1 Introduction

Zvi Griliches

In the spring of 1960 a conference was held at the University of Minnesota under the auspices of the Universities-NBER Committee for Economic Research, resulting in the publication of *The Rate and Direction of Inventive Activity* (Nelson 1962), a volume that still serves as a major statement and source book of economic ideas in this field. In the fall of 1981 the Productivity and Technical Change Studies Program of the National Bureau of Economic Research (NBER) organized a conference on R & D, Patents, and Productivity at Lenox, Massachusetts. This was the first NBER conference devoted entirely to R & D related topics since 1960. This volume which is an outgrowth of that conference, contains revised versions of papers presented at the conference plus a number of additional related papers which were distributed at the conference as background papers. Most of the latter papers report on ongoing research projects at NBER.

The major themes of research in this field were already clear at the Minnesota conference: The belief that invention and technical change are the major driving forces of economic growth; that economists have to try to understand these forces, to devise frameworks and measures which would help to comprehend them and perhaps also to affect them; that much of technical change is the product of relatively deliberate economic investment activity which has come to be labeled "research and development" and that one of the few direct reflections of this activity is the number and kind of patents granted to different firms in different years. At the earlier conference there had already been much discussion of the validity and utility of patent statistics (cf. the papers by Kuznets, Sanders, and Schmookler and the associated discussion in Nelson 1962) and

Zvi Griliches is professor of economics at Harvard University, and project director, Productivity and Technical Change, at the National Bureau of Economic Research.

attempts to relate R & D investments to their subsequent effects on the growth of total factor productivity (cf. Minasian 1962).

In the two decades that passed, the field of economics expanded tremendously with a concomitant increase in specialization in this and related fields. The Lenox conference, therefore, had a much narrower focus. While the Minnesota conference included papers by historians, social psychologists, and economic theorists, the Lenox conference concentrated primarily on applied econometric papers describing data building and data analysis efforts which focus on the role of R & D investments as generators of economic progress at the firm and industry level and on the role of patent statistics in helping to illuminate these issues.

The volume before us does not cover the whole field of R & D and innovation studies. It contains few purely methodological or taxonomic papers, and it ignores much of the interesting literature on the success and failure of individual R & D projects (such as the work of Freeman 1974; Teubal 1981; and von Hippel 1982), the work of economic historians and historians of science and technology about individual inventions or industries (cf. Nelson 1962 or Rosenberg 1976), the work of sociologists of science about the operation of big and little science (Cole and Cole 1973; Price 1963) and the construction of science indicators (Elkana et al. 1978; National Science Board 1981), and studies of how optimally to plan, organize, and evaluate R & D projects. It also connects only very loosely with the growing theoretical literature in economics on the interaction of R & D and market structure, the impact of R & D investment on the relative market positions of firms, and the impact of the market structure on the firm's incentives to innovate (see Dasgupta and Stiglitz 1980; Kamien and Schwartz 1982; and Spence 1982 for examples).

The focus of this volume is much narrower. Most of the papers deal with one of the following issues: What is the relationship of R & D investments at the firm and industry level to subsequent performance indicators such as patents, productivity, and market value? How does one formulate and estimate such relationships? What makes them vary across different contexts and time periods? To what extent can one use patent counts as indicators of R & D output? Can one detect the output of R & D in the market's valuation of the firm as a whole? What determines how much R & D is done and how many patents are received? A large number of papers in this volume grew out of several National Science Foundation (NSF) sponsored NBER projects which have been motivated in part by two data developments: (1) the increasing propensity of U.S. firms to report their R & D expenditures publicly (especially since 1972) and the availability of such data in machine readable form (Standard and Poor's Compustat tapes); and (2) the recent computerization of the U.S. Patent Office records and their public availability at reasonable cost. Much of the effort in these projects has been devoted to the acquisition, cleaning, merging, and preliminary analysis of these two data bases. The first two papers (Bound et al.; Pakes and Griliches) report the results of such preliminary analyses, as do several other papers in this volume (Ben-Zion; Griliches; Griliches and Mairesse; Mairesse and Siu; Pakes).

Bound et al. describe the construction and merging of the various data sets, present the basic statistics for the resulting samples (essentially most of the universe of large- and medium-size, publicly traded firms in U.S. manufacturing), and illustrate their potential use by presenting relatively simple analyses of the R & D and firm size and the patents and R & D relationships. They find that the elasticity of R & D expenditures with respect to firm size (measured by sales and gross plant) is close to unity with some indication of slightly higher R & D intensities for both very small and very large firms in the sample. The estimated R & D to patents relationship implies a decreasing propensity to patent with growth in the size of the R & D program throughout the range of their sample, though the conclusion is clouded somewhat by the way the issue of a large number of zero patents observations (for the smaller firms) is treated econometrically. The question of selection bias also remains in their sample. To be publicly traded, small firms must be above averagely successful, which may explain their higher patents to R & D ratios.

The Pakes-Griliches paper is based on an earlier, smaller (but longer) sample and tries to estimate something like a patent production function, focusing especially on the degree of correlation between patent applications and past R & D expenditures and on the lag structure of this relationship. Their main finding is a statistically significant relationship between R & D expenditures and patent applications. This relationship is very strong in the cross-section dimension. It is weaker but still significant in the within-firm time-series dimension. Not only do firms that spend more on R & D receive more patents, but also when a firm changes its R & D expenditures, parallel changes occur in its level of patenting. The bulk of the relationship in the within-firm dimension between R & D and patent *applications* appears to be close to contemporaneous. The lag effects are significant but relatively small and not well estimated. They interpret their estimates as implying that patents are a good indicator of differences in inventive activity across firms, but that short-term fluctuations in their numbers within firms have a large noise component in them. They also find that, except for drug firms, there has been a consistent, negative trend in the number of patents applied for and granted relative to R & D expenditures during their time period of observation (1968–75). Some of the econometric issues raised by their work have been pursued further by Hausman, Hall, and Griliches (1984) and Pakes and Griliches (1982).

Pakes and Schankerman use older European data on patent renewal fees and rates of renewal to infer the rate of depreciation in the private value of patents. This is an important question both for the computation of net rates of return to R & D and for the construction of the various R & D "capital" measures. Since the value of a patent may decline both because of the inability of its owner to maintain complete proprietary rights over the information contained in it and because of the appearance of new information which may supplant it in part or entirely, the resulting decay rate need not be similar to the deterioration rates observed or assumed for physical capital. In fact, they find it to be quite a bit higher, on the order of 25 percent per year, and argue that this may explain, in part, the conflict between relatively high estimates of *gross* private rates of return to R & D investments by Griliches and Mansfield and the rather sluggish growth in R & D during the late 1960s and 1970. The actual net rate of return may not have been all that high.

Evenson examines a variety of international data on trends in patenting, R & D, the number of scientists and engineers, and the ratio of patents granted to nationals versus foreigners by country of origin and country of patent grant. From the many interesting numbers reported in his tables, two major facts emerge: (1) Most of the foreign patents granted in all countries originate in the industrial countries. This is consistent with the hypothesis that developing countries have the advantage of being able to draw on and adopt the technologies of the developed world and that it is in their comparative advantage to do so. Moreover, the data show patenting patterns in the developing countries that are similar to the production patterns in these countries. (2) Patenting per scientist and engineer (and per R & D dollar) declined from the late 1960s to the late 1970s in almost all of the fifty countries for which data are available. Whether this indicates a significant decline in the real productivity of invention was one of the more debated topics at the conference. I shall return to this below.

In his paper, Mansfield reviews a wide-ranging research program on R & D, innovation, and technological change. The studies summarized deal with the composition of R & D, price indexes of R & D, international and "reverse" technology transfer and its impact on U.S. productivity growth, imitation costs, the role of patents and their effects on market structure, and other topics. The first study reviewed indicates that the relationship of productivity increases in an industry to R & D depends on the extent to which it is long-term and basic. Larger firms (but not in more concentrated industries) did devote more of their research funding to basic research but did less than their share of risky and long-term R & D projects. In the second study, data were collected and price indexes of R & D were constructed for a number of industries for 1979 (1969 = 100), showing that the price of R & D rose by more than the GNP deflator during this period. The third study dealt with the effects of federally financed R & D on private R & D investments. The findings could be interpreted as indicating that federally supported R & D expenditures substituted for about 3 to 20 percent of private expenditures and induced an additional 12 to 25 percent increase in private R & D investments. While the direct returns from federally financed R & D projects may be lower, the projects do seem to expand the opportunities faced by firms and induce additional R & D investments by them. Several studies of international technology flows showed significant "reverse flows." Overseas R & D by U.S. corporations resulted (in over 40 percent of the cases examined) in technologies that were transferred "back" to the United States. Moreover, for a small sample of chemical and petroleum firms, the estimated impact of overseas R & D investments on their own overall productivity is quite large and statistically significant. The last set of studies reviewed by Mansfield found that imitation costs were close to two-thirds of the original innovation costs; that they were higher for patented innovations but not by much (except in the drug industry); that patented innovations were imitated surprisingly often (60 percent within four years of their initial introduction); that less than one-fourth of the patented innovations outside the drug industry would not have been introduced without patent protection; and that, in some major industries, concentration-decreasing innovations are a very substantial proportion of all product innovations.

Beggs examines historical data on patenting for twenty selected industries and thirteen census years between 1850 and 1939. By looking at the relationship of movements in relative patenting to movements in relative value added across industries, he arrives at conclusions opposite to those of Schmookler: industries that do relatively better can afford to and do slacken off their inventive efforts. (This interpretation is problematic and is inconsistent with the Griliches-Schmookler 1963 results which also can be interpreted as "relative" comparisons.) Beggs also examines the long annual series on total patents issued from 1790 to date with the help of spectral analysis techniques and finds a five-year cycle and a secondary eight-year cycle, which he interprets as several rounds of "follow-up" inventions which follow major technological breakthroughs. His results are very interesting and suggestive but not very firm both because of the high level of aggregation in this analysis and because they do not appear to be reproducible on the shorter but possibly more relevant series on patents by date of application (rather than issue).

The next set of three papers deals with different aspects of the interrelationship between R & D and the market structure of industries that firms find themselves in and have an influence on. Levin and Reiss present a detailed industry equilibrium model where concentration, R & D intensity, and advertising intensity are all jointly determined by the profit-maximizing actions of individual firms. While emphasis is put on both differences in technological opportunity across industries and differences in the conditions of appropriability of the results of technological innovations, very little directly relevant data are available on these topics. Levin and Reisss use a variety of proxies and estimate their model using simultaneous equations techniques and data for twenty manufacturing industries and three years (1963, 1967, and 1972). They get a strong positive effect of R & D on industry concentration and a negative effect of concentration on R & D intensity which becomes positive for industries with a high share of product rather than process R & D. They also find that government-supported R & D does increase private R & D investments, though the estimated effect is not very large (on the order of .1). The main difficulty with their findings (discussed also in the Tandon comment) is that while the model is fundamentally dynamic, it is estimated on industry *levels* without leaning more heavily on the changes that occurred in these industries over time.

The second Pakes-Schankerman paper also deals with the determinants of R & D intensity, primarily at the firm level but with some attention to the issue of aggregating to the industry level. They set up a model in which the optimal (from a firm's point of view) R & D level depends on the expected market size for the firm's products, the conditions of appropriability of the firm's technological innovations, and the technological opportunities facing the firm. The latter two, which are the determinants of market structure and R & D spillovers and the driving forces in the Levin-Reiss model, are largely subsumed here in an unobservable parameter which varies across firms. They reexamine some of the earlier firm R & D data reported in Grilliches (1980) and conclude that, though the coefficient of variation of research intensity is much larger than that of more traditional factors of production, very little of the observed differences in R & D intensity across firms can be explained by either past or even expected rates of growth in sales or by transitory fluctuations (errors) in the various variables. They attribute the remaining differences to interfirm variance in appropriability environments and technological opportunities. An interesting empirical fact emerges in the second part of their paper, when they expand their data base and add industry sales growth rates to their analysis. These turn out to be more important in explaining firm differences in R & D intensity than the firm's own sales growth rates, which helps explain the empirical fact that at the industry level of aggregation the variance in growth rates does account for much of the variance in R & D intensity. One interpretation of the finding, which they emphasize and which is also consistent with the Levin-Reiss results, is that industry differences in technological and market opportunities predominate in a firm's decision processes. The industry potential matters much more than the firm's own specific past history.

Scott uses the newly collected FTC line of business level data to investigate several interesting hypotheses below the firm level. Given the fact that many of the major R & D performing firms in the United States are large, diversified, and conglomerate, it is interesting to ask: Is their R & D behavior primarily determined by the industrial location of their "lines of business" (division or establishment) or does a common "company" R & D policy exist? Without an affirmative answer to the last part of this question there would be grave doubts about the applicability of various R & D optimizing models which relate to such firmwide variables as the cost of capital or their managerial style. Luckily Scott does provide an affirmative answer. In his data (473 companies, 259 different four-digit level FTC lines of business, and a total N of 3387) he can observe the variation in the R & D to sales ratio (R/S) within firms across their various lines of business. He finds that approximately half of the overall variance in R/S can be accounted for by common company effects, common industry effects, and their interaction, in roughly equal parts. Thus, there appear to be significant differences in company R & D policy above and beyond what would have been predicted just from their differential location within the industrial spectrum. Scott also finds that government-supported R & D encourages company-financed R & D. The effect here is statistically significant but small (on the order of .1). Unfortunately, no clear behavioral model is developed here (or for Mansfield's similar finding) to explain these results.

The next five papers (Griliches, Pakes, Abel, Mairesse and Siu, and Ben-Zion) are connected by their use of the stock market value of the firm either as an indicator of the success of R & D programs or as a measure of expectations and a driving force of subsequent R & D investments. In a brief note (that is republished here), Griliches sets the stage for some of this research by using the market value of a firm as an indicator of the market's valuation of the firm's intangible capital, especially its R & D program and accumulated patent experience. Using combined cross-section time-series data, he finds significant effects for both R & D and patent variables in a market value equation. These effects persist when both individual firm constants and lags in market response are allowed for, but the interpretation of the coefficients become obscure. In particular, holding previous market values and current dividends constant, there should be no additional effects from past or current R & D expenditures on subsequent market valuation unless something unanticipated happens. That is, only the news component in the R & D and patent series should affect changes in market value, and this is, in fact, what some of the results seem to imply.

In a theoretically and econometrically much more ambitious study, only whose abstract is included in this volume, Pakes tries to unravel such effects using modern time-series analysis methods. He uses the reduced form of an intertemporal stochastic optimizing model to interpret the time-series relationships between patents, R & D, and the stock market rate of return. In this interpretation, events occur which affect the market value of a firm's R & D program and what one estimates are the reduced form relationships between percentage increases in this value and current and subsequent changes in the firm's R & D expenditures, its patent applications, and the market rate of return on its stock. His empirical results indicate that though most of the variance in the stock market rate of return is noise (in the sense that it is not related to either the firm's R & D expenditures or its patent applications) there is still a significant correlation between movements in the stock market rate of return and unpredictable changes in both patents and R & D expenditures (changes which could not be predicted from past values of patents and R & D). Moreover, the parameter values indicate that these changes in patents and R & D are associated with large movements in stock market values. The R & D expenditure series appear to be almost error free in this context. Patents, however, contain a significant noise component (a component whose variance is not related to either the R & D or the stock market rate of return series). This noise component accounts for only a small fraction of the large differences in the number of patent applications of different firms, but plays a much larger role among the smaller fluctuations that occur in the patent applications of a given firm over time. The timing of the response of patents and R & D to events which change the value of a firm's R & D effort is quite similar. One gets the impression from the estimates that such events cause a chain reaction, inducing an increase in R & D expenditures far into the future, and that firms patent around the links of this chain almost as quickly as they are completed, resulting in a rather close relationship between R & D expenditures and the number of patents applied for. Perhaps surprisingly, he finds no evidence that independent changes in the number of patents applied for (independent of current and earlier R & D expenditures) produce significant effects on subsequent market valuations of the firm. The data cannot differentiate between different kinds of events that change a firm's R & D level. If his model were expanded by adding sales and investment data, it may prove possible to differentiate between pure technological shocks and demand-induced shifts in the R & D and patenting variables. Even without this distinction and without a precise structural interpretation of the estimated relationships, the current model does yield a useful description of the relationships between the various variables and their timing and a very suggestive interpretation of them.

In his note, which developed out of his comment on the Pakes paper at the conference, Abel shows how to construct a structural model in which there is an explicit connection between the market value of the firm and its current and past expenditures on R & D. This model, which is rather simple and primarily illustrative of how one might go about constructing such structural models, yields two interesting conclusions: (1) It is possible, in his one-capital world, to write the value of the firm as a linear function of the stock of R & D capital (which provides some comfort to the empirical approach outlined in the Griliches note). (2) The value of the firm depends directly on demand shocks (the innovations in the output price process) and their square and so does R & D activity (and presumably also on technological shocks which are not contained in his model), but R & D *expenditures* are a function only of the square of these shocks, of the variance in the output price process which measures the degree of uncertainty in the firm's environment. The last result may help to explain the relatively low observed correlation between the two (R & D expenditures and changes in market value).

Mairesse and Siu analyze the time-series interrelationships between changes in the market value of the firm, sales, R & D, and physical investment using what they call the extended accelerator model. This paper follows the Pakes paper both in approach and in the use of essentially the same data. It differs by not focusing on patents, adding instead sales and investment to the list of series whose interrelationship is to be examined. Two additional differences between these papers should be noted: (1) Pakes defines q (the stock market rate of return) as of the year preceding the R & D expenditures, while Mairesse and Siu define it as concurrent. (2) Parkes relates q to the (logarithmic) level of R & D and patents, while Mairesse and Siu use first differences of the logarithms of levels for their R & D, investment, and sales variables. They find that a relatively simple "causal" model fits their data: "innovations" in both market value and sales "cause" subsequent R & D and investment changes without further feedback from R & D or investment to either the stock market rate of return or sales. There is little evidence of a strong feedback relationship between physical and R & D investment, though there is some evidence of contemporaneous interaction. An interesting conclusion of their paper is that independent changes in sales explain a significant fraction of the changes in R & D (and physical investment) above and beyond what is already explained by changes in the market value of the firm and by lagged movements in R & D itself, implying that using different variables one might be able to separate out the effects of different kinds of shocks in the R & D process. This finding could, of course, be just a reflection of a substantial noise (error) level in the observed fluctuations of the stock market rate of return pointed out earlier.

Ben-Zion examines the cross-sectional determinants of market value, following an approach similar to that outlined in the Griliches note. It differs by not allowing for specific firm constants (except in the last table) and by including other variables, such as earnings and physical investment, in the same equation. He also finds that R & D and patents are significant in explaining the variability of market value (relative to the book value of its assets), in addition to such other variables as earnings. His most interesting finding, from our point of view, is the relative importance of total patents taken out in the industry on the firm's own market value. In his interpretation, patents applied for indicate new technological opportunities in the industry, and these overall opportunities may be more important than a firm's own recent accomplishments, though here again this could arise just from the high error rate in the firm's own patent counts. In the last table, Ben-Zion comes close to reproducing the market value change or stock market rate of return equations in the Griliches, Pakes, and Mairesse-Siu papers. He finds that "unexpected" changes in R & D affect market value significantly, in addition to the major impact of unexpected changes in the firm's earnings.

This set of papers clearly opens up an interesting research area but still leaves many issues unresolved. Like the proverbial research on the characteristics of an elephant, different papers approach this topic from slightly different points of view. Pakes analyzes movements in patents, R & D, and market value; Mairesse and Siu investigate the relationship between R & D, investment, sales, and market value; while Ben-Zion (in his change regressions) looks at R & D, earnings, and market value. We should be able to do a more inclusive analysis in the future, incorporating the various variables in one overall model (or at least description) of these processes.

The Schankerman and Nadiri paper is motivated by the availability of R & D investment anticipations in the McGraw-Hill Surveys data base. It sets up an optimal R & D investment model and derives the equations of motion for actual R & D, anticipated R & D, and their difference—the realization error. Given the anticipations data, the paper shows a way to formulate cost of adjustment models which permit the testing of various expectational schemes: rational versus adaptive versus static expectations. The results are somewhat inconclusive. The pure rational expectations model is rejected by the data. As is true in many such endeavors, it is not clear whether the specific hypothesis is rejected or whether the model is failing for other reasons (errors in variables or wrong functional form). The adaptive expectations version fares better. The main driving variables in their model are current and past sales and the price index of R & D. Since there are no data on specific firm prices, an aggregate and rather smooth price index is used as an approximation, resulting, unfortunately, in mostly nonsignificant estimates of the various price coefficients. The paper shows how difficult it is to formulate a rigorous theory of R & D investment and to derive explicit functional forms for its estimation. It is an important first step on a rather long road with at least three major tasks still ahead of us: (1) developing a rigorous and effective model of two types of investments with two capital stocks (R & D and physical); (2)

treating output (or sales) as endogenous, being planned simultaneously with the R & D program (after all, it is the R & D that is supposed to generate new sales in the future); and (3) finding more relevant proxies for the "price of R & D" than different versions of aggregate wage and price indexes.

The last set of papers considers the impact of R & D on total factor productivity. The first three consider the issue primarily at the micro firm level, while the last two analyze the R & D-productivity relationship at the more aggregated industrial level. The relationship of R & D to productivity growth is of interest for at least two reasons: (1) Productivity growth is a major source of overall economic growth and the understanding of its sources has been a major research goal of economists over the last three decades. (2) Many of the effects of R & D may be social rather than private, in the sense that they are not appropriated by the unit that produced the particular R & D results and may not show up in its profits or output measures. To the extent that they are of benefit to other firms and industries, they may show up in the more aggregate industry and economywide productivity numbers. Unfortunately, the productivity measures themselves are subject to much error and may not reflect many of the technological changes that are not ultimately embodied in easily measurable products. For example, the gains in space exploration and medical research do not show up directly in the productivity figures as currently constructed, nor do most of the gains in complex commodities, such as consumer electronics, for which no good price indexes are available to use in the construction of real product measures. (See Griliches 1979 for further discussion of some of these issues.)

The Griliches and Mairesse paper uses a production function framework to analyze the impact of past cumulated R & D expenditures on the output (deflated sales) of over a hundred large U.S. firms, covering a twelve-year time period (1966–77). They find that there is a strong relationship between firm productivity and the level of its past R & D investments in the cross-sectional dimension, but that in the within-firm time-series dimension of the data, this relationship almost vanishes. This may be the result of a higher degree of collinearity between the time trend which is used as a proxy for more general outside sources of technical change and the growth in physical and R & D capital stocks in the within-firm dimension, and to the greater importance of measurement errors and other transitory fluctuations in these data. When they constrain the other coefficients to reasonable values, the R & D coefficients are sizeable and significant even in the within-firm dimension of the data. They also develop a simultaneous equations interpretation of their model and estimate what they call "semireduced form" equations, which again yield rather high estimates of the contribution of R & D capital to productivity growth relative to that of physical capital. Two important topics are raised but not pursued very far: (1) In these data some of the return to R & D appears to come in the form of mergers. (2) The conventional model used in this type of research does not allow for a noncompetitive firm environment. They do reinterpret their numbers by an individual firm monopoly model but do not pursue its implications for estimation.

In their companion paper, Cuneo and Mairesse reran the Griliches-Mairesse models on similarly constructed French data with largely similar results. While their sample covers a shorter time period, they had access to better data and this seems to matter. In particular, they show that having data on materials use (or value added) does reduce the difference between the estimated coefficients in the total and the within-firm dimension. They can also take out the R & D related components of labor, materials, and capital from the conventional input measures and thereby avoid the usual double counting problem: treating R & D expenditures as a separate variable while the actual inputs bought with these expenditures are already counted once in standard measures of labor, etc. They show that this type of double counting biases the estimated R & D coefficients downward, more so in the total than in the within-firm dimension. They also confirm a not so surprising fact already discovered by Griliches (1980) and Griliches-Mairesse: the R & D intensive industries, the "scientific" industries, are also the industries with higher estimated R & D coefficients. They do more R & D and therefore, on the margin, have rates of return similar to those found in the other less R & D intensive industries. Thus there is evidence for differences in R & D elasticities (across industries) and for the view that R & D is pushed far enough in the more "scientific" industries to come close to equalizing the private rates of return to R & D across industries.

This assumption of equalization of private rates of return (rather than equality of elasticities) across firms is the motivation behind the approach taken by Clark and Griliches in their productivity study based on the PIMS (profit impact of market strategies) data base. They use a measure of R & D intensity (R/S) instead of the change in the R & D capital stock as their basic variable, postulating that private gross rates of return to R & D are more nearly equalized across firms in different industries than would be the case for production function elasticities. Their paper is also interesting for the use of data on "businesses" below the level of a company. These data are based on concepts that are similar to the FTC "line of business" classification used in the Scott and Scherer papers. This is a much more relevant level for data analysis when companies are large and conglomerate. They also had access to data on the composition of R & D expenditures (product vs. process R & D), on the importance of patents and other proprietary processes for the firm, and on the prevalence of technological change in the firm's particular line of business. Their main finding is a statistically significant relationship between R & D intensity and the growth in total factor productivity, implying a gross excess rate of return to R & D of about 20 percent. This return is bigger for process R & D than for product R & D. The effects of the latter are not all caught, presumably, by the firm's own deflated sales measures but are passed on in some part to the other product-using industries. They find no significant decline in the productivity of R & D during the 1970s, but they do find that these returns depend crucially on the presence of previous major technological changes in the respective industries, implying (without being able to measure it) a major role for spillovers from the previous R & D efforts of other firms and industries.

This is also the topic that motivates Scherer's important contribution. His paper describes in detail a major and valuable data construction effort whose basic purpose was to reallocate R & D expenditures from an industrial "origin" classification (where they are done) to a classification of ultimate "use" (where they will have their major productivity-enhancing impact). This was accomplished by examining over 15,000 patents in detail and assigning them to both industrial origin and industrial use categories and categorizing them into product and process patent categories. The detailed R & D by line of business data collected by the FTC were then reallocated from industry of origin to industries of use in proportion to the "use" distribution of their patents, thereby generating a kind of technological flow table. The many conceptual and practical difficulties in such an enterprise are discussed by Scherer in some detail. The appendix to his paper presents the most detailed data on R & D by three- and four-digit Standard Industrial Classification (SIC), by origin, and by use ever made available. These data will prove invaluable in future studies of productivity growth and differential industry R & D activity. Scherer reports briefly on an analysis of productivity growth in which, once the quality of the output growth data is controlled for, the newly generated R & D by industry of use data prove superior to the industry of origin data in the explanation of interindustry productivity growth differences.

In the final paper in this volume, Griliches and Lichtenberg examine the relationship between R & D and productivity growth in U.S. manufacturing industries at the two- and three-digit SIC levels. They use the NSF classification of R & D by product group rather than by industry of origin to approximate better the ultimate industrial location of the effects of these R & D expenditures and the Census–Penn–Stanford Research Institute (SRI) data base to construct detailed total factor productivity indexes. They look at the question of a possible secular decline in the fecundity of R & D and find no evidence for it. While there has been an overall decline in productivity growth, including R & D intensive industries, the statistical relationship between productivity growth and R & D intensity did not disappear. If anything, it grew stronger in the 1970s.

The conference concluded with a discussion session in which different speakers expressed their perception of the state of research in this field. That discussion and the discussions following the presentations of the various papers (parts of which are reproduced in this volume) ranged over a variety of topics with the following receiving the most attention: (1) the ambiguities of the patent data; (2) the aggregation level at which the R & D process should be studied: project, establishment, firm, industry, or economywide; (3) the absence of data on what really drives R & D—the changing state of technological and market opportunities; (4) the low quality and the dubious relevance of the available productivity data and the absence of alternative indicators of social returns to R & D; and (5) the difficulties in and the importance of modeling the spillover of knowledge and technology from one firm or industry to another.

The critics emphasized the fact that patents differ greatly in their economic significance and play different roles in different industries. The authors of some of the papers found it difficult to reconcile their basic reliance on the law of large numbers with their desire to analyze microdata. Even though the meaning of any individual patent may be highly variable, one hopes that large differences in the number of patents applied for across firms or over time do convey relevant information about the underlying trends and fluctuations in inventive output. This reasoning is not very helpful, however, when applied to data on individual firms, many of whom take out only an occasional patent or two, especially if patent counts are to be used as an independent variable, helping to explain some further measure of the consequences of inventive activity. It is doubtful whether small fluctuations in patent counts convey much information. Even though they should have known better, the authors were surprised, I think, by the large amount of randomness and by the low fits that they encountered in trying to analyze such data.

Nevertheless, the main conclusion that did emerge (though not unanimously) was that something is there, something worth working on and analyzing. Patents and patent counts are, after all, one of the few direct quantitative glimpses into the innovation process available to us. The studies do show a strong cross-sectional relationship between R & D and patents and a weaker, but still statistically significant one between their fluctuations over time. Thus, to a first approximation, one can use patent data as indicators of technological activity in parallel with or in lieu of R & D data. This is of significant practical import since in many contexts detailed patent data are more readily accessible than R & D data.

The work reported to date has yet to establish that there is net information added in patent counts, that patents as a measure of output of the

R & D process provide superior explanatory power in modeling productivity change or other performance indicators. Some scattered results implied an independent contribution of patents to the explanation of differences in the market value of firms (above and beyond what was already accounted for by R & D variables), but no study had connected it yet to productivity growth. Perhaps the greatest promise of the patent data is in the level of detail contained in them and in the potential for using this detail to reclassify and illuminate other data. Scherer's paper is a prime example of such work where the information contained in the patent documents was used to reclassify and reallocate R & D expenditures into more relevant industrial boundaries. Another use discussed at the conference and currently being pursued at NBER is to study the overlap in the patenting of different firms across different patent classes in an effort to develop measures of technological similarity or distance between pairs and groups of firms. The notion here would be to use such distance measures in the analysis of spillovers of R & D effects between firms and industries, assuming that they are more prevalent at shorter "distances," and to derive better, technologically more homogeneous, industry groupings for the various firm samples.

There was much debate about the appropriate level and detail for the study of the R & D process. Many of the interesting questions and decisions are taken at the "project" level but little data are available on this level. Nor is it clear how generalizable some of the cases and smaller survey studies really are. It would be valuable, however, as is illustrated by Mansfield's work reported in this volume, to have greater detail on the composition of R & D itself. It matters to the analysis whether most of R & D is spent on basic and long-term research or almost entirely on adaptive and short-term research. Most econometric research uses, however, the available data, even if they are not entirely relevant. Clearly we would like to have better and different data, but it is my belief that we have not yet digested and understood much of what is already available. Even so, the currently available data have already produced some interesting findings, confirming, for example, the relationship between R & D and productivity growth and indicating no apparent deterioration in it recently.

One of the issues discussed at the conference was the apparent worldwide decline in patenting in the 1970s. This emerged from the negative trend coefficients in the Pakes-Griliches paper, a perusal of Evenson's tables, and also a look at the raw overall data (as summarized, for example, in *Science Indicators 1980* [National Science Board 1981]. Since the resources of the various patent offices are largely fixed and since worldwide cross-patenting has increased, some of the apparent decline may be an artifact of the "crowding out" of applications, but patent office data on applications filed (not just granted) do suggest that there may have been a real decline in the 1970s in patents applied for by corporations. Patent applications by all corporations (which were ultimately granted) peaked around 1969, roughly coincidentally with peaks in real R & D expenditures in industry and in the employment of scientists and engineers. Since 1969 the level of corporate patenting has somewhat declined on the order of 10 percent. Whether this should be interpreted as reflecting an exhaustion of technological opportunities is doubtful. It is more likely a reflection of the deteriorating macro conditions around the world and of a possible decline in the value of patenting, because of rising costs of litigation to keep them in force and faster rates of disclosure and subsequent imitation following patent granting and publication, because of improved communication systems and more continuous surveillance of the patents being granted by computer accessible literature search services. Unfortunately, detailed data on applications filed but not granted are inaccessible, and the currently available patent data sets do not go back far enough (before 1965) to allow us to distinguish a trend from a cycle.

Several papers set up models of the determinants of R & D investment and others talk about the necessity of considering the R & D decision within some wider, multi equational, simultaneous framework. Unfortunately, the standard variables that are brought in to "explain" movements in R & D and the other factors of production do not appear to be all that relevant to the R & D story. First, little good data are available on the "price" of R & D, but even if we had them, they would move largely in parallel with the major cost component of R & D—the cost of labor. Second, the price of capital story is likely to be similar for physical and R & D capital except for some differences in the tax treatment of depreciation, the effects of which one is unlikely to be able to detect well in the kind of data available to us. Thus, I am somewhat pessimistic about the promise of approaches which can be caricatured as defining and treating R & D as just another "capital stock" and reducing the analysis to the previous case. This misses whatever it is that makes R & D a different endeavor from just buying another plant or a new set of machines.

Unfortunately, when one starts thinking about what is special about R & D—the importance of technological opportunities, scientific know-how levels, and expectations about eventual market size for particular products—it is difficult to see how these characteristics can be quantified and forced into the Procrustean bed of our standard models. The most that one can do at the moment is to model them as unobservable "shocks," along the lines of Pakes's paper and the earlier "unobservables" literature, and trace out their effects on and interaction with the other variables of interest, such as patents, physical investment, and market value. While it may prove possible to distinguish between demand and supply (technological opportunity) shocks in such models and provide insight into and an interpretation of the interdependence between these variables, this line of research is unlikely to lead to models with clear policy handles.

The productivity slowdown and possible reasons for it created much interest at the conference, but there were expressions of pessimism about our ability to detect the major effects of R & D in such data. Several problems are evident: (1) the poor quality of output price indexes in the R & D intensive industries, such as computers and electronic components; (2) the long and variable lags in the impact of particular technological developments on subsequent productivity growth; (3) the diffuse nature of such effects and the arbitrariness of many of the industrial boundaries in our data; and (4) the absence of good measures of real product in some of the final demand sectors with important R & D effects such as health, defense, and space exploration. In spite of the uncertainties introduced by measurement issues and the fact that different total factor productivity measures (by different researchers) do not agree closely when it comes to an examination of individual industry trends, there was general concurrence in the notion that the R & D slowdown is not implicated directly in the sharp and worldwide productivity slowdown which started in 1974–75. Less clear-cut evidence and less agreement are found for the milder and longer term total factor productivity growth slowdown which may have started in the late 1960s. Here the slowdown in R & D growth may have been a contributing factor. Whether it reflects an exhaustion of technological opportunities is not clear, but it is likely to contribute to a slower growth rate in the underlying potential of the world economy in the future.

From a methodological point of view, in spite of all the talk about technology flows and externalities, the main unsolved research problem is still how to handle the interdependence between projects, firms, and industries. Almost all our methodology is based on the individual unit of observation, be it firm or industry. We do have models which allow for a one- or two-dimensional interdependence (such as serial correlation or variance components), but we have little experience and skill in modeling the clustering of and interaction between a relatively large number of actors. Even the work based on transaction flow tables does not do the job since it is fundamentally unidirectional. Our existing methodological tools cannot handle this and predispose us to ignore such problems. I believe that the development of methods for the analysis of large group interactions will be the major task and challenge during the next decade of research in this area.

What have we learned since the last NBER R & D conference more than twenty years ago? What were the substantive findings reported at this conference? I do believe that we have made progress in understanding the questions, in developing more rigorous models and better tools of analysis, and in accumulating and analyzing much larger data sets. We have described and documented a significant relationship between R & D investments and subsequent productivity growth. This relationship remains even in these trying times, though shrouded by data and measurement uncertainties. We have also concluded that federal R & D expenditures do not have much of a direct effect on productivity as it is conventionally measured but do stimulate private R & D spending and may thereby have a nonnegligible indirect effect. We are more aware of both the conceptual and the measurement difficulties involved in productivity measurement and less sure about the relevance of the existing measures to the issues at hand. We do have evidence that there may be something interesting in the patent data after all. They do appear to be useful indicators of innovative activity (though less so in the small and over short time periods), and there may be fruitful ways of using them in further analysis. We have also learned that the relationship between R & D, firm size, concentration, and all the rest is much looser and more obscure than is implied by the usual statements of the Schumpeterian hypothesis. While much of the R & D effect is concentrated in large firms, it is more likely that they have become large because of their R & D successes rather than that they do more and more fruitful R & D because they are large.

However, we have not provided, except indirectly, many policy handles. Nor is it likely that we will do so in the future. This is not because we do not want to be helpful to the National Science Foundation or the rest of the policymaking establishment, but because what we are studying is not really amenable to short-run policy intervention or manipulation. R & D investment and performance are largely determined by the evolution of scientific opportunities in a field and by peoples perceptions and expectations of the future economic climate within which new products or processes are to be sold or used. These can only be affected indirectly and imperfectly by supporting science in general and basic research in particular and by pursuing wise macroeconomic policies. All else, I believe, is of secondary importance.

References

Cole, J. R., and S. Cole. 1973. *Social stratification in science.* Chicago: University of Chicago Press.

Dasgupta, P., and J. Stiglitz. 1980. Industrial structure and the nature of innovative activity. *Economic Journal* 90:266–93

Elkana, Y., J. Lederberg, R. K. Merton, A. Thackray, and H. Zuckerman, eds. 1978. *Toward a metric of science.* New York: Wiley.

Freeman, C. 1974. *The economics of industrial innovation.* Baltimore: Penguin.

Griliches, Z. 1979. Issues in assessing the contribution of research and development to producivity growth. *Bell Journal of Economics* 10, no. 1 (Spring): 92–116

———. 1980. Returns to research and development expenditures in the private sector. In *New developments in productivity and analysis measurement and analysis*, ed. J. W. Kendrick and B. N. Vaccara, 419–54. Conference on Research in Income and Wealth: Studies in Income and Wealth, vol. 44. Chicago: University of Chicago Press for the National Bureau of Economic Research.

Griliches, Z., and J. Schmookler. 1963. Inventing and maximizing. *American Economic Review*, September 53(4):725–29

Hausman, J., B. H. Hall, and Z. Griliches. Econometric models for count data with an application to the patents R & D relationship. 1984. *Econometrica* (in press).

Kamien, M. I., and N. L. Schwartz. 1982. *Market structure and innovation.* New York: Cambridge University Press.

Minasian, J. P. 1962. The economics of research and development. In *The rate and direction of inventive activity*, ed. R. R. Nelson, 93–142. Princeton: Princeton University Press for the National Bureau of Economic Research.

National Science Board. 1981. *Science indicators 1980.* Washington, D.C.: National Science Foundation.

Nelson, R. R., ed. 1962. *The rate and direction of inventive activity: Economic and social factors.* Universities-NBER Conference Series no. 13. Princeton: Princeton University Press for the National Bureau of Economic Research.

Pakes, A., and Z. Griliches. 1982. Estimating distributed lags in short panels. NBER Working Paper no. 933. Cambridge, Mass.: National Bureau of Economic Research.

Price, D. J. 1963. *Little science, big science.* New York: Columbia University Press.

Rosenberg, N. 1976. *Perspectives on technology.* New York: Columbia University Press.

Spence, A. M. 1982. Cost reduction, competition and industry performance. Harvard Institute of Economic Research Discussion Paper no. 897, Cambridge, Mass.

Teubal, M. 1981. The R & D performance through time of young, high technology firms. Maurice Falk Institute for Economic Research in Israel Discussion Paper 814, Jerusalem.

von Hippel, E. 1982. Appropriability of innovation benefit as a predictor of the source of innovation. *Research Policy* 11(2):95–115.