hypothesis—amply supported in the resulting case studies—that the kinds of government programs that have shown themselves feasible and effective vary greatly among industrial sectors, depending upon the nature of the governmental involvement and the nature of competition in the industry.

The selection of industries for study was made with an eye toward obtaining a sample with a broad spectrum of characteristics: industries with fragmented as well as with concentrated structures and industries subject to much government intervention and to relatively little. The design of the study was also informed by a desire to attract recognized scholars knowledgeable about the technology in each industry; the prior interests and areas of expertise of these scholars were therefore a factor in the selection of industries. The study was then carried out as a cooperative research effort (*2*).

The Unraveling Consensus

In treating the questions of innovation policy as warranting detailed empirical exploration, we were acknowledging, reluctantly, that the general theoretical analyses and statistical observations of economists provide only limited and incomplete guidance for policy. We are not alone in this perception; the most significant aspect of the recent economic literature on innovation is its progressive inconclusiveness about the appropriate role for government.

It was not always that way. Economic research a decade or more ago had settled on two closely related sets of propositions about industrial innovation. The first of these was that technological change is an important source of productivity growth and, simultaneously, that R & D expenditure is a principal determinant of technological advance. The implication drawn from this was that R & D spending is a kind of "control variable" through which one would affect macroeconomic productivity.

The second set of arguments derived from theoretical rather than statistical work. Economists during the 1950's and '60's developed models in which private firms possessed an inherent tendency to

The authors are, respectively, a professor of economics at Yale University, New Haven, Connecticut 06520, and a research assistant professor in the Center for Science and Technology Policy, Graduate School of Business Administration, New York University, New York 10006.

"underinvest" in R & D. There were several reasons for this unhappy circumstance: (i) R & D creates knowledge, and knowledge is a "public good" in the sense that the firm cannot fully appropriate to itself the benefits of the knowledge it creates (*3*); (ii) the payoff to R & D is uncertain, and risk-averse firms will therefore wish to do less of it than a (risk-neutral) society would prefer (*4*); and (iii) the fragmented structure of certain industries militates against sufficient R & D spending, the firms being too small to undertake certain kinds of projects (*5*).

Taken together, these sets of arguments strongly supported the notion that the government's role lay in correcting a global problem of inadequate private R & D. But over the last decade the consensus surrounding this conclusion has essentially come undone.

First, and perhaps most important, the experience of the 1970's cast doubt on the presumed tight link between a nation's overall R & D spending and its rate of productivity growth. Although it is true that high rates of R & D spending attended the rapid growth of productivity in the United States, Western Europe, and Japan during the 1950's and '60's, spending for R & D continued to be high in most countries during the ubiquitous slowdown in productivity growth after 1973. Only in the United States and France was the slowdown presaged by a decrease in R & D expenditures, and in both those countries the decrease was almost exclusively in defense and space rather than in civilian areas. Moreover, recent studies of the differences in productivity growth among countries suggests that, even in the 1950's and '60's, the countries with the highest ratio of R & D spending to gross national product—the United States and Britain—had among the lowest rates of productivity growth (*6*).

Part of the message seems to be that it matters where a country is relative to the frontier of technology and productivity. Countries not on the frontier can "play catch up" fairly easily without much R & D spending so long as their rates of physical investment are high. Countries nearer the frontier have to work harder for each percent increase in productivity. This begins to suggest that it is not necessarily to a country's great advantage to be alone on the frontier. There is evidence that, as before World War II, the United States is again benefiting from technological ideas developed elsewhere; and there may be much to recommend a world in which many countries share the technological frontier and therefore share a common economic environment in an interdependent way (*7*).

At the theoretical level, many economists have also begun to believe that the relation between competition and innovative behavior is more than a matter of some tendency to "underinvest" (or even to "overinvest") in R & D. The relationships of innovation to information "externalities," to risk, and to market structure are increasingly seen to be subtle and complex (*8, 9*).

In the first place, the implications for total R & D spending of imperfect appropriability are now understood to be less clear-cut than they once seemed. Economists are coming to realize that, in a world of patents and industrial secrecy, firms in some instances have an incentive to engage in duplicative or near-duplicative R & D in an effort to copy a rival's technology or "invent around" its patents. This at once calls into question the idea that firms necessarily engage in too little R & D; more important, it begins to focus our attention not on the level of R & D but on the types of R & D projects the industry engages in.

The economist's view of the "uncertainty" issues had taken a similar turn. Rather than focusing on the amount of R & D that uncertainty is likely to draw forth, economists are now recognizing that what uncertainty really demands is the exploration of a diverse set of approaches. This way of looking at things suggests that heavy commitments to any one approach are dangerous in the early stages of development of a technology and should be avoided until the uncertainties (both market uncertainties and technological ones) are significantly reduced.

What that suggests in turn is that the focus of the policy issue should shift from R & D levels to the portfolio of R & D projects an industry tends to generate. This entails turning our attention toward the incentive effects of policies and institutional structures and toward considerations of access, secrecy, and information flow. In practice, some kinds of R & D projects will tend to be "underfunded" and some to be "overfunded"; and a simple R & D subsidy is not the sort of policy such a situation demands.

The Case Study Approach

Although many economists are exploring new approaches to the theory of innovation and technical change, analyzing the effects of government policy on industrial innovation must still be seen as largely an empirical problem—which policies have worked, which have not, and why. This was the premise behind our seven industry case studies. One conclusion, which we develop below, is that what the government can do effectively differs from industry to industry. There are nevertheless certain features common to technological change in all industries that should be kept clearly in mind in designing government policies.

One theme that unites the history of technological change is the pervasiveness of uncertainty. Although it takes a form in (say) the pharmaceutical industry very different from that in the commercial aircraft industry, uncertainty seems nevertheless to be endemic. A quick reading of the case studies is enough to dash any supposition that technological change is somehow a cleanly plannable activity. In fact, it is an activity characterized as much by false starts, missed opportunities, and lucky breaks as by brilliant insights and clever strategic decisions. Only in hindsight does the right approach seem obvious; before the fact, it is far from clear which of a bewildering array of options will prove most fruitful or even feasible. Strange as it now seems to us, aviation experts were once divided on the relative merits of the turboprop and turbojet engines as power plants for the aircraft of the future; and the computer industry was by no means unanimous that transistors—or, later, integrated circuits—were to be the technology of the future. Policy must recognize uncertainty as a fact of life, and must not try to repress it or analyze it away.

A second and related universal theme is the importance of detailed knowledge of the technology, of its strengths and weaknesses, and of user needs in guiding the innovation process. In all the case studies, either the producer-provider or the user-demander played a major role in generating and screening technological advances. Whenever major innovation was attempted without access to their knowledge, the results were disastrous. This fact imposes severe constraints on what government can do effectively.

The implications become clear if we consider four general kinds of government support for civilian R & D: (i) programs attendant on government procurement or some other well-defined public objective; (ii) programs to support research on "generic" technologies—research in the gray area between basic scientific research and applied R & D; (iii) programs to support applied R & D in the service of well-defined clientele demands; (iv) programs that insist on "picking winners" in commercial applied R & D. Programs of each sort show up in several guises in the industry case studies.

Procurement-Related R & D

In three of the industries studied—aviation, computers, and semiconductors—the government was heavily involved as a user-demander of the technology. This kind of government involvement has two important policy implications.

The first has to do with the ability of the government to guide R & D effectively. In cases of government procurement for defense, space, or similar clearly defined public projects, the government is itself the user-demander. It thus has knowledge of its own needs and, usually, at least a modicum of expertise in the technology it proposes to use. Motivation and knowledge line up fairly well in such circumstances, and the government is frequently able to sponsor effective R & D on the relevant technology. To the extent that the technology can be easily transferred to commercial application, the result is the well-known "spillover" into civilian technology.

Second, a public belief in the legitimacy of the government's primary mission—defense, for example—smooths the political waters for any related program of government R & D support. In the semiconductor industry the Department of Defense and the National Aeronautics and Space Administration (NASA) played a crucial role. Both the transistor itself and the integrated circuit were first developed with private funds; but, in the latter case at least, military and space demand was certainly an important motivation, and once those innovations were clearly identified, government support of product and process engineering helped speed their advance. Similarly, the government funded much of the early research, and provided much of the early market, for the digital computer. Defense procurement and in recent years heavy government R & D funding have played a major role in the evolution of aircraft technology.

It is important to recognize that the efficacy of government procurement–related R & D depends on the knowledge-advantage that comes from the government's position as user and on the political legitimacy of its mission as justified on grounds other than spillover benefits. The conclusion thus does not extend to government procurement projects, the justification of which is the spillover itself or in which the procurement is intended to make a market for the technology. [The supersonic transport (SST) project remains the best case in point.]

Moreover, our case studies suggest that the potential for the generation of spillover by procurement-related government R & D support may be limited to the early stages of a technology's development, when government and civilian demands are not yet specialized. As a technology matures, the requirements of the government and the private sector normally diverge. This means not only that spillover diminishes but also that military and commercial R & D increasingly compete for resources. In the mature phases of a technology's development, spillover may be as much to the military from the commercial sector as the other way round.

Generic Technologies

When there is no recognized public-sector demand for a technology, the government's ability to fund R & D effectively and to guide the development of that technology is more limited. The government does not then have natural access to the sorts of information necessary to guide allocation, and may in fact be blocked from getting the information. Furthermore, the legitimizing effect of a public sector purpose is not there to protect a support program from strong political opposition.

Nonetheless, these problems may be attenuated if the government restricts its attentions to areas, such as so-called generic technology, that are a step or two removed from specific commercial application. The reason is that, at this "directed basic" level of research, the knowledge involved has a large public component: much of it is the sort of nonpatentable and nonspecific knowledge—broad design concepts, properties of materials, and testing concepts—that is generally shared among scientists and does not pose a strong threat to proprietary interests.

In a sense, such generic work falls in between the sorts of work that an academic researcher, pursuing fashionable questions within the bounds of a standard scientific field, would tackle and the kinds of result-oriented research that would interest most corporate R & D laboratories. Of course, some companies do support generic research, and the findings are very often treated as public rather than as proprietary. In many instances, the funding for such research comes at least in part from governmental sources. In either case, the keys to success seem to be, first, involving the relevant scientific and technical communities in the allocation process and, second, recognizing that research ought to be influenced both by the purely scientific disciplines and by those interested in applications; indeed, a tension between the pure and the applied is generally salutary.

Our case studies provide examples of generic research, some associated with government procurement (in aviation, computers, semiconductors) and some more commercially oriented work (certain aspects of agricultural and pharmaceutical science).

The agricultural sciences, viewed as a generic research system, seem to have defined and filled their niche appropriately. Such work fits in between the academic basic sciences (like chemistry and biology) and the applied R & D carried out in private firms and in the experiment stations (like the development of new seeds or fertilizers). Interests on both sides of the line pull and tug to influence the kinds of research that are done as well as to monitor its quality and efficacy. The biomedical research community is another example. Research here too is pulled by applied interests (the physicians) and tugged by scientists in the more basic fields. Interestingly, both the agricultural and the biomedical sciences typically reside in university settings, but in separate professional schools rather than main-line departments.

Another similarity between agricultural science and biomedical research is the way funding allocation is carried out. Both disciplines take the majority of their support from the government, but the funding agencies keep their distance, allowing the allocation machinery to be manipulated by the research communities themselves.

The National Advisory Committee on Aeronautics (NACA), the forerunner of NASA, is another example of a generic research system. Here the setting was a freestanding organization, not a university. But NACA's research concerns were certainly generic—broad-gauged aviation problems rather than specific designs—and the relevant engineering societies played a significant role in guiding and monitoring work at NACA. After World War II, the military increasingly assigned to private contractors the sort of work NACA had carried out, a trend that both reflected and abetted the divergence of military from civilian aircraft technology.

Because generic research poses a diffuse rather than a visible threat to established competitive positions, it may be possible to mount such a program successfully in any industry. But the size of the gray area between the basic and the applied may vary greatly among industries. Of course, the extent of this generic range may itself be influenced by the presence of a government program: the public financing of R & D often proves contagious, luring business scientists into a wider communications network and increasing the public flavor of private work. This is certainly desirable to the extent that it does not diminish the private incentive or ability to seize upon new ideas and develop them for market.

The manner in which a generic program develops may also be of critical importance. The aborted Cooperative Automotive Research Program (CARP) and the Cooperative Generic Technologies Program (COGENT) of the last administration seem to fit the description of generic systems, yet neither attracted the enthusiasm of the industries it proposed to aid, perhaps because the initiative and the design of the programs came strictly from Washington, with little participation by industry.

Clientele-Oriented Applied R & D

The most important characteristic of government support for basic and generic research is that it does not require government administrators to make decisions that involve considerations of profitability and commercial potential. A basic or generic research program seeks to advance scientific and technical knowledge, and the decisions involved require primarily scientific and technical knowledge. While acquiring this sort of knowledge is not a trivial matter, it can at least conceivably be marshaled by program administrators, especially since the relevant scientific community can often be enlisted to help guide allocation. Moreover, basic and generic research, which involves exploring widely applicable technological options, seldom poses a concentrated threat to proprietary interests.

When we move closer to the level of applied research, however, the problems of government involvement multiply. The knowledge involved is both specific and idiosyncratic in form and may be proprietary in character. This instantly puts the government administrator at an informational disadvantage vis-à-vis firms that have no incentive, and sometimes no ability, to transmit what they know to Washington. Moreover, it is difficult to maintain a political constituency for a program that poses visible threats to established competitive positions.

Now, there is one much-discussed example of a strikingly successful government program of applied R & D: the agricultural research system. As noted above, much of the research supported in this program has been basic and generic, but a sizable proportion has been extremely applied in character, focusing on particular objectives like better seed varieties or more effective pest control. The interesting question is: what special conditions have made this applied R & D program feasible and productive?

A crucial feature of the agricultural industry is that it is largely atomistic in form. The competition among farms is something near to the "perfect competition" described in economics textbooks rather than the more rivalrous kind that characterizes most manufacturing industries. For this reason, fellow competitors are seen as inherently less threatening in farming than in most other industries; technological knowledge is therefore far less proprietary, and there is a public cast to the results of even very applied R & D.

The federal-state system of agricultural experiment stations evolved in a way that took advantage of the market structure in agriculture, marshaling the support of the farmers and giving them an important position in the evaluation and selection of projects. Coupled with the regional nature of agricultural technology, this led to a system in which farmers see it as advantageous to them to advance even very specific technologies as quickly as possible. As a model for the administration of a government-supported applied R & D program, the agricultural system is quite instructive. It is highly decentralized, and specific resource allocation decisions are made at state and county levels. Those decisions respond with some sensitivity to the demands of two constituencies: farmers (given voice through state legislatures) and the agricultural science community.

In the language of the social scientist, we might call this a "captured" system, in much the same sense that transportation, communication, and other industries are said to have captured their regulators. Capture of this sort is not very often congruent with the general interest of consumers; but in the case of agriculture the system seems to have evolved in a salutary fashion.

The residential construction industry is also relatively atomistic in structure, and has therefore long been seen as conformable to the agricultural model. Yet several government efforts to spur housing R & D have not worked; more accurately, the housing industry beat back or cut back the government attempts to mount such programs.

There are probably many reasons for the nontransferability of the agricultural model to housing. Building is somewhat more rivalrous in character than farming, at least at the level of materials suppliers. More important, the atomistic home builder is very likely more conscious of a threat from housing innovation than is (or was) his atomistic counterpart in farming. Although agricultural innovation did in the long run lead to the demise of the small farmer by increasing the scale of farming operations, each small farmer could nonetheless see the benefit to him (in the short and medium run, at least) of improved farming methods. The builder may well be more aware that any exploitation of scale economies brought about by innovation would very likely redound to his disadvantage, posing a clear threat to the system of small local firms in which he operates. Another factor is that building codes—a very old form of "new social regulation"—are intractable and entrenched. And there is not the background of good basic science in housing that there is in agriculture.

Beyond that, however, it may have been crucial that the agricultural research system evolved slowly over time and was not constructed de novo or centrally designed. It may even be that its success derives critically from the particular path the system's development followed (for example, its growth out of what was essentially a training program for farmers) and particular historical circumstances (such as the characteristics of the 19th-century industry) that are not easily replicated.

Picking Winners

Which brings us to the final approach to government R & D support—"picking winners" in commercial competition. Here the historical record seems, for a change, unequivocal. Unequivocally negative.

The SST project and Operation Breakthrough were two examples touched upon in our case studies. In both cases, the government did not attempt to create a framework in which scientific and user interests could guide allocation; rather, the federal agencies attempted to insert themselves directly into the business of developing particular technologies for a commercial market in which they had little or no procurement interest.

In the case of Operation Breakthrough, the Department of Housing and Urban Development had been neither a major builder nor a buyer of nonsubsidized housing. It had neither the technical expertise nor the market experience that commercial success demands. Similarly, the government was not in the business of making or buying commercial airplanes; those who were, the commercial aircraft companies, showed little interest in such a plane until the prospects for high subsidies appeared. Very few of the housing designs that came out of Operation Breakthrough have since had any commercial value; and the lesson from the British-French Concorde experience is that we are lucky we went no further than we did with plans for an American SST.

The lesson here is not specific to these cases; it is not that these particular government agencies lacked some necessary expertise that could in principle have been remedied by hiring a larger or better cadre of experts. The lesson is a general one, about the location of knowledge and the mechanism of its transmission in the R & D system.

European experience testifies to the generality of the lesson (*10*). In many cases, government attempts to enter the business of commercial applied R & D led to (i) duplicating private efforts or (ii) subsidizing those efforts and thereby replacing private with public funds or (iii) investing in designs the private sector had long abandoned as unpromising. There is certainly an argument that the government can be more forward looking than a private firm, supporting projects that are unpromising today but may be promising tomorrow. But the most effective way to perform such a next-generation function is not by competing in the commercial marketplace but through research of a more generic sort.

Lessons for Policy

The conclusions of a comparative historical analysis can only be qualitative and judgmental. But perhaps the lesson of economic theory and political practice during the last couple of decades is precisely the importance of this sort of empirical analysis.

Our central conclusion might be summed up in one word: complexity. The wide diversity of technological and institutional details, of knowledge structures and incentive structures, among American industries recommends against an industrial policy to boost "industrial innovation" in some global sense in the hope of affecting macroeconomic problems. Broad-brush measures like tax policy, antitrust policy, and patent policy, which affect each of the various industries in a very different way, should be assessed on their own merits and should not be viewed as "control variables" to stimulate innovation.

We do not propose a general rationale or justification for active government support of R & D. Applying the lessons of history to create programs that are both politically viable and socially desirable is no straightforward task. But the historical experience we have examined reveals three approaches that have worked in the past: support associated with government procurement or some other well-defined public sector objective; support of defined nonproprietary research, with allocation funds guided by the appropriate scientific community; and provision of an institutional structure that allows potential users to guide the allocation of applied R & D funds.

A fourth kind of policy, whereby government officials themselves try to identify projects that will be winners in a commercial market competition, is always seductive, but the evidence, from our studies and others, suggests that such strategy is to be avoided.

References and Notes

1. R. R. Nelson, Ed., *Government and Technical Change: A Cross-Industry Analysis* (Pergamon, New York, 1982). The authors responsible for the industry studies are: agriculture, Robert Evenson, Yale University; automobiles, Lawrence J. White, New York University; commercial aviation, David Mowery and Nathan Rosenberg, Stanford University; computers, Barbara Katz, New York University, and Almarin Phillips, University of Pennsylvania; housing, John Quigley, University of California, Berkeley; pharmaceuticals, Henry Grabowski and John Vernon, Duke University; and semiconductors, Richard Levin, Yale University. The study was supported by grant NSG7636 from the National Aeronautics and Space Administration.
2. Each author (or pair of authors) was responsible for an individual chapter, but the group agreed on a common format for the cases, and it met several times (sometimes with other scholars and representatives from government and industry) to discuss and integrate the chapters. R. Nelson was responsible for a synthetic "cross-cutting" chapter, which was in turn discussed and commented upon by the group.
3. The most influential of these models was by Kenneth Arrow [K. Arrow, "The allocation of resources to invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, R. R. Nelson, Ed. (Princeton Univ. Press, Princeton, N.J., 1962)]. But see also H. Demsetz, "Information and efficiency: Another view," *J. Law Econ.* **12**, 1 (1969).
4. K. Arrow and R. Lind, "Uncertainty and the evaluation of public investment," *Am. Econ. Rev.* **60** (No. 3), 364 (1970).
5. The notion that "competitive"—that is, atomistic—firms innovate less than larger, "oligopolistic" firms has been called the Schumpeterian hypothesis, since those who read the work of Joseph Schumpeter through a certain set of spectacles discerned an argument of this sort in his *Capitalism, Socialism, and Democracy* (Harper, New York, 1942). The argument was stripped down to a more elemental form and popularized by John Kenneth Galbraith, especially in his *The New Industrial State* (Houghton-Mifflin, Boston, 1967).
6. Organization for Economic Cooperation and Development, *Technical Change and Economic Progress* (Organization for Economic Cooperation and Development, Paris, 1980). It may be significant that both the United States and Britain devoted an unusually high fraction of R & D to defense. But even when differences in defense R & D spending are accounted for, intercountry differentials in productivity growth seem better correlated with such matters as physical investment than with overall R & D spending.
7. N. Rosenberg, "U.S. technological leadership and foreign competition: 'De te fabula narratur?' " (photocopy, Stanford University, November 1981).
8. R. R. Nelson and S. Winter, "In search of a useful theory of innovation," *Res. Policy* **6**, 36 (1977).
9. R. R. Nelson, "Research on productivity growth and productivity differences: Dead ends and new departures," *J. Econ. Lit.* **19**, 1029 (1981).
10. See generally K. Pavitt and W. Walker, "Government policies toward industrial innovation: A review," *Res. Policy* **5** (No. 1) (1976).