GCR 99-780

PB2000-104377

Estimating Social and Private Returns from Innovations Based on the Advanced Technology Program: Problems and Opportunities

Prepared by

Edwin Mansfield

University of Pennsylvania

for

U.S. DEPARTMENT OF COMMERCE Technology Administration The Advanced Technology Program National Institute of Standards and Technology Gaithersburg, MD 20899

REPRODUCED BY: NTIS. U.S. Department of Commerce National Technical Information Service Springfield, Virginia 22161



National Institute of Standards and Technology Technology Administration U.S. Department of Commerce

GCR 99-780

Estimating Social and Private Returns from Innovations Based on the Advanced Technology Program: Problems and Opportunities

Prepared by

Edwin Mansfield University of Pennsylvania

for

U.S. DEPARTMENT OF COMMERCE Technology Administration The Advanced Technology Program National Institute of Standards and Technology Gaithersburg, MD 20899

January 1996

PROTECTED UNDER INTERNATIONAL COPYRIGHT ALL RIGHTS RESERVED NATIONAL TECHNICAL INFORMATION SERVICE U.S. DEPARTMENT OF COMMERCE





U.S. DEPARTMENT OF COMMERCE Michael Kantor, Secretary

TECHNOLOGY ADMINISTRATION Mary L. Good, Under Secretary for Technology

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY Arati Prabhakar, Director

PREFACE

Most who are knowledgeable about the economic evaluation of industrial innovation are familiar with Professor Edwin Mansfield's 1977 landmark set of case studies which measured the private and social returns on investment in 17 industrial research projects. Oft cited and followed as a model for subsequent case studies, Professor Mansfield's work continues to represent an important body of economic case studies of industrial innovation.

The Advanced Technology Program (ATP), the U.S.'s civilian technology program begun in 1990, has developed and is implementing an ambitious program to evaluate the impacts of its investments in high-risk, enabling technologies proposed, led, cost shared, and carried out by U.S. companies. One of the components of ATP's comprehensive evaluation program is to conduct case studies of the projects it funds. In keeping with ATP's approach of involving the nation's leading economists in its evaluation studies, it has engaged the services of Professor Mansfield to direct a series of case studies of ATP-funded projects.

Preparatory to undertaking the ATP-project case studies, Professor Mansfield carried out a preliminary exploration of a sample of ATP-funded projects to lay the groundwork for the case studies. His purpose was to learn more about the ATP and the projects it has funded; to assess the willingness of company award recipients to cooperate directly with him; to determine the availability of data; and to consider the applicability of his previous case-study methodology to ATP case studies. This report is the result of Professor Mansfield's preliminary investigation.

It is written in an informal style, and presumes some familiarity of the reader with Professor Mansfield's previous work and a working knowledge of private and social rate-of-return methodology. We welcome your comments.

i

Rosalie Ruegg Chief, Economic Assessment Office Advanced Technology Program

January 1996

·

ii

TABLE OF CONTENTS

PREFACE i
EXECUTIVE SUMMARY v
1. Introduction
2. Industry's Estimation of Private Returns from New Technology 2 Scoring systems 2 Programming techniques 3 Decision analysis 4
3. Benefits of Early Integration of R and D and Commercialization Plans
4. Errors in Forecasting Development Cost, Time, and Profits 8 Development Cost and Time 9 Probability of Technical Completion 11 Discounted Profits from New Processes and Products 12
5. Estimation of Social Returns from New Technology
6. Further Complications and Adjustments 21
7. Other Types of Innovations
8. Follow-On Studies by Foster and Nathan
9. The Relationship between Social and Private Returns
10. Forecasting Social Rates of Return
11. Models of the Diffusion Process
12. Extending and Implementing the Model
13. Data Availability and Research Strategy
References
Endnotes

·

.

EXECUTIVE SUMMARY

- 1. Even the most sophisticated firms have found it very difficult to forecast the private returns from R and D projects. In part, this is because it is so hard to predict how effectively the results of such a project will be exploited. In addition, because of the uncertainties inherent in R and D (as well as the inevitable biases due to organizational pressures and the perceived need to sell projects), the forecasts of development cost and time, probabilities of success, and profits seem to contain substantial errors and biases (at least based on the available, but old, data). Choosing projects which are close to commercialization, or where the technology is being commercialized, for study will improve the quality of the forecast.
- 2. Firms that attempt forecasts of the economic potential of the R and D at an early stage of the R and D process, integrating marketing, production, and finance variables, tend to be more successful with using forecasting in their decision making than others.
- 3. The model used by my students and myself, and by Foster Associates and Nathan Associates in their replication and extension of our work, seems to be a promising framework for estimating the social (and private) returns from the investments in new technology made by ATP and its award recipients. However, while this simple framework has been flexible and general enough to be useful for dozens of innovations, modifications and extensions may be needed to handle the large variety of ATP-funded projects and the knowledge spillovers that ATP expects to result from many of the technologies it funds.
- 4. To illustrate how the social benefits from an industrial innovation can be derived, suppose that the innovation is a new product used by firms, and that it reduces the costs of the industry using the innovation. The social benefits from the innovation can be measured by the profits of the innovator from the innovation plus the benefits to consumers -- that is, consumers of the good produced by the industry using the innovation -- due to whatever reduction occurs in the price of this good due to the innovation. (Because a variety of adjustments have to be made, the calculations are not this simple, but this conveys the spirit of the analysis.) To the extent that the innovation is adopted (or adapted) in other applications, a similar approach would be taken in each application and the total social benefits (less costs) would be aggregated.

- 5. Based on a sample of 16 ATP award recipients, it appears that most of them would be willing to cooperate by providing descriptive material, analytical assistance and data.
- 6. NIST should view the estimation of social (and private) rates of return as a dynamic, continuing process which initially results in admittedly crude estimates but which gradually homes in on reasonably accurate findings. It is my view that this could result in a wealth of useful information. Unless NIST can somehow obtain much more reliable forecasts of many of the inputs into this model (or any other such model) than were available in the 1970s, it seems doubtful that the initial estimates based on forecasted data at the beginnings of programs and projects will be very accurate. But with updating as commercialization and diffusion occurs, valuable information can be obtained concerning social rates of return, as well as the size of forecast errors and how one can devise and use early estimates in a civilian technology program like ATP.
- 7. The effectiveness and credibility of the results will depend very heavily on the reliability and unbiasedness of the basic data, which can be obtained only from firms. There are advantages in having at least some of the work done by non-government personnel who collect data on a confidential basis from firms. This may reduce fears that firms may feel pressure (even if subtle, unintentional, or imagined) to paint a rosy picture to increase chances of future funding and for other reasons.

Estimating Social and Private Returns from Innovations Based on the Advanced Technology Program: Problems and Opportunities

Edwin Mansfield University of Pennsylvania

1. Introduction

The purpose of this paper, commissioned by the U.S. Department of Commerce, is to bring together the results of previous research of my own and others, to supplement these results with information obtained from a sample of ATP award recipients, and to discuss in a preliminary way some methods that the managers of the Advanced Technology Program might use to estimate and forecast the social and private returns from the R and D projects they support. This paper is meant to be a preliminary, small-scale attempt to put the forecasting problem in perspective and to get some idea of what can reasonably be expected in this regard. It is not meant to come up with detailed recommendations regarding which forecasts of this kind for particular projects. Instead, it is an effort to determine what can be learned from experience in leading business firms and government agencies, as well as from academic studies, regarding this problem.

In 1977, my students and I published both a paper in the <u>Quarterly Journal of</u> <u>Economics</u> and a book on the estimation of social and private rates of return from industrial innovations. The models and techniques that were used have been applied frequently, and seem to have proved helpful to policy makers and analysts (Mansfield et al, 1977 a,b). This paper considers how these models and techniques could be adapted to help ATP with the forecasting and estimation problems cited above. As part of this discussion we also look at the sorts of data that might be collected by ATP and/or its contractors in order to make subsequent forecasts (and eventually retrospective estimates) of the social and private rates of return from the R and D projects supported by ATP.¹

2. Industry's Estimation of Private Returns from New Technology

To begin with, let's consider the estimation of private returns from new technology. While this is not the primary interest of the ATP program, it is the area where the most attention (by far) has been devoted. For decades, leading high-tech firms have been faced with the problem of estimating the returns from their R and D (and related) investments. While some firms seem to have felt that retrospective studies of the returns from such investments have been useful, the available studies indicate that few firms, if any, are confident of their forecasts of the profitability of particular R and D projects.

<u>Scoring systems</u> At General Electric, for example, a simple scoring system has been used to estimate the relative profitability of various R and D projects.² Two values have been calculated, the first being

$M \times S \times G \times T$,

where M is the estimated size of market (for a new product), S is probable GE share, G is estimated rate of growth of market, and T is a measure of its sensitivity to technological advance. Whereas the first value pertains to the impact of the program if it succeeds, the second value pertains to the probability of success. Specifically, this second value is

D x B x F x O,

where D is a measure of the difficulty of the technical problem, B is a measure of the competitive status of the lab effort, F is a measure of the fit with the laboratory's resources, and O is a measure of the ease of transition to operations. These two values have been used to try to determine which R and D projects should be carried out (Steele, 1988).

Programming techniques Given the crudeness of such scoring methods, it is not surprising that firms have experimented with more sophisticated techniques like linear programming. To illustrate how linear programming can be used, suppose that a firm has a list of n possible R and D projects that it might carry out and that to undertake project I would cost C_i dollars. Moreover, suppose that project I has a probability of success of P_i and that, if successful, it will result in a discounted profit (net of commercialization costs but gross of R and D costs) of π_i . Then, if the firm can spend no more than C dollars on R and D and if it wants to maximize the expected value of profit, its problem can be represented as follows:

Maximize $\sum_{i=1}^{n} X_i (P_i \pi_i - C_i)$ subject to the constraint that

$$\sum_{i=1}^{n} X_i C_i \le C$$

3

and

 $X_{i} = 0, 1.$

In other words, the firm's problem is to choose the X_i -- where $X_i = 1$ if project I is accepted and 0 if it is rejected -- in such a way that it maximizes the expected value of profit, subject to the constraint that the total amount spent on R and D does not exceed C.

Using programming techniques, this problem can be solved. However, this does not mean that techniques of this sort are widely used. One reason why programming techniques have not found more extensive use is that data concerning P_{i} , π_{i} , and C_{i} are often very rough. Much more will be said on this score in Section 4. For now, the important point is that surveys indicate that relatively sophisticated quantitative techniques generally have not been used by firms to select R and D projects. Instead, relatively simple (and crude) techniques generally are applied (Steele, 1988; Mansfield et al, 1971).

Decision analysis Besides using quantitative techniques to try to select R and D projects, firms sometimes look at the prospective returns from a sample of the potential innovations emerging from their R and D to provide a rough indication of the potential value of their research. For example, a leading oil firm has carried out a detailed series of decision analyses for this purpose.³ Each year, a number of potential innovations was chosen by the managers of the firm's central research laboratory. These projects had to be at a sufficiently

advanced stage so that reasonable estimates concerning prospective costs and benefits could be made. Basically, the present value of the cash flow from the innovation was estimated, given each of a set of possible circumstances. Probabilities were attached to each of these circumstances, and the expected value of the discounted cash flow was calculated for the innovation. The computation of discounted cash flow began with the present year and included R and D expenditures through the pilot-plant stage. (Benefits are discussed below.) Deflated dollars were used throughout each analysis. A discount rate of 5 percent was employed because, according to the firm's personnel, this "corresponds to [the firm's] expected rate of return on its investments, when measured in the context of constant (i.e., deflated) dollars."⁴

According to an executive of this firm who was intimately involved with these decision analyses, many projects displayed a similar logical structure: "The initial step in the analysis is the identification of the qualitative benefits to be provided by the innovation. Usually this is a straightforward issue; for example, the project may involve a process improvement, the production of a more valuable product, or the establishment of a completely new business area. Concomitant with this step is a clarification of the 'next best' existing alternative to the new technology. By definition, this alternative is the one that [the firm] would practice in the absence of the innovation. While this initial step may seem obvious, it is critical to stress the importance of clearly identifying the alternate case so that a suitable reference case can be established and that the appropriate incremental benefits can be taken into account."

5

More specifically, according to this executive, there generally have been several models that have been part of the analysis of each innovation. "At the beginning is the 'Technical Model' in which the key scientific and technological issues are enumerated. Included are topics which need to be quantified before the project can proceed to a pilotplant or precommercial stage. Often these issues are phrased as knockouts, that is, critical path items that must be resolved favorably for the project to remain viable. Other issues may be characterized as 'figures of merit' and are associated with parameters whose numerical values determine the actual quantitative benefit derived from the innovation."

"The 'Economic Model' follows and includes a detailed computation of the net savings provided by the new technology. This model requires the parameters from the preceding technical model. Often an algorithm is created that calculates the incremental benefits relative to a reference case based on existing technology. The computational guidelines for cost estimation are consistent with [this firm's] standard engineering practice. In fact, direct involvement by engineering personnel in this phase of the analysis is usually required. The output from the model consists of generalized expressions for the incremental capital and operative savings per unit application of the innovation."

"The 'Market Model' is undoubtedly the most difficult model to construct. This deals with how the innovation will affect [this firm's] future business activities over the period of the technology's lifetime. Typically, the development of the model begins with categorizing the types of application areas and identifying the corresponding receiving organizations (e.g., affiliates). An effort is then made to obtain broad inputs ... to construct a scenario representative of an implementation schedule. The model may ultimately predict how many refineries would utilize a new process technology as a function of time. Alternatively, the model could predict a market penetration scenario for a new business example using cost projections and assumptions regarding future demand and competitive technologies."

3. Benefits of Early Integration of R and D and Commercialization Plans

Even the most sophisticated techniques aimed at forecasting the returns from R and D projects and programs seem to have had limited success. For example, the oil firm in the previous section abandoned the model described above. One of the key problems has been the difficulty in forecasting how R and D, if successful, will be utilized. Whereas economists' models often put R and D at the center of the stage, the truth is that R and D is often useless unless it is combined properly with marketing, production, and finance. Thus, to forecast the returns from investments in new technology, one must forecast how the new technology will be used, which seldom is easy.

Nonetheless, the available evidence (which unfortunately dates back to the 1970s) seems to suggest that firms that attempt such forecasts at an early stage do better than others.⁵ The probability that an R and D project will be an economic success is the product of three probabilities -- the probability of technical success, the probability of commercialization (given technical success), and the probability of economic success (given commercialization). According to econometric results based on data from 20 firms in the chemical, drug, petroleum, and electronics industries, all three of these probabilities of

7

success were directly related to how quickly R and D projects have been evaluated from the point of view of economic potential.⁶ Apparently, this reflects the fact that firms sometimes allow projects to run on too long before evaluating their economic (as distinct from their technological) potential. Even though it is very difficult to forecast the benefits and costs of a particular R and D project, firms seem to do better if they try to do so early in the game, rather than wait until substantial amounts have been spent.

To illustrate the significance of organizational factors, the probability of commercialization (given technical completion) seems to have been directly related to the degree to which R and D and marketing have been integrated. In some firms, the R and D department has not always worked very closely with the marketing staff, the result being that the R and D output has not been as well mated with market realities as it might have been. A substantial percentage of a firm's R and D results may lie fallow because other parts of the firm do not make proper use of them. According to estimates made by executives of 18 of these firms, the percentage of R and D projects that were economic successes would have increased by about one-half if the marketing and production people had done a proper job in exploiting them.⁷ (And it is important to note that the non-R and D executives seemed to agree on this point with the R and D executives.)

4. Errors in Forecasting Development Cost, Time, and Profits

Forecasts of the private rate of return from an investment in new technology depend on, among other things, how much will be spent on development and how long development will take, as well as on the probability of technical success, commercialization, and economic success. Further, it depends on the profitability of the new products and processes, if any, that stem from the investment.

Probably the most important reason why sophisticated models have proved so difficult to apply is that firms have not been able to make reasonably reliable forecasts of development cost and time, or the probability of success, or the profitability of new products or processes. During the 1960s and 1970s, my students and I collected a substantial amount of detailed data regarding the size of errors in forecasting these items. Since these data seem to be all that are currently available, I have no choice but to use them, although the need for updating them is clear. However, for what it is worth, a small sample of R and D executives (that I interviewed in connection with this paper) felt that the situation in this regard had not changed very much.⁸ Indeed, many of them thought that such forecasts tend to be poorer now than in the 1960s and 1970s, because down-sizing at many R and D laboratories has lowered the capability to do such forecasting.

Development Cost and Time

To illustrate our findings regarding development cost and time, consider the results of an early study in the drug industry (Mansfield et al, 1971). Very detailed data were obtained for a major ethical drug firm and a major proprietary drug firm concerning the errors in the cost and time estimates made at the beginning of drug-development projects. For over 80 percent of the projects in the ethical drug firm, the actual cost and time exceeded the estimated values. The average ratio of actual to estimated cost was 1.78; the average ratio of actual to estimated time was 1.61. Cost and time estimates were less reliable for new chemical entities than for compounded products and alternate dosage forms. In the proprietary drug firm, the average ratio of actual to estimated cost was 2.11, and the average ratio of actual to estimated time was 2.95. Again, the overruns were greater for more ambitious projects.

When we compared the overruns in these two drug firms with those in weapons development, we found that the cost overruns for new drug products were less than those in weapons development and that the time overruns were greater than in weapons development. However, it is important to note that the cost overruns in the drug firms began to approximate those for military projects when entirely new types of projects or larger technical advances were attempted. For example, the average ratio of actual to expected cost was 2.25 for new chemical entities in the ethical drug firms, 2.75 for new products in the proprietary drug firm, and 3.2 for a sample of airplane and missile projects. Turning to time overruns, we found the average ratio of actual to expected time was 1.89 for new chemical entities, 3.24 for new products in the proprietary drug firms, and 1.4 for the airplanes and missiles.

In the ethical drug firm, we tested various hypotheses concerning the effects of various factors on the size of a project's cost overrun. In accord with these hypotheses, it turned out that technically more ambitious projects tend to have greater cost overruns than technically less ambitious projects. Also, products with wider spectra of activity tended to have larger cost overruns than single-market products, and projects with small estimated

costs or longer duration tended to have larger cost overruns than projects with large estimated costs or shorter duration. In the proprietary drug firm, there was also a significant tendency for technically more ambitious projects and projects with smaller estimated costs and longer duration to have larger cost overruns.

When the same kind of model was used to analyze development time, the results were rather similar to those for development cost. In the ethical drug firm, there was a significant tendency for products with wider spectra of activity and projects with smaller estimated lengths to have greater time overruns. In the proprietary drug firm, there was also a significant tendency for projects with small estimated lengths to have greater time overruns; moreover, there was a nearly-significant tendency for new products to have larger time overruns than product improvements.

Probability of Technical Completion

Let's turn now to estimates of the probability of technical completion of R and D projects. To illustrate our findings, consider the proprietary drug laboratory cited in the previous section, where records of many completed projects included an estimate of the probability of achieving the technical objectives, as stated in the project proposal. This estimate was made at the time of formal project proposal. Data for 79 completed projects indicated that the estimated probability of technical completion was <u>on the average</u>, a very good indicator of actual outcome, the average estimated probability of technical completion (0.81) being very close to the actual proportion of projects that were completed (0.76).

However, the fact that the average estimated probability of technical completion was close to the actual proportion of projects that were completed does not mean that the estimates were useful in predicting which projects were more likely to be completed. In fact, the estimated probabilities of technical completion were of some use in predicting which projects would be completed and which ones would not. But they were not of much use. Even if they were employed in such a way that the probability of an incorrect prediction was minimized, they predicted incorrectly in about 30 percent of the cases. (One would have expected to have made incorrect predictions in only 36 percent of the cases by chance.) We also compared the actual proportion of projects that were technically completed with the average estimated probability of technical completion for projects attempting small, medium, and large technical advances. For those attempting small technical advances, the estimated probability of completion, on the average, overstated the risk of failure. On the other hand, for those attempting medium or large technical advances, the estimated probability of completion, on the average, understated the risk of failure.

Discounted Profits from New Processes and Products

To complete this brief survey of our findings during the 1960s and 1970s regarding forecasting errors, let's turn to the accuracy of forecasts of discounted profits from new processes and products. Consider one of the largest firms in the country.⁹ During the 1960s and 1970s, this firm made a careful inventory of the major new products and processes it developed each year and forecasted the discounted profits from each such technological development. Moreover, in each subsequent year, it revised these forecasts in the light of new information. For example, after making its initial estimate in year 2 of the discounted profits from each process or product developed in year 1, revised estimates of this sort were made in year 3, year 4, year 5, and so on. Because these estimates were systematically and carefully updated, they provided a relatively unique opportunity to study how quickly forecasts of this sort converge on their true value. Unfortunately, because of reductions in available resources, this analytical effort was reduced by the firm during the 1980s, and terminated in the early 1990s.

Based on an early study of this firm's experience (Beardsley and Mansfield, 1978), there frequently have been rather significant revisions of the profit forecasts during the first five years after a new product or process has been developed, but, as one would expect, these revisions have become more minor as time goes on (and more and more of the uncertainties have been resolved). By nine years after the development process ends, it appears, in this firm at least, that a reasonably definitive estimate can be made of the discounted profits from the new technology. In very few cases were any significant revisions made in this profit estimate in the tenth to thirteenth year after termination of development. Thus, we could safely use the estimate of a new product's or new process's discounted profits made nine years after its development as an adequate approximation to its actual discounted profits.

To see how rapidly the forecasting errors diminish as time goes on (after the development of a new product or process), we divided the forecast made one year later, two years later, and so on, for each product or process by its actual discounted profits (i.e., the

forecast nine years after development). Then, as a simple measure of the size of the forecasting errors, we calculated the proportion of cases where this ratio was greater than or equal to 2.0 or less than or equal to 0.5. Figure 1 shows the decrease that occurred in this measure of forecasting error as one revision after another was made in the forecasts. The initial forecasts -- those made one year after development -- were generally quite poor, the proportion in Figure 1 being 0.50 for processes and 0.62 for products. During the first four years after the initial forecast, the size of the forecasting error, as measured by this proportion, decreased at a relatively constant rate. By five years after development, this proportion was 0.06 for processes and 0.15 for products.

Figure 1 shows that the initial profitability estimates for individual product and process innovations developed by this firm were not very accurate. To see whether the firm's initial estimates improved over time, we categorized the 57 new processes and products by the year they were developed, and calculated the frequency distribution of the ratio of initially forecasted to actual discounted profits in each such category. The results indicated that there was no tendency (at least during this period) for the initial estimates to improve over time.

To see how closely correlated the initial forecasts were with the actual discounted profits from each of the new processes or new products and to estimate the relationship between them, we regressed the initial forecasts on the actual discounted profits. For new products there was a surprisingly low correlation between the initial forecasts and the actual discounted profits, \overline{R}^2 being only .14. In the case of new processes, the correlation was much

higher, \overline{R}^2 being about .87. Thus, the initial forecasts for processes tended to relate more closely to the actual outcomes than those for products. Both for processes and products, the initial forecasts tended to be relatively optimistic in cases where actual profits were small and relatively pessimistic in cases where actual profits were large. Specifically, the forecasts underestimated the profitability of new processes that had discounted profits exceeding about \$3 million and of new products where they exceeded about \$1 million, and over-estimated the profitability of less profitable new products and processes.

5. Estimation of Social Returns from New Technology

Having looked briefly at the problems firms have encountered in trying to forecast the private returns from their investments in new technology, let's turn to the estimation of the social returns from such investments. Economists have tended to shy away from forecasting in this area; they generally have been content to do retrospective studies. To illustrate, consider the early study that my students and I carried out (Mansfield et al, 1977 a,b). This was the first such study of industrial innovations. Valuable previous work had been done in agriculture (see Evenson, Waggoner, and Ruttan (1979) and particularly Griliches (1958)), but it was necessary to extend this earlier work to handle many basic features of innovation in industry, such as the pricing policies of innovators, and to recognize the effect on displaced products, and the costs of un-commercialized R and D and R and D done outside the innovating organization, as well as environmental effects. Having done so, it was possible for the first time to make comparisons of social and private rates of return.

Figure 1. - Proportion of new processes or new products in which the ratio of forecasted discounted profits to actual discounted profits was greater than (or equal to) 2.0 or less than (or equal to) 0.5, 57 new products and processes, 1-9 years after product or process development



16

The innovations in our sample can be divided into three classes: product innovations used by firms, product innovations used by households, and process innovations. Based on an intensive study of each of the innovations in each of these classes, it appeared that the same general kind of model was applicable to all of the innovations in our sample in a particular class. This section describes the model that was used to measure the social benefits in a particular period from product innovations used by firms.

Each of these new products resulted in a potential saving to users. For example, the product innovation in the primary metals industry resulted in a potential saving to makers of household appliances. Thus, each of these innovations could shift downward the supply curve of the industry using the innovation. How far downward this supply curve will shift depends, of course, on the pricing policy of the innovator. If the innovator charges a relatively high price for the new product, the supply curve may shift only slightly. Indeed, if the innovator charges a high enough price, the supply curve will not shift downward at all.

Assume that the innovator decides to set a price for its new product which yields a profit¹⁰ to the innovator equivalent to r dollars per unit of output of the industry using the innovation (for example, r dollars per appliance in the case of the new type of metal). Also, assume that the industry using the innovation is competitive and that its supply curve is horizontal in the relevant range. In particular, assume that, before the advent of the innovation, this supply curve was S₁ in Figure 2, and the price charged by the industry using the innovation was P₁. After the advent of the innovation, this supply curve is S₂, and the price is P₂.

Under these circumstances, the social benefits from the innovation can be measured by the sum of the two shaded areas in Figure 2. The top shaded area is the consumer surplus due to the lower price (P_2 rather than P) resulting from the use of the innovation. In addition, there is a resource saving, and a corresponding increase in output elsewhere in the economy, due to the fact that the resource costs of producing the good using the innovation are less than P_2Q_2 . Instead, they are P_2Q_2 minus the profits of the innovator from the innovation, the latter being merely a transfer from the producers of the good using the innovation to the innovator. Thus, besides the consumer surplus arising from the price reduction, there is a resource saving amounting to the profits of the innovator.

To illustrate, suppose that the innovator responsible for the new product used by the appliance industry reaps a \$100 million annual profit from this innovation. This means that P_2Q_2 is an over-estimate of the value of the resources used by the appliance industry. It is too big by \$100 million, the amount the appliance industry pays the innovator in profits, because this payment is not in exchange for resources: it is a transfer of profit to the innovator.

Two adjustments must frequently be made in the estimate corresponding to the lower shaded area in Figure 2. First, if the innovation replaces another product, the resource saving cited in the paragraph before last does not equal the profits of the innovator (from the innovation), but these profits less those that would have been made (by the innovator and/or other firms) if the innovation had not occurred and the displaced product had been used instead. This is the correct measure of the resource saving. Second, if other firms imitate the innovator and begin selling the innovation to the industry that uses it, their profits from the sale of the innovation must be added to those of the innovator to get a full measure of the extent of the resource saving due to the innovation.

Using this model, an estimate was made of the social benefit in each period from the investment in each of these innovations. For each innovation, the top shaded area in Figure 2 equals

$$(\mathbf{P}_1 - \mathbf{P}_2) Q_2 (1 - \frac{1}{2} Kn), \tag{1}$$

where $K = (P_1 - P_2)/P_2$ and n is the price elasticity of demand (in absolute value) of the product of the industry using the innovation. To estimate $P_1 - P_2$, we obtained as much information as we could concerning the size of the unit cost reduction due to the innovation in the industry using the innovation. Based on interviews with executives of the innovating firm, executives of a sample of firms using the innovation, and reports and studies made by these firms for internal purposes, it was possible to obtain reasonably reliable estimates of $(P_1 - P_2)$. And once we had an estimate of $(P_1 - P_2)$, it was simple to compute K. Also Q was generally available from published records. Rough estimates of n were obtained from published studies and from the firms. Since K was generally very small, the results were not sensitive to errors in n. Indeed, the expression in equation (1) could be approximated quite well in most cases by $(P_1 - P_2)Q_2$, which is the total savings to consumers due to the lower price, if they buy Q_2 units of the product of the industry using the innovation.





To estimate the additional resource saving from the innovation, which equals the bottom shaded area in Figure 2 (if the adjustments described in the paragraph before the last are unnecessary), the innovator's profit from the new product was obtained from detailed discussions with the firm's executives, as well as study of relevant financial records. For each year, the costs of marketing and producing the innovation, as well as the costs of carrying out the innovation (R and D, plant and equipment, manufacturing start-up, and marketing start-up), were deducted from the innovator's revenues from the innovation. Note that the R and D costs were adjusted to allow for the fact that the innovator invested R and D resources in uncommercialized R and D projects. To make this adjustment, we obtained estimates from each of the innovating firms of the average number of dollars spent on uncommercialized R and D projects per dollar spent on a commercialized R and D project during the relevant period. Then we multiplied the innovator's R and D outlays (in each year) on the innovation by this number in order to get an estimate of the total R and D investment, including a pro-rated allowance for uncommercialized projects. In cases where the adjustments described in the paragraph before the last were necessary, estimates of the foregone profits from displaced products were deducted, and the profits of imitators were added, to the innovator's profits. These estimates were obtained from the relevant firms.

6. Further Complications and Adjustments

We also had to recognize that besides the innovator, other firms or research organizations may have invested resources (prior to the introduction of the innovation in question) in R and D and related innovative activities aimed at innovations of essentially the same kind as this one. Clearly, it is not easy to obtain data on the extent of such investments, but fortunately the difficulties seemed to be less than they might appear, for two reasons. First, only a limited number of organizations could reasonably have been expected to have been doing R and D in the relevant area, and if they had been devoting any substantial amount of resources to such work in the relevant time frame, it is inconceivable that the current executives of these organizations (and all their competitors) would have been unaware of it. Second, since these innovations occurred some time before our study, firms generally were quite willing to discuss whether they were carrying out work of the relevant kind at that time. Moreover, and this is particularly fortunate, our results are quite insensitive to errors in the estimated investment in R and D carried out by others. Even if the true social research expenditures were ten times our estimate, and the true social development expenditures were double our estimate, the results change remarkably little. (To a considerable extent, this is due to the fact that R and D costs are often much less than the total cost of launching an innovation.)

Interviews with executives of the innovating firm as well as of other firms that could reasonably be expected to do (and be aware of) R and D of the relevant kind, indicated that, in most cases in our sample, no other firm or research organization was doing work aimed at roughly the same kind of innovation. Thus, in these cases, the private investment seemed to be a good approximation of the social investment. In the remaining cases, other firms or individuals were engaged in R and D aimed at the same kind of innovation. In the cases where this R and D was *unsuccessful*, we obtained as accurate an estimate as possible of the cost of this unsuccessful R and D, and added this figure to the R and D costs described in the previous section to get an estimate of the total social cost of the relevant R and D carried out by the innovator and others.

In the single case (among product innovations used by firms) where other firms or individuals were engaged (prior to the introduction of the innovation in question) in R and D aimed at roughly the same kind of innovation, and where this R and D was successful, it was necessary to recognize that the innovator's investment only resulted in the innovation's availability at an earlier point in time, not in all of the social benefits from the innovation up to the relevant time horizon. In other words, the proper comparison was between what would have occurred if the innovator had not carried out the innovation (but other firms were free to do so) and what in fact occurred. In this case, we obtained as accurate an appraisal as we could of the date when the innovation would have appeared if the innovator had not carried out the innovation, and we calculated the social benefits only during the period between the date when the innovation occurred and the date when it would have appeared if the innovator had not carried it out. But in this particular case a realistic (if perhaps somewhat conservative) estimate is that the second firm to produce the innovation would have come up with it when in fact it did so, regardless of whether the innovator preceded it or not.

For most innovations, these calculations were carried out for each year from the beginning of work on the innovation until 1973. Thus, our estimates of the social benefits were conservative, since all benefits after 1973 were ignored. But in some cases, this would have introduced a serious distortion, since the innovation was relatively new. In these cases, forecasts were made of the consumer surplus and the innovator's profits (adjusted for imitators' profits and for profits on older products) in each year up to 1980. These forecasts

were based on firms' expectations concerning $(P_1 - P_2), Q_2$, and the relevant profits in the next few years. They were intentionally very conservative, so whatever bias there was in the resulting rates of return should be downward.

Having made the calculations described in this and the previous section for each year, we had an estimate of the net social benefits (which may be positive or negative) from the innovation for each year. Then we could compute the internal social rate of return, the interest rate that makes the present value of the net social benefits equal to zero. In other words, it is the interest rate, I, that results in the following equality:

$$B(t) + \frac{B(t+1)}{1+i} + \frac{B(t+2)}{(1+i)^2} + \dots + \frac{B(t+n)}{(1+i)^n} = 0, \qquad (2)$$

where B(t) is the net social benefit in year t, t is the first year in which the net social benefit is nonzero, and (t + n) is the last year in which the net social benefit is nonzero.

We computed the private rate of return from the innovator's investment in each innovation. To do so, we calculated the cash flow to the innovator from the innovation during each year. This calculation involved the subtraction of all costs incurred by the innovator to carry out, produce, and sell the innovation (including the allowance described in the previous section for R and D on uncommercialized projects) from the innovator's revenues from the innovation. Also, profits that the innovator would have earned on products displaced by the innovation had to be subtracted. The time period over which these computations were made was generally up to 1973, but in some cases (as in the case of the social rate of return) it extended to 1980. Again, the forecasts in the latter cases were decidedly conservative. The net private benefits in each year, like the Bs in equation (2), were deflated. The Consumer Price Index, which generally was used, is not ideal, but it seemed very unlikely that the results would be affected in an important way by this choice of a deflator.

Finally, one misconception concerning this study should be dispelled. Contrary to some assertions, the sample of innovations included in this study was not confined to "winners". We went to considerable trouble to get as representative a sample as possible. The innovations were chosen at random from those carried out recently by the cooperating firms. A very substantial number turned out to have low or negative private returns. (One interesting finding was that the social rate of return tended to be very high for them, as well as for the others). One of the contributions of this study, in our opinion, was that it included a much broader and more representative sample than any in the past. However, it was not, strictly speaking, a random sample since some firms refused to cooperate because they did not want to disclose such detailed data regarding their innovative activities.

7. Other Types of Innovations

The model in Figure 2 could also handle the process innovations in our sample. In the case of three of the four such innovations in our sample, there was no apparent effect on product prices. By lowering the costs of the innovators, these process innovations increased the innovator's profits. Also, since they were imitated (at nominal cost) by other firms, they soon increased the profits of other firms as well. The total decrease in costs (which equals the increase in profits) of all the relevant firms equals the social saving in resources due to the innovation. For the fourth process innovation, the product price was reduced, and the situation was essentially the same as in Figure 2.

Turning to product innovation used by households, the same model could be applied because all of the innovations of this type in our sample were meant to reduce the cost of some household activity. The extent to which this will be true for ATP-funded projects of this kind is hard to say.

8. Follow-On Studies by Foster and Nathan

During the late 1970s, the National Science Foundation, which financed our work described in previous sections, commissioned two studies, one by Foster Associates and one by Robert Nathan Associates. The purpose was to replicate our study and to enlarge our sample. The findings of these two follow-on studies, based on about 40 innovations in all, were very similar to our own, and the model described above was judged to be serviceable in a wide variety of circumstances. (See Foster Associates, 1978 and Nathan Associates, 1978.) These studies by Foster and Nathan, together with our own, remain the only investigations of the social rates of return of innovations in industry (other than isolated case studies). Clearly, the need for updated analysis is substantial.

9. The Relationship between Social and Private Returns

From the point of view of the ATP program, the gap between social and private rates of return from investments in new technology is of central importance. After all, a major rationale for the ATP program is that some R and D projects have social rates of return far in excess of their private rates of return. What determines the gap between the social and private rates of return? One relevant factor is the market structure of the innovator's industry. If the innovator is faced with a highly competitive environment, it is less likely that it will be able to appropriate a large proportion of the social benefits than if it has a secure monopoly position or if it is part of a tight oligopoly. Of course, the extent to which the innovator is subjected to competition, and how rapidly, may depend on whether the innovation is patented. Another consideration of at least equal importance is how expensive it is for potential competitors to "invent around" the innovator's patents, if they exist, and to obtain the equipment needed to begin producing the new product (or using the new process). In some cases, like duPont's nylon, it would have been extremely difficult to imitate the innovation (legally). In other cases, a potential competitor could obtain and begin producing a "me too" product (or using a "me too" process) at relatively little cost.

Another factor that economists have emphasized as a determinant of the size of the gap between social and private rates of return is whether the innovation is major or minor. According to R.C.O. Matthews (1973), the "degree of appropriability is likely to be less...in major innovations than in minor ones...," since major innovations are more likely, in his view, to be imitated quickly. Similarly, on the basis of a model stressing the indivisibility of information, Kenneth Arrow (1962) concluded that "the inventor obtains the entire realized social benefit of moderately cost reducing inventions but not of more radical inventions."

Still another factor that is sometimes cited is whether the innovation is a new product or process. Thus, Matthews hypothesized that the degree of appropriability might be less for process innovations than for product innovations. On the other hand, Nelson, Peck, and Kalachek (1967) stressed that new processes can often be kept secret and that it frequently is difficult for one firm to find out what processes another firm is using. This, of course, suggests that the gap between social and private rates of return might be greater for products than for processes.

Although most of these hypotheses seem quite plausible, the unfortunate fact seems to be that they have been subjected to only one systematic empirical test, which was based on data for only about 20 innovations (Mansfield et al, 1977). The results seem to support the hypotheses that the gap between social and private rates of return tends to be greater for more important innovations and for innovations that can be imitated relatively cheaply by competitors. Apparently, when the cost of imitating the innovation is held constant, it makes little or no difference whether the innovation is patented -- which seems reasonable because whether or not a patent exists is of relevance largely (perhaps only) because of its effects on the costs of imitation. It is worth noting that this simple model can explain about two-thirds of the observed variation in this gap among the product innovations in our sample. However, at the same time, it is important to bear in mind the smallness (and age) of the sample.

10. Forecasting Social Rates of Return

The U.S. Department of Commerce would like to be able to forecast the social rates of return from the investments in new technology supported by the ATP program. Clearly, there is no reason in principle why such forecasts cannot be made if one is willing to estimate the probability of technical success and the probability of commercialization (given technical success), development and commercialization cost and time, $(P_1 - P_2)$, Q_2 , and the relevant profits. However, due to scanty data at the early proposal review time, such forecasts may be quite inaccurate. The later in the R and D/commercialization stage forecasting is attempted, the better the data are likely to be, and the more reliable the forecast.

One point to bear in mind is that the rate of growth of Q_2 will often have a major impact on the social rate of return. Economists have done a great deal of research on the factors influencing the rate of growth of Q_2 . The diffusion of an innovation is often a slow process. For example, measuring from the date of first commercial application, it generally took more than ten years for all of the major American firms in the relevant industries to begin using industrial robots (Mansfield, 1989, 1993). Similar results have been obtained for the chemical and other industries as well. Also, the rate of diffusion varies widely. Sometimes it took decades for firms to install a new technique, but in other cases they imitated the innovator very quickly. The rate of diffusion of an innovation depends on the average profitability of the innovation, the variation among firms in the profitability of the innovation, the size of the investment required to introduce the innovation, the number of firms in the industry, their average size, the inequality in their sizes, and the amount that they spend on research and development. Using these variables, it is possible to explain a large proportion of the variation among innovations in the rate of diffusion. Moreover, this seems to be the case in a wide variety of industries and in other countries as well as in the United States.

Firms where the expected returns from the innovation are highest tend to be quickest to introduce an innovation. Also, holding constant the profitability of the innovation, big firms tend to introduce an innovation before small firms. In some industries, this may be due to the fact that larger firms -- although not necessarily the largest ones -- are more progressive than small firms. But even if the larger firms were not more progressive, one would expect them to be quicker, on the average, to begin using a new technique for reasons discussed in our earlier studies.¹¹ Also, holding other factors constant, firms with younger and better-educated managers tend to be quicker to introduce new techniques -- or at least, this seems to be the case in industries where firms are small.

Companies also differ greatly with regard to the intrafirm rate of diffusion -- the rate at which, once it has begun to use the new technique, a firm substitutes it for older methods. A considerable amount of this variation can be explained by differences among companies in the profitability of the innovation, the size of the firm, and its liquidity. Also, there is a tendency for late starters to "catch up". That is, firms that are slow to begin using an innovation tend to substitute it for older techniques more rapidly than those that are quick to begin using it. It is also relevant to note that the same sort of process occurs on the international scene; countries that are slow to begin using an innovation tend to substitute it for older techniques more rapidly than countries that are quick to begin using it. The reasons for this tendency, both at the firm and national levels, seem clear enough.

Sociologists have studied the nature and sources of information obtained by managers concerning new techniques. Judging from the available evidence, firms, once they hear of the existence of an innovation, may wait a considerable period of time before beginning to use it. In many cases, this is quite rational. But to some extent this may also be due to incomplete or erroneous information, prejudice, and resistance to change. (For example, see Mansfield (1993).) The sources of information sometimes vary depending on how close the manager is to adopting the innovation. For example, in agriculture, mass media are most important sources at the very early stages of a manager's awareness of the innovation, but friends and neighbors are most important sources when a manager is ready to try the innovation. Also, there is evidence of a "two-step flow of communication". The early users of an innovation tend to rely on sources of information beyond their peer group's experience; after they have begun using the innovation, they become a model for their less expert peers, who can imitate their performance.

In addition, the diffusion process may be slowed by bottlenecks in the production of the innovation -- as in the case of the Boeing 707. Also, the extent of advertising and other promotional activities used by producers of the new product or equipment will have an effect. So too will the innovation's requirements with respect to knowledge and coordination, the diffusion process being impeded if the innovation requires new kinds of knowledge on the part of the user, new types of behavior, and the coordinated efforts of a number of organizations. If an innovation requires few changes in sociocultural values and behavior patterns, it is likely to spread more rapidly. Also, the policies adopted by relevant labor unions influence the rate of diffusion. For example, some locals of the painter's union have refused to use the spray gun. There is, of course, a considerable literature on the effect of collective bargaining on the rate of adoption of new techniques.

11. Models of the Diffusion Process

While economic models of the <u>diffusion</u> process seem to be far more advanced than economic models of the <u>invention</u> process, this does not mean that they are other than crude devices. Particularly when a radical innovation is in its infancy, it is notoriously difficult to predict how rapidly and how far its use will spread. For example, when computers were first introduced, the extent of their diffusion and significance was often under-estimated.

Nonetheless, economic models of the diffusion process, while still in the experimental stages, may prove useful in forecasting the growth of Q_2 . At a minimum, they should be a helpful check on the realism of firms' forecasts. For examples of their use in this regard, see Mansfield (1989, 1993).

12. Extending and Implementing the Model

In section 10, we cautioned that forecasts of social rates of return at the project proposal stage or development stage are likely to be quite inaccurate. No matter what model is used, this would seem to be true. After all, if you have only a hazy idea of what it will cost to complete and commercialize a certain R and D project (or set of such projects), and if you have even a hazier idea of what the outcomes will be, it obviously is very difficult to obtain reliable estimates of social rates of return from these projects. Moreover, these problems are exacerbated by the fact that the ATP program is focused on high-risk projects (which are ones where the forecasting problems are most difficult) and the knowledge spillovers are expected to be substantial.

NIST may, of course, decide to have such forecasts constructed, internally or externally (or both), using the most serviceable techniques at hand. The closer to commercialization the projects are in most of the potential application areas, the more dependable the forecasts are likely to be. Allowing time for the innovation to be subjected to a market test and basing the studies at least partially on actual market performances is preferable. Based on the Foster and Nathan studies, as well as our own, the analytical framework described in Sections 5 and 6 is likely to be as serviceable as any, although it is hardly a panacea. In fact, it is a very simple, straightforward analytical technique that is based largely on principles dating back a good many generations (to Alfred Marshall and beyond). The trick is to get an adequate understanding of what each innovation is (through

discussions with the relevant parties), what its effects were, and how to model these effects faithfully. Imagination and respect for reality are more important than technical flourishes.

But this does not mean that this simple model cannot and should not be extended.¹² For example, if an innovation results in significant spillovers of information that are directly responsible for other innovations in other industries (not merely imitation by the innovator's rivals), an attempt should be made to include this in the calculations. Since none of the innovations in our sample was of this type, we did not have to worry about this possible complication in our 1977 study, but ATP expects that many of the innovations it funds will be of this type. Another kind of extension of this basic model was carried out by Trajtenberg (1990) in his study of CT scanners, which may prove useful in the study of some ATP innovations.

One could develop in advance and try to implement a uniform social rate of return model with extensions covering a broad spectrum of innovations. A more realistic approach is to use the model described above as a flexible framework for examining individual cases as they arise, having gained an understanding of the appropriate business realities of the case. Then one can extend the model accordingly, in those cases where it is necessary.

13. Data Availability and Research Strategy

To get some feel for the extent to which this model might be applicable to innovations derived from ATP-financed work, I contacted 16 ATP award recipients, chosen more or less at random from the list I was furnished by ATP staff, directly without ATP's intermediation.

34

While these conversations were far too brief to provide more than first impressions, there was no obvious indication that this model -- or some simple variant thereof -- would not be applicable, at least as a first approximation.

Equally important, most of the award recipients I spoke with seemed quite willing to cooperate with a study of the social and private returns from innovations based on their results. Several of them were very enthusiastic. But in several cases (all large corporations), there was undisguised hostility to any such effort. As in the case of our 1977 study, I think that some firms are extremely concerned about the leakage of sensitive information. It is almost a cultural phenomenon. I don't think that this attitude is correlated with the social rate of return from innovations based on the firm's results, so the omission of such firms from a study may not be a serious problem. But based on this very small sample, perhaps one-quarter of the contractors are likely to be uncooperative (though there may be some pretense of cooperation, particularly if ATP invokes the terms of the cooperative research agreement which all the firms signed, that, among other things, requires their participation in evaluation studies).

The majority of the firms contacted said that it would be at least two years before the results of their ATP-funded projects would be commercialized. Many felt that, although rough estimates might be made at present, it would take some years before hard numbers would be available regarding profits and widespread savings generated by innovations based on their projects. This, of course, is entirely in accord with our discussion in Section 4.

Perhaps the procedure of estimating the relevant social rates of return should be viewed as a dynamic, continuing process. Like the firm in Figure 1, the investigators responsible for the estimates might update their figures as new data on $(P_1 - P_2)$, Q_2 , and so on become available. It is not certain whether progress on many of these projects to date is adequate to support meaningful analyses of social rates of return. As time goes by, and as the innovations based on ATP-financed R and D are commercialized and their diffusion rates become clearer, estimates of social rates of return should become feasible in more cases and increasingly reliable. The results might provide a wealth of information of use to decision makers, policy analysts, and economists.

Turning to NIST's data collection activities regarding the ATP program, it seems logical to ask award recipients concerning probabilities of technical success and of commercialization (given technical success), as well as development and commercialization cost and time, and $(P_1 - P_2)$, Q_2 , and profits, among other things. It would also be very desirable to launch studies carried out by non-government personnel permitted to obtain such data on a confidential basis from the award recipients. If these personnel are highly reputable, the results would be likely to enjoy greater credibility than findings based on estimates submitted by award recipients to the agency. Rightly or wrongly, data obtained from award recipients by agencies on matters relating to the effectiveness of the agency's programs are sometimes viewed with suspicion.

References

- Arrow, K., "Economic Welfare and the Allocation of Resources for Invention", in National Bureau of Economic Research, <u>The Rate and Direction of Inventive Activity</u> (Princeton: Princeton University Press, 1962).
- Beardsley, G., and E. Mansfield, "A Note on the Accuracy of Industrial Forecasts of the Profitability of New Products and Processes", Journal of Business, January 1978.
- Evenson, R., P. Waggoner, and V. Ruttan, "Economic Benefits from Research: An Example from Agriculture", <u>Science</u>, September, 14, 1979.
- Foster Associates, "A Survey on Net Rates of Return on Innovations", Report to the National Science Foundation, May 1978.
- Griliches, Z., "Research Costs and Social Returns: Hybrid Corn and Related Innovations", Journal of Political Economy, 1958.
- Jaffe, A., M. Trajtenberg, and R. Henderson, "Geographical Localization of Knowledge Spillovers as Evidenced by Patent Citations", <u>Quarterly Journal of Economics</u>, August 1993.
- Mansfield, E., "The Contributions of New Technology to the Economy", Brookings Institution and American Enterprise Institute, forthcoming (a).
 - _____, "A Note on Estimating the Returns from New Technology: How Much Learning?", International Journal of Technology Management, forthcoming (b).
- _____, <u>Innovation, Technology, and the Economy</u> (Aldershot: Elgar, 1995).
- _____, "The Diffusion of Flexible Manufacturing Systems in Japan, Europe, and the United States, Management Science, February 1993.
- _____, "The Diffusion of Industrial Robots in Japan and the United States", <u>Research</u> <u>Policy</u>, 1989.
- _____, "Technical Change and the Rate of Imitation", Econometrica, October 1961.
- _____, Industrial Research and Technological Innovation, (Norton, 1968).

- _____, J. Rapoport, A. Romeo, S. Wagner, and G. Beardsley, "Social and Private Rates of Return from Industrial Innovations", <u>Quarterly Journal of Economics</u>, May 1977.
- _____, F. Husic, J. Rapoport, A. Romeo, E. Villani, and S. Wagner, <u>The Production and</u> <u>Application of New Industrial Technology</u> (New York: Norton, 1977 b).
- , and S. Wagner, "Organizational and Strategic Factors Associated with Probabilities of Success in Industrial R and D", <u>Journal of Business</u>, April 1975.
- _____, J. Rapoport, J. Schnee, S. Wagner, and M. Hamburger, <u>Research and Innovation in</u> the Modern Corporation (New York: Norton, 1971).
- Matthews, R.C.O., "The Contribution of Science and Technology to Economic Development", in B. Williams, ed., <u>Science and Technology in Economic Growth</u> (London: Macmillan, 1973).
- Nathan Associates, "Net Rates of Return on Innovations", Report to the National Science Foundation, July 1978.
- Nelson, R., M. Peck, and E. Kalachek, <u>Technology, Economic Growth, and Public Policy</u> (Washington, D.C.: Brookings Institution, 1967).
- Steele, L., "Selecting R and D Programs and Objectives", <u>Research-Technology</u> <u>Management</u>, March-April 1988.
- Trajtenberg, M., <u>Economic Analysis of Product Innovation</u>: <u>The Case of CT Scanners</u> (Cambridge: Harvard, 1990).

Endnotes

1. This paper draws freely on my publications listed in the References.

2. See Steele (1988).

3. R and D executives of this firm provided a description of these analyses. The identify of the firm is withheld to insure that no confidential information is divulged.

4. The material in quotation marks in this section comes from internal memoranda of the firm provided by R and D executives.

5. See Mansfield and Wagner (1975).

6. Ibid.

7. <u>Ibid</u>.

8. There is some fragmentary evidence that firms' estimates of the returns from the <u>adoption</u> of a new technology may have improved in the past thirty years. See Mansfield (forthcoming b). But this, of course, does not mean that the kinds of forecasts discussed in this section have improved significantly.

9. The identity of this firm is withheld to insure that no confidential information is divulged.

10. Profit here is defined to be gross of any depreciation on the investment in the innovation. Basically, it is a cash-flow concept.

11. See Mansfield (1968).

12. For a list of limitations of our model, see Mansfield et al (1977 a, b).