

NBER WORKING PAPER SERIES

RESEARCH AND DEVELOPMENT AS AN INVESTMENT

Bronwyn Hall

Fumio Hayashi

Working Paper No. 2973

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1989

We are grateful to Zvi Griliches for discussions and encouragement that led to this paper being written. The first author thanks the Olin Foundation and the Committee on Research, University of California, for support. Some of the data preparation effort was supported by the National Science Foundation (PRA 81-08635) and the National Bureau of Economic Research. Seung Hee Choi provided able research assistance. This paper is part of NBER's research program in Productivity. Any opinions expressed are those of the authors not those of the National Bureau of Economic Research.

NBER Working Paper #2973
May 1989

RESEARCH AND DEVELOPMENT AS AN INVESTMENT

ABSTRACT

About 20 percent of the gross investment expenditures of U.S. manufacturing firms is expenditures on research and development. Like investment in physical capital, R&D also responds to news about future prospects of the firm, such as profitability, technological opportunities, or changes in factor prices. Using data from a panel of large U.S. manufacturing firms that was developed within the Productivity Program of the NBER, we investigate the differential responses of these two types of investment to changes in the value of the firm's assets as perceived by financial markets and the interaction of these responses.

In order to study this topic empirically, we develop a stochastic dynamic programming model of a firm with two types of capital (physical and knowledge capital) which are used to produce profits. A feature of the model is the distinction between the accumulation of the two kinds of capital: expenditures on the physical capital stock are incurred one or more years before the capital actually becomes productive, whereas R&D capital is produced jointly as a function of current expenditure and the past technological position of the firm. Two individual firm-specific shocks are considered: one to the overall profitability of the firm, and one to the "productivity" of R&D. In the empirical estimates, we find that these two shocks account for about 20 percent of the total variance in net investment, 15 percent of the variance in the firm-level R&D to capital ratio, but only about 5 percent of the annual rates of return. The profitability shock is well described by a moving average process of order three, while the technology shock process is more nearly permanent: first order autoregressive with parameter near unity.

Bronwyn Hall
NBER
204 Junipero Serra Blvd.
Stanford, CA 94305

Fumio Hayashi
Economics Department
University of Pennsylvania
Philadelphia, PA 19104

RESEARCH AND DEVELOPMENT AS AN INVESTMENT

Bronwyn H. Hall and Fumio Hayashi¹

1. Introduction

Economists and policymakers have long been interested in the response of firm-level investment to changes in the macroenvironment. In the recent past, studies relating investment to such variables as output, anticipated or desired output, or the market value of a firm have been carried out by many researchers. (See Chirinko 1986 for a recent survey of investment models). Almost all of these studies--with the notable exception of Epstein and Denny (1983), who use macro data, and Nadiri and his co-workers, who use firm-level data--have focused on a single type of investment, usually an aggregate of plant and equipment expenditure. Although the physical capital stock of a firm may be the most important generator of net cash flows and hence profits or returns to the holders of equity in the firm, other forms of capital are also capable of generating these returns and may in fact be more likely to

generate more long-lasting, supranormal returns.

An important type of capital of this sort, at least for a large part of the manufacturing sector, is knowledge or R&D capital--the accumulated know-how, technical expertise, trade secrets, patents, etc., that are embodied in the firm and its employees. Investment in this type of capital by industry has become an increasingly large component of total industrial investment in the post-war period (see Table 1). The level of R&D capital held by firms is also an important component of the market value of those firms (Griliches 1981, Hall 1987). There has recently been considerable discussion (e.g., Mansfield 1986; Eisner, Albert, and Sullivan 1986) over whether the level of R&D investment by firms is sufficiently high, even considering only the private rate of return, and whether tax credits are the appropriate tool for subsidizing it (on the grounds that the social rate of return is even higher). We are therefore prompted to ask whether R&D investment at the firm level behaves like ordinary investment in the level and timing of its response to changes in the firm's environment signalled by changes in the market value of the firm. We also investigate how this type of investment interacts with the investment in physical capital, since presumably substitution possibilities do exist for the firm engaged in maximizing the returns generated from both types of capital.

In previous work, we and others have used a traditional Tobin's Q approach to value the existing stocks of capital in the manufacturing firms in our sample (Griliches 1981, Griliches and Cockburn 1988, Hall 1987, Hall 1988). This methodology yields coefficients that may be interpreted as the average shadow value of capital of a certain type in a firm that also has average levels of the other types of capital. As emphasized by Hayashi (1982) and others, finding that these coefficients

are less than unity may or may not signal that disinvestment should take place, since investment decisions should be based on the marginal value of new capital, not the average value of old capital, and these two could differ widely. This phenomenon is probably what explains the paradoxical result that the current level of R&D spending is valued twice as highly as the associated capital stock, even allowing for the conversion of a stock into a flow (see Hall 1988). Under the assumption of present value maximizing behavior on the part of the firm, the value of R&D capital is the value of the knowledge produced by past investments plus the value to the firm of the optimal trajectory of future R&D investments. Because of the random-walk nature of R&D spending (Hall, Griliches, and Hausman 1987), current expenditures on R&D are a good measure of the whole trajectory; when they are converted from a flow to a stock, they perform better in the regression than the stock that has been created from the past history of investments.

This empirical finding is an example of how the value function has to do double duty empirically, in that it must simultaneously value the future revenues to be obtained from the old capital and the revenues to be obtained from investments not yet undertaken by the firm. Firms that do not execute the optimal decision path should see their values decline in the next period, other things equal, but we have no way of separating this effect from the effects of bad news about the firm's prospects in the next period. We must simply assume that, on average, firms make the right decisions so that our coefficients measure the true shadow prices. The use of stock price data as an indicator of the value of a firm's investment program requires that one assume both market efficiency (so that the measured value is correct) and optimizing

behavior on the part of the firm (so that we can solve for the value function using dynamic programming techniques).

The approach of Pakes (1985), among others, sidesteps some of the difficulties of distinguishing marginal and average values of capital by focusing on the marginal responses of investment to news in market value. By holding the capital stocks fixed for a short interval, we can identify changes in market value entirely with news about demand, technology, etc., which influences future development. This is the approach followed in this paper: we measure the magnitude and timing of the changes induced in investment spending by shocks to the market value of the firm. The model is that of a firm buffeted by shocks to the macroeconomy, its industry, the (non-forecastable) behavior of competitors, factor prices, and, in particular, the output of its R&D program. These shocks cause changes in the market value of the firm and thus in Tobin's Q . The firm responds to this shocks by adjusting the level of spending on the two types of capital, physical and research and development.

Recent work by Lach and Schankerman (1988) investigated some of the interactions between the two kinds of investment (i, r) and the one-period rate of return (q) to the firm using dynamic factor methods and a dataset that overlaps considerably with ours. Their model is explicitly non-structural, but informative as to timing, and provides an interesting descriptive summary of the stylized facts. They found that the data are adequately described by three factors, one of which is idiosyncratic to q , one of which affects both i and r rapidly and permanently, and one of which is a transitory factor that is idiosyncratic to i . Neither of the two investment factors is an important explanator of the variance in q (they account for at most 5 percent of that variance), and the

factors cannot be identified exclusively with shocks to demand or to factor prices.

Using a slightly different but essentially equivalent methodology, Griliches, Hall, and Pakes (1987) have found three factors linking sales, investment, R&D, patents, and q for a similar sample of firms. In addition to the idiosyncratic factor in q (again accounting for about 95 percent of the variance), there is one factor that we label demand linking sales, investment, and R&D and a much weaker factor that we interpret as shifts in technological opportunity linking R&D, patents, and q . This second factor may also appear very weakly in investment, but it then becomes doubtful to identify it with technological shocks. Our results are inconsistent with Lach and Schankerman's since we do find it necessary to include a second factor that is not in investment and sales to explain the movements of R&D expenditures. Also, the additional identifying power we hoped to bring to bear on the problem via patents turned out to add almost no information to the problem in a panel data (within-firm) setting so that the inclusion of the second factor is not caused by the presence of patents in the model.

This paper takes a different approach to the problem from the two just described in an effort to provide interpretation of the observed stylized fact within the context of an explicit optimizing model of the investment behavior of the firm. To do this, we must confront the problem that although work on theoretical models of investment decision where there is more than one type of capital stock has existed at least as far back as Lucas (1967), it is difficult to extract from these models a closed-form solution for the investments as a function of the various prices and initial conditions. Thus there has been relatively

little empirical investigation of such models, partly due to their complexity in all but simple cases. Such work as does exist is in the cost or profit function framework, which takes output as predetermined, and introduces lags in investment via adjustment costs (Bernstein and Nadiri 1982), owing to the relative tractability of that type of model. Our model is of a different kind since it takes the firm's output and price-setting behavior as endogenous by focusing on the present discounted value of the profits yielded by the capital assets of the firm without being specific about the type of market in which the firm operates. The profits the firm earns from its choice of assets may be due either to specific factors, know-how, some level of market power arising from an oligopolistic structure, or to a combination of all of these. Although this kind of model cannot be used to investigate market structure and the competitive behavior of firms, it is useful for interpreting the timing and interaction of investment decisions without having to impose a particular kind of competition.

In the simple flexible accelerator model or the q model of investment, investment is assumed to respond to shocks in the environment that shift the level of desired output (and in the case of the q model, to other shocks such as changes in the relative price levels of inputs or outputs). The models are silent on the origin of these shocks, whether demand or supply (technological). The purpose of the more elaborate model presented in this paper is to identify the relative importance of technological shocks in influencing investment decisions using a comprehensive panel dataset of publicly traded manufacturing firms.

In performing this analysis, our identifying assumption is that shifts in technological opportunity appear first in the investment

choices of the firm and only later in increased output. That is, there is a gestation lag between the expenditure on investment and the movement of the capital thus obtained into productive use. Because we use a specific functional form for the profit function (corresponding to a constant returns Cobb-Douglas for the production function), we are able to distinguish profitability shocks (which appear first in sales and only later in the productive capital) from those shocks that are unrelated to current sales but cause contemporaneous changes in capital spending and R&D investment. We call this kind of shock a "technological opportunity" shock since it appears in the model as a revaluing of R&D capital. Because R&D and physical capital are used jointly to produce profits, this shock affects both types of investment.

Note that it is not possible to identify either of our two shocks as demand or supply (a la Lach-Schankerman), since the profitability shock represents news in demand or in factor prices, which affects the optimal choice of capital with a lag, whereas the technological shock embodies all the factors that force a revaluation of R&D capital, including those that are demand-related. The interpretation of our factors is only really possible in the framework of our specific model of the firm.

Griliches, Hall, and Pakes (1986) used patent data to identify this supply shock to the productivity of R&D explicitly and found that its relative importance in the variance of q is anywhere from one percent to seventy percent (in the drug industry) of demand-related shocks to q and very poorly measured. In this paper, we are using evidence on the timing of R&D and capital expenditures relative to sales to disentangle the technological and non-technological effects.

2. Model

In building a model of two kinds of investment (in R&D and physical capital), we chose not to use the full adjustment cost methodology familiar from the investment literature but instead introduce gestation lags into both R&D and investment decisions along with an intertemporal link for R&D investment. This is more tractable than an adjustment cost model involving two capital stocks but still preserves the determinacy of firm size and the relationship between investment and lagged sales that are requirements imposed by the data.

Our model posits a price-taking firm that maximizes the present discounted value of the cash flow (profits) generated by the two assets K_t (physical capital) and G_t (knowledge capital). We write profits in period t as a function of the quasi-fixed factors (variable factors have already been solved out):

$$(1) \quad \pi_t = \pi (K_{t-\tau}, G_{t-\theta}, u_t)$$

where K_t and G_t denote stocks at the end of period t and τ and θ are gestation lags. That is, the capital stock that determines profits in period t was chosen, and the appropriate investment performed τ years prior.

In constructing this model of the firm choosing two kinds of investment, a stylized fact we wish to interpret is the distinctive behavior of R&D investment dynamics as opposed to those of investment in physical capital. This is the fact alluded to above, that the growth rates of investment have a sample variance much larger than that of R&D expenditures (Mairesse and Siu 1984, Lach and Schankerman 1988). Although it may be due to a large difference in adjustment costs for the

two kinds of capital, we believe that it may also be due to a difference in the way the two kinds of capital are produced. That is, the production of R&D or knowledge capital is intrinsically different from the procurement and installation of capital goods within the firm since it is likely to involve projects of duration longer than one year that cannot be speeded up greatly without a considerable increase in cost per unit of output. Ten man-years of R&D engineers may not have the same productivity when considered as two men for five years as when it is five men for two years.

We capture this idea by assuming the conventional form of capital accumulation for physical capital:

$$(2) \quad K_t = I_t + (1-\delta) K_{t-1}$$

but assuming a Cobb-Douglas form for R&D capital:

$$(3) \quad G_t = \xi_t R_t^\sigma G_{t-1}^{1-\sigma}$$

where

I_t = physical investment,

R_t = R&D investment, and

ξ_t = technological shock.

Both equations (2) and (3) are homogeneous of degree one, but the first has a marginal product of investment equal to unity, while for the second the marginal product depends inversely on the rate of investment ($\sigma G_{t-1}/R_t$), so there are diminishing returns to R_t . If we consider the problem of maximizing the amount of capital stock available at time t subject to a fixed level of investment I (with non-negative investment in all years), the first formulation will yield an optimal investment

pattern of $(0, 0, 0, \dots, 1)$, whereas the second will in general yield a pattern that is spread across all the years (if $\sigma < 1$).

The firm's problem is to choose two decision rules at time t , h^I and h^R , that maximize the expected discounted present value of the cash flows generated by those investment paths:

$$(4) \max_{h^I, h^R} V_t = E_t \sum_{j=0}^{\infty} \gamma^j (\pi_{t+j} - a_{t+j} I_{t+j} - b_{t+j} R_{t+j})$$

subject to the profit function (1) and the capital accumulation constraints (2 and 3). We now proceed to attempt a solution to the firm's problem that can be expressed in terms of the observable quantities, which are the stock of physical capital, the market value of the firm, the output (sales), and the investment in research and development. Note that an important difference between our two types of capital is that only the stock of physical capital is observable independently from the investment history (in principle). Although we could construct a stock of knowledge capital from the history of R&D expenditures, this would add no new information to the problem. Accordingly, in the search for tractable estimating equations, we will need to eliminate G_t from the model.

Since there are no adjustment costs in K , V_t can be written as

$$(5) V_t = a_t(1-\delta)K_{t-1} + E_t \sum_{j=0}^{r-1} \gamma^j \pi_{t+j}(K_{t+j-r}, G_{t+j-\theta}, u_{t+j}) + W_t$$

where

$$(6) W_t = E_t \sum_{j=0}^{\infty} \gamma^j [\gamma^r \pi_{t+r+j}(K_{t+j}, G_{t+j-(\theta-r)}, u_{t+r+j}) - c_{t+j} K_{t+j} - b_{t+j} R_{t+j}]$$

and c_t is the user cost of capital:

$$(7) \quad c_t = a_t - \gamma(1-\delta) E_t (a_{t+1})$$

The second term in V_t represents the profits still to be earned by the physical capital stock that has already been chosen and over which the firm has no control. This is a consequence of the gestation lag assumption. If we assume that $\theta \geq \tau$, $(G_{t-\theta}, G_{t-\theta+1}, \dots, G_{t-\theta+\tau-1})$ will also be beyond control of the firm at time t , and this term will be . exogenous to the firm at time t .

Thus we have eliminated the investment from the optimization problem in (4) and reduced the problem to one of choosing K_t (for production in period $t+\tau$) and a stochastic R&D investment program h^R that maximizes W_t subject to the accumulation constraint (3).² The first order condition for K_t is easily obtained:

$$(8) \quad E_t \left[\gamma^\tau (\partial \pi(K_t, G_{t-(\theta-\tau)}, u_{t+\tau})) / \partial K_t \right] = c_t$$

Characterization of the optimal decision rule for R_t is less straightforward. If $\theta - \tau > 0$, not only G_{t-1} (which shows up in the R&D accumulation constraint) but also $\{G_{t-2}, \dots, G_{t-(\theta-\tau)}\}$ influence W_t

.....

2. Note here that the role of R&D capital stock in this model is crucial since in its absence the value function would be a linear function of the K_t s (under the constant returns assumption of equation 10) with

weights equal to $\gamma^\tau u_{t+\tau} - c_t$. The optimal plan in any period would be either to increase the capital stock infinitely or to set it to zero, depending on whether the it earns more or less than its rental rate. This is a consequence of the no-adjustment cost assumption. In our model, since R&D capital is jointly produced over several periods and there is substitution between the two kinds of capital, the size of the firm is determinate.

because they enter $\pi_{t+r}, \dots, \pi_{t+\theta-1}$. This will create difficulties for estimation due to the non-observability of $\{G_t\}$, and so we set $\theta-r = 0$, which implies that W_t depends on K_t , G_{t-1} , and the investment path $\{R_t\}$. K_t and the optimal investment path $\{R_t\}$ (conditioned on information available to the firm at time t) are functions in turn of G_{t-1} after solution of the first order conditions for equation (4). A familiar property of the value function (see, e.g., Sargent 1986, chapter 1, or Duffie 1988, pp. 196-97) implies that

$$(9) \quad \mu_t = \frac{\partial W_t(G_{t-1})}{\partial G_{t-1}} = -b_t \frac{\partial R_t}{\partial G_{t-1}} \Big|_{G_t = \text{const.}}$$

$$= -b_t \frac{1-\sigma}{\sigma} \frac{R_t}{G_{t-1}}$$

To obtain an explicit solution to these first order conditions, we assume that the technology is Cobb-Douglas with constant returns to scale. Then nominal sales are proportional to profits, and profits are

$$(10) \quad \pi(K_t, G_t, u_{t+r}) = u_{t+r} K_t^\alpha G_t^{1-\alpha}$$

so that the first order condition (8) becomes

$$(11) \quad E_t (u_{t+r} (K_t/G_t)^{\alpha-1}) = c_t \gamma^{-r} \alpha^{-1}$$

Now we use the fact that a constant fraction (ϕ_i) of sales is profits and take logs of (10), (11), (3), and (9), denoting the logs of S , K , G , and R by lower case letters and the individual firm observations by i :

$$(12) \quad s_{it} - k_{it-r} = \ln u_{it} - (1-\alpha_i)(k_{it-r} - g_{it-r}) + \phi_i$$

$$(13) \quad k_{it} - g_{it} = (1-\alpha_i)^{-1} [\ln E_t u_{it+r} - \ln c_t + \ln \alpha_i \gamma^r]$$

$$(14) \quad g_{it} - g_{it-1} = \ln(\xi_{it}) + \sigma(r_{it} - g_{it-1})$$

$$(15) \quad r_{it} - g_{it-1} = -\ln b_t - \ln((1-\sigma)/\sigma) + \ln \mu_{it}$$

There are two remaining problems with this formulation. First, g_{it} is not observable, as noted before, and the marginal value product of g_{it} , μ_{it} is also unobservable. The first problem is solved by eliminating g_{it} and $g_{i,t-1}$ from equations (12) - (15):

$$(16) \quad s_{it} - k_{it-\tau} = \ln(u_{it}) - \ln E_{t-\tau}(u_{it}) + \text{firm \& time dummies}$$

$$(17) \quad k_{it} - k_{it-1} = \ln(\xi_{it}) + \sigma \ln \mu_{it} \\ + (1-\alpha_i)^{-1} [\ln E_t u_{i,t+\tau} - \ln E_{t-1} u_{i,t+\tau-1}] + \text{time dummies}$$

$$(18) \quad r_{it} - k_{it-1} = \ln(\mu_{it}) - (1-\alpha_i)^{-1} \ln(E_{t-1} u_{i,t+\tau-1}) + \text{time dummies}$$

Equation (16) specifies that the capital-output ratio in period t is equal to an individual firm constant plus year effects that reflect changes in the macro-environment plus the idiosyncratic news in the firm's demand, which arrived since the capital stock was chosen (at $t-\tau$). Equation (17) describes the choice of the capital stock made in year t that will be used for production at year $t+\tau$: it depends on the shock to the R&D accumulation equation at year t , the shadow price of R&D capital at year t , and the expected change in demand between year $t+\tau-1$ and year $t+\tau$ as well as on a set of time dummies. Finally, equation (18) describes the optimal choice of R&D intensity in the same year, which is a function of time and firm dummies, the shadow price of R&D capital μ_{it} , and the expected demand τ periods hence.

Now we confront the problem of how to measure the shadow price of

R&D capital μ_{it} , which is unobservable. If we specify the stochastic processes that generate u_{it} and ξ_{it} , then in principle the model could be solved to obtain an expression for W_{it} , the value of the firm's problem from $t+r$ onward as a function of G_{it-1} and the distribution of $\{u_{it}\}$ and $\{\xi_{it}\}$. However, there does not exist a closed-form solution of this kind. We therefore adopt the following strategy. First, assume that $\ln(u_{it})$ and $\ln(\xi_{it})$ follow mutually uncorrelated moving average processes:³

$$(19) \ln(u_{it}) = \lambda_0 x_{it} + \dots + \lambda_n x_{i,t-n}$$

$$(20) \ln(\xi_{it}) = \delta_0 y_{it} + \dots + \delta_m y_{i,t-m}$$

Then μ_{it} depends on $\alpha_i, x_{it}, \dots, x_{it+r-n}$, and $y_{it}, \dots, y_{i,t-m}$ but not on $x_{i,t+r-n-1}, \dots, x_{i,t-n}$ since the probability distribution of $u_{it}, \dots, u_{i,t+r-1}$ does not affect W_{it} . Then use a loglinear approximation for $\ln \mu_{it}$:

$$(21) \ln \mu_{it} = \eta_r x_{it} + \dots + \eta_n x_{i,t+r-n} + \gamma_0 y_{it} + \dots + \gamma_m y_{i,t-m} \\ + \text{firm and time dummies.}$$

When equations (19)-(20) are substituted into equations (16)-(18), we obtain a system of equations for the output capital ratio, the growth in net capital stock, and the R&D-capital ratio in terms of the unobservable shocks $\{x_{it}\}$, $\{y_{it}\}$, and firm and time dummies. For

3. Note that because of the shortness of our panel, this specification is general enough to accommodate any ARMA process, not just moving average. We simply need n and m greater than the number of lags in the data to allow for an AR process. Later on we present evidence that an AR(1) process is a more parsimonious specification for $\ln \xi_{it}$.

estimation and identification purposes, we augment these equations with an equation for the one-period rate of return to the firm's common stock. The use of this variable is based on an idea in Pakes (1985): although we cannot explicitly solve for the value of the optimal investment program in this model, we do know that changes in this value should depend only on the "news" that arrives between period t and period $t-1$, under the assumption of efficient stock markets and optimal behavior by the firm. This leads us to write the one-period rate of return to equity (the news in market value) as a linear function of the contemporaneous shocks:

$$(22) \quad q_{it} = \text{time dummies} + \sigma_x x_{it} + \sigma_y y_{it} + \sigma_z z_{it}$$

where $\text{Var}[x_{it}] = \text{Var}[y_{it}] = \text{Var}[z_{it}] = 1$, and x , y , and z are mutually uncorrelated. z_{it} is the part of the excess rate of return that is not related to news in demand or the "technology production function."

In the next section, we discuss the data we will use for estimation of this model. This will lead us to modify the specification just derived slightly before estimating the model in section 4.

3. Data and Preliminary Results

The data come from the Compustat files, cleaned and merged as described in Hall et al. (1988). The universe from which we draw our sample consists of all the manufacturing firms that existed for at least four years between 1976 and 1985 on the Annual Industrial or Over-the-Counter file, which includes all firms traded on the major stock exchanges and NASDAQ.

In constructing the samples used in this paper, we removed firms

that experienced major mergers or other reorganizations or whose capital stock changed in any one year by more than one hundred percent during the periods we were examining (1973 to 1981, or 1977 to 1985). We also removed firms for which we could not construct a stock market rate-of-return that coincided with their fiscal year, usually because they did not appear on the Quarterly Compustat file. Finally, we required that each panel be balanced so that sales, capital stock, q , and R&D were available for all nine years from 1973 to 1981 or 1977 to 1985. Only the last requirement (on R&D) caused us to lose a significant number of firms that did not enter or exit during the period. It was obviously necessary since our model is about firms that have R&D programs, which are over half of those in the sample.

Table 2 exhibits some summary statistics for the two samples. The variables of interest are sales; the stock of physical capital, measured at the end of the period (as in our model) and adjusted for the effects of inflation as described in Hall et al. (1988); research and development expenditures; and the one-period rate of return to holding a share of the firm's common. This latter variable is defined as

$$(23) \quad q_{it} = (p_{it} - p_{it-1} + d_{it-1})/p_{it-1}$$

where p_{it} is the price at the beginning of period t and d_{it} are the dividends paid during period t ; and all quantities have been adjusted for splits, stock dividends, etc.⁴ The other variables shown in the

4. The alert reader will note that this definition of q_{it} implies a timing convention where news arrives during the previous period ($t-1$) and investment plans are made at the beginning of year t . This convention was also followed in Pakes (1985), Griliches, Hall, and Pakes (1987), and Lach and Schankerman (1988) for the following reason: the data indicate that the strongest correlations between investment variables and q are

table are the growth rates of sales, net capital, R&D investment, and gross investment, the log of the sales-capital ratio with r set to one, $s_{it}^{-k_{i,t-1}}$, and the log of the R&D-capital ratio $r_{it}^{-k_{i,t-1}}$. These variables exhibit the usual characteristics of growth rates and ratio variables in the panel data of firms: most of the variance of the growth rates is within firm, while most of the variance of the ratios is between firms (see the last row of each section of the table).⁵ Note in particular the large amount of unexplained variance in the annual stock market rate of return, corresponding to a standard deviation of forty percent.

Although our model is in terms of real variables, we use (the logs of) nominal variables for estimation and include a set of free time dummies and a set of twenty industry dummies at the roughly two-digit level in each equation; the only remaining deflation that has not been performed is that which is firm-year specific. Thus the variance and covariance relationships which we will be fitting are based on the row of Table 2 labelled "Var. within year & ind."

Before attempting to estimate a model like that in equations (16)-(18) and (22), in Table 3 we examine the variances and covariances of

between this year's investment and last year's q . Of course, some of last year's investment might also be a response to news in q last year but the contemporaneous q also includes news to which investment in the first part of the year is unable to respond. The necessity for a choice of timing convention arises from the time aggregation inherent in this kind of data.

5. See Mairesse (1987) for further discussion of this phenomenon and more examples.

the data in an effort to ascertain the model's plausibility and the likely values of τ , n , and m . We use a smaller subset of the firms (those with data from 1973 through 1985) in Table 3 so that we can look at the covariances over a longer period (out to six lags); the results were also checked on the two slightly larger samples on which Table 2 is based. From evidence cited in Kydland and Prescott (1982) and R. E. Hall (1977), the most reasonable choices of τ seem to be somewhat less than one for equipment (two quarters) and about two for plants. Our capital stock figure combines both types of stocks; in the manufacturing sector as a whole, equipment forms about two-thirds of the capital stock (BLS 1979), so that a τ of about one would be the most reasonable. Using a completely different methodology and data similar to ours, Pakes and Griliches (1984) find that the response of profits to investments made in prior years increases to about lag three and then is level for several years thereafter. This would imply an average gestation lag of about two years or slightly more. Of course, it is much more likely that capital moves into production with a lag that varies across industry and across the cycle, but this is difficult to model (see Chirinko (1988)), so we choose instead to compare the use of a fixed value of τ equal to 1, 2, or 3 in our estimation. For the presentation of the variances and covariances in Table 3 we have chosen τ equal to one (capital stock acquired this period is used in production next period).

The analysis of variance in Table 2 and set of auto-covariances and cross-covariances in Table 3 reveal several interesting facts. First, as mentioned in the introduction, the variance of the growth rate of gross investment is much larger than that for R&D investment. Going within firm increases this effect, implying that R&D is somewhat more

persistent than investment. Second, the logs of the sales-capital ratio and R&D-capital ratio vary little within each firm over time; most of the variation is permanent and across firms. This can be seen both in the analysis of variance in Table 2 and in the persistence of the autocovariances at long lags in Table 3. From the point of view of our model specification, this means that firm-specific effects are likely to be important in these equations.

Recalling the derivation of the model, the firm-specific effects will arise from the following sources: differences in the firms' production or profit functions of the two capitals, α_i ; differences in the firms' profit-sales ratios, due to differences in value added for market power, ϕ_i ; differences across firms in the mean of the demand shock u_{it} ; and differences across firms in the mean of the technological shock or R&D productivity ξ_{it} . Both u_{it} and ξ_{it} appear only in logarithmic transformation so their overall levels for each firm do not matter. It is plausible to assume that u_{it} , which is the unexpected return to capital under the Cobb-Douglas linear homogeneity assumption of equation (10), has an expected value that is roughly equal across firms. Similarly, the expected value of the news in the R&D productivity shock is probably the same across firms. However, these firms come from many different industries and perform quite different activities, so it is reasonable to assume that there are permanent differences in the shares of capital, R&D capital, labor, and raw materials devoted to the production of output and hence differences in α_i and ϕ_i .

Under these assumptions, the firm specific effects are the following:

$$(24) \quad sk: \phi_i - \ln \alpha_i$$

$$\Delta k: (1 - \alpha_i)^{-1} (\ln E_t u_{i,t+r} - \ln E_{t-1} u_{i,t+r-1} + \ln c_{t-1} - \ln c_t)$$

$$rk: \ln E_t \mu_{it}$$

Note that we expect the expected value of μ_{it} , the shadow value of an additional unit of R&D capital, to vary across firms since it depends on α_i , so we denote this expected value by μ_i . The evidence in Table 3 suggests that while the first and third of these firm effects have good-sized variances across the sample of firms, the variance of the second is small, which is not surprising since it is a function of the differences between two adjacent periods of the expected net return to capital in excess of its rental rate.

The easiest predictions of the model to check are those involving q , since it contains only contemporaneous shocks. The first implication of this is that q should exhibit no serial correlation, but the table shows a small amount of negative serial correlation at most lags. This may be due either to measurement error in the stock prices (which is likely to affect a small number of q values), or to some sort of mean-reverting behavior of the type documented by Fama and French (1987) or Poterba and Summers (1987). Since the effect is small (in fact, the joint hypothesis that all lagged covariances are zero yields an $F(6,14903) = 2.36$) and other research in this area using different periods has found even smaller effects of this kind (Pakes 1985; Griliches, Hall, and Pakes 1986; Lach and Schankerman 1987), we chose to ignore this in what follows.

The second implication of the efficient markets assumption together with the moving average shocks is that the covariances of q with lagged

values of s_k , Δk , and r_k should all be zero but that its covariance with the leading values of s_k , Δk , and r_k should be non-zero out to the order of the moving average processes. Because of the time aggregation problem alluded to before, we do not really expect this to hold for the covariance of q with the first lag of each variable (since q is the rate of return in the year before t). Joint tests of the hypotheses that the covariances of q with the second through sixth lags of s_k , Δk , and r_k were conducted and all were easily accepted. With the exception of s_k , the tests were also accepted when the first lag was included. This provides support for the idea that investment plans are made early in the year and tend to lag behind the news in q .

For s_k , the sharp gestation lag assumption implies that its covariances with q should drop to zero after $\tau-1$ periods (or after τ periods, allowing for time aggregation). There is a sharp drop in the covariances from lead 2 to lead 3, and the auto-covariances of s_k also appear to reach their permanent level by lag 3, so $\tau=2$ appears to be the best choice. The forward covariances of q with Δk and r_k are less conclusive: those with Δk appear to die away after about two lags, while those with r_k are somewhat more persistent, but small, for all three of the samples.⁶ It is likely that at least one of the two

6. This persistence is another demonstration of the time series characteristics of R&D investment that were noted in Griliches, Hall, and Pakes (1987): the logarithm of R&D expenditures for any firm can be well approximated as a random walk with a small error variance. This implies that the shocks to r are permanent and would suggest an autoregressive specification for $\ln \xi_{it}$ rather than moving average.

In fact, our preliminary estimates suggested this was the case and we modified our model to take account of this in Section 4.

processes $\ln u_{it}$ or $\ln \xi_{it}$ is of high order.

Finally, note that the covariance pattern of sk and rk is consistent with permanent differences in the R&D-to-sales ratio coupled with a gestation lag of 2 on the capital stock. This is because sk is a moving average of the shocks x_{it} of order r plus a permanent firm effect so that the covariance of sk and rk should be a constant (arising from permanent differences across firms) after the first $r-1$ lags. Allowing for some spillover due to time aggregation implies a r of two. The same is not true for the covariances of sk and Δk which persist out to the longest lag shown. One possibility is that the model for the capital stock k should be augmented to include measurement error, both for this reason and on a priori grounds, since this is the variable that we expect is most poorly measured in our dataset. The most natural way to do this is to assume that the growth of the capital stock from year to year contains a serially uncorrelated measurement error:

$$(25) \quad \Delta k_t = \Delta k_t^* + \varepsilon_t$$

where $\text{Var}[\varepsilon_t] = \frac{2}{\varepsilon}$ and $E[\varepsilon_t \varepsilon_s] = 0$ for $s \neq t$. This will imply a cumulative measurement error in sk and rk , which are in levels, and this is approximately consistent with the observed variances over time (not shown). In practice we found that the model fit much better when this type of measurement error was included. However, a disadvantage of this kind of formulation is that a stationary measurement error on growth rates induces nonstationary behavior of the variances and covariances of the level variables. This fact plus the importance of the permanent firm effects leads us to rewrite our model in terms of first differences (growth rates) before estimating it in the next section.

4. Estimation of the Model

The model we estimated is based on that given in equations (16) through (18) and (23), with the addition of measurement error in the capital stock growth rate as discussed above. In order to eliminate the firm-specific effects and to create a model whose variances and covariances are stationary over time, we rewrite the equations for sk and rk in terms of the growth rates of the variables:

$$(26) \Delta s_{it} = \Delta \ln u_{it} + (\alpha_i / (1 - \alpha_i)) (\ln E_{t-\tau} u_{it} - \ln E_{t-\tau-1} u_{i,t-1}) \\ + \ln \xi_{i,t-\tau} + \sigma \ln \mu_{i,t-\tau} + \text{time dummies}$$

$$(27) \Delta k_{it} = \ln \xi_{it} + \sigma \ln \mu_{it} + (1 / (1 - \alpha_i)) (\ln E_t u_{it+\tau} - \ln E_{t-1} u_{i,t+\tau-1}) \\ + \varepsilon_{it} + \text{time dummies}$$

$$(28) \Delta r_{it} = \ln \xi_{i,t-1} + \ln \mu_{it} + (\sigma - 1) \ln \mu_{i,t-1} + \text{time dummies}$$

The model of Section 2 assumed moving average processes of low order for both $\ln u_{it}$ and $\ln \xi_{it}$ (and therefore for $\ln \mu_{it}$), but our preliminary estimates for $\ln \xi_{it}$ (not shown) suggested an order at least as high as the number of lags of data observed. The coefficients for $\ln u_{it}$ on the other hand appeared to imply a moving average process of order less than or equal to 3. Thus we were lead to model $\ln \xi_{it}$ as an autoregressive process of order one, which we write

$$(29) \ln \xi_{it} = \rho \ln \xi_{i,t-1} + \delta_0 y_{it}$$

where y_{it} is normalized to have a unit variance. Thus, current $\ln \xi_{it}$ is a sufficient statistic for the joint distribution of future $\ln \xi$'s, and since $\ln \mu_{it}$ depends on current and past values of y only through $\ln \xi_{it}$,

we assume that $\ln\mu_{it}$ also evolves as an AR(1) process in y_{it} (and that the processes are proportional):

$$(30) \ln\mu_{it} = \eta_r x_{it} + \dots + \eta_n x_{i,t+r-n} + \gamma_0 \ln\xi_{it} + \text{firm dummies}$$

For the present, we assume that the firm dummies are small after differencing (since the long lag covariances are small). Under these assumptions, the equations of the estimating model may be expressed as functions of the x process, the y process ($\ln\xi_{it}$), and two idiosyncratic sources of error, ε_{it} and z_{it} (assuming that overall year means have been removed from the growth rates, so time dummies are no longer needed).

$$(31) q_{it} = \sigma_x x_{it} + \sigma_y y_{it} + \sigma_z z_{it}$$

$$(32) \Delta s_{it} = \lambda_0 x_{it} + (\lambda_1 - \lambda_0) x_{i,t-1} + \dots + (\lambda_{r-1} - \lambda_{r-2}) x_{i,t-r+1} \\ + (\beta_r - \lambda_{r-1} + \sigma\eta_r) x_{i,t-r} + (\beta_{r+1} - \beta_r + \sigma\eta_{r+1}) x_{i,t-r-1} + \dots \\ + (\beta_n - \beta_{n-1} + \sigma\eta_n) x_{i,t-n} - \beta_n x_{i,t-n-1} + (1 + \sigma\gamma_0) \ln\xi_{i,t-r}$$

$$(33) \Delta k_{it} = (\beta_r + \sigma\eta_r) x_{it} + (\beta_{r+1} - \beta_r + \sigma\eta_{r+1}) x_{i,t-1} + \dots \\ + (\beta_n - \beta_{n-1} + \sigma\eta_n) x_{i,t+r-n} - \beta_n x_{i,t+r-n-1} + (1 + \sigma\gamma_0) \ln\xi_{i,t} \\ + \varepsilon_{it}$$

$$(34) \Delta r_{it} = \eta_r x_{it} + (\eta_{r+1} + (\sigma-1)\eta_r) x_{i,t-1} + \dots + (\eta_n + (\sigma-1)\eta_{n-1}) x_{i,t+r-n} \\ + (\sigma-1)\eta_n x_{i,t+r-n-1} + (1 + \gamma_0\rho + (\sigma-1)\gamma_0) \ln\xi_{i,t-1} + \gamma_0\delta_0 y_{it}$$

where $\beta_r = (1-\alpha)^{-1} \lambda_r$, ... $\beta_n = (1-\alpha)^{-1} \lambda_n$. \ 7

From the fact that the covariogram of $\ln \xi_{it}$ is equal to $\delta_0^2/(1-\rho^2)$, $\rho\delta_0^2/(1-\rho^2)$, $\rho^2\delta_0^2/(1-\rho^2)$..., we can derive the implied covariances of q_{it} , Δs_{it} , Δk_{it} , and Δr_{it} .

Equations (31) through (34) imply variances and covariances for q and the observed growth rates of sales, capital, and R&D, and their lags as functions of the unknown parameters and variances of the factors x , y , z , and ϵ . We use a minimum distance method described in Appendix A to estimate the parameters of the model from the variances and covariances of the data. The method uses the sample fourth moments of the data to form an estimate of the variance of these second moments and hence is robust to non-normality and serial correlation across lags.

We used the three overlapping data samples described earlier for estimation: 496 firms with data from 1978 through 1981 plus four lags, 539 firms with data from 1982 through 1985 plus four lags, and 367 firms with data from 1980 through 1985 plus six lags (as a check on our specification of the order of the x and y processes). All three samples produced qualitatively the same results.

Table 4 shows estimates of our model for gestation lags τ equal to

7. Note that we have now assumed α , the share of physical capital in the total capital stock, to be the same across all firms for tractability. The consequence of this assumption is undoubtedly to underestimate the contribution of the x factor to the variance in the Δs and Δk equations.

2 and 3.⁸ Changing the gestation lag changes the implied restrictions on the coefficients in a non-nested way so it is difficult to test the two models against each other directly, although one can compare their performance under the implied restrictions. In fact, the $r=3$ model, which has fewer parameters, actually fits quite a bit better, in spite of the discussion in the previous section. Since capital is measured at the end of the period in our model, this result implies that the sales capital ratio takes at least two years and closer to three years to respond to demand or profitability shocks completely. It is more consistent with the results of Pakes and Griliches cited earlier, which are based on similar data, than it is with the survey data. This is partly because we are measuring the upper bound on the length of time it takes to adjust the capital stock to shocks, while the survey-style results cited give the average length of time. An additional problem with the $r=2$ results is that the coefficients of the technological accumulation shock process (ξ_{it}) are very poorly determined and unstable across the different samples.

Focusing on the $r=3$ estimates, we can see that although many of the coefficients have plausible values, even this relatively simple structural model is hard to identify in the data. The most disturbing (but interesting) result is that the autoregressive coefficient (ρ) for the technology process is near unity. Only in the longer sample were we

8. Estimates for $r=1$ were uniformly far worse than the others in terms of the trace criterion and frequently failed to converge, so these are not shown.

even able to estimate this coefficient; in the shorter samples it tended to go toward one when it was freed up. Thus we have two processes driving sales, investment, and R&D: a transitory profitability (demand?) shock which is third order moving average--that is, it decays rather quickly--and a technology shock which is nearly permanent. The transitory shock is identified by the fact that it dies out quickly in the investment and R&D processes: because the shock lasts only slightly longer than the gestation lag of capital, capital and R&D respond only to the changes implied by the last couple of lags in the moving average process. The technology shock, on the other hand, is highly persistent and is used by the model to explain the slowly moving individual trends in the growth rates of sales, capital stock, and R&D.

The estimated coefficients on the x process are quite reasonable: all are positive and they imply that news about the profitability of the firm peaks in the first two years and becomes close to zero by year four (for example, β_3 equal to 0.005 in the last column implies λ_3 roughly equal to 0.001, if we assume that the average α is 0.8). The coefficients on y are another matter: again looking at the last column, imagine a technology-associated shock to the firm's value which is equal to the standard deviation of y . That is, the market value of the firm rises 2.5 percent due to this shock. Through δ_0 this will translate into a 21.5 percent fall in the productivity of existing R&D capital, and through γ_0 into a 20 percent increase in the shadow value of the R&D capital to the firm (μ_{it}). This latter effect is not unreasonable (since the sign of the value effects is the same), but the former is a bit mysterious and suggests problems in our specification. It may be plausible that declines in the amount of R&D capital obtained for a given investment be associated with an increase in its value, but it is

not obvious why this should be true.

At the beginning of this investigation, the structural parameter in which we had the most interest was σ , the "adjustment cost" parameter for knowledge capital. σ specifies how R&D investments combine from year to year to produce an intangible capital stock for the firm. Since it is the parameter of a constant returns Cobb-Douglas specification for capital, it should lie between zero and one. Our results for σ turn out to be somewhat inconclusive: the gestation lag assumption makes a large difference in the estimated value of σ , from near zero for $\tau=2$ to unity when $\tau=3$.

This latter estimate (from our preferred set of estimates) may not be as unreasonable as it first appears and can be interpreted in the following way: An estimate of σ equal to unity means that the R&D expenditures made in the most recent year are all that matter for determining G_t , the stock of R&D capital that enters the profit function. Now assume that R&D expenditures are constant in real terms over time within the firm and ξ_t does not vary much. Then the G_t function would simply be a multiple of current R&D expenditures, regardless of the value of σ :

$$(35) \quad G_t = \xi R^\sigma R^{\sigma(1-\sigma)} R^{\sigma(1-\sigma)^2} \dots \\ = \xi R^\sigma \left(\sum_{s=0}^{\infty} (1-\sigma)^s \right) = \xi R$$

If σ were unity, the same property would hold, independent of the pattern of R&D expenditure. Since in fact we know that the R&D for each firm does not vary much over the short time periods in this data (see the evidence in Hall, Griliches, and Hausman 1986 on the small within to total variance ratio for the logarithm of R&D expenditures), we do not have

much power in the data to identify σ , and a value of unity does as well as anything else. Again, we have the finding that current R&D expenditure suffices to describe the level of technological activity of the firm once permanent firm differences have been removed (again, compare the evidence in Hall, Griliches, and Hausman 1986 on the patent productivity of R&D).

The model and estimates here also allow us to say something about the sources of the variance in the growth of sales, capital, and R&D expenditures across firms. The allocation of these variance components implied by the $\tau=2$ and $\tau=3$ estimates on the longer sample is shown in Table 5. As we have mentioned before, the results on q are basically consistent with those of previous authors, although the variance attributed to the profitability shock is a bit low when $\tau=2$. The interesting result is the allocation of the within-firm variance of the growth of capital stock and R&D investment: most of the variance in the growth of net capital stock can be attributed to measurement error (since the associated shock is neither persistent nor correlated with those in the other variables), and what remains is evenly distributed across profitability and technology shocks. On the other hand, most of the variance in the growth of R&D investment comes from the shock we have labelled technological. The identification of this factor comes from the transitory movements linking q , net investment, and R&D but is not associated with immediate changes in the output-capital ratio. It represents the part of changes in expectations about the future of the firm that is orthogonal to changes in current profitability, and it seems natural to label these as shifts in technological opportunity or in the productivity of R&D.

5. Conclusions

In this paper, we have attempted to specify and estimate a structural economic model of a forward-looking manufacturing firm making investment decisions on two interrelated margins, so that we could characterize the observed dynamics of sales, investment, R&D, and market value in an interpretable fashion. Although the project has not been a complete victory for the structural approach, we can draw a few conclusions from the results. First, these manufacturing firms do experience shocks to their environment that do not immediately affect current output but which cause changes in R&D investment and investment in physical capital that are effectively permanent. These shocks are important in explaining the variance in the R&D growth rate, accounting for over 90 percent of it. However, they are associated with at most one-half of one percent of the variance of the year-to-year fluctuations in the stock market value of the firms. The association is initially positive for R&D: a technology-associated shock in the rate of return to the firm is associated with much larger increase in current R&D spending holding capital stock constant and with a slight (but insignificant) decrease in net capital stock. Although there is a large burst in R&D growth rates in the first year (on the order of eight times the growth rate in market value), this is followed by very small declines in the growth rates every year (since the shock is essentially permanent); by about 25 years out, the actual change in R&D spending is zero, and then the growth rate goes negative. But since the longest time period in our data is about 12 years, this result is based on very tenuous extrapolation of the time series process and is not reliable.

The effects of the profitability shock are more sensible: When

there is good news in sales, this immediately causes increases of roughly the same order of magnitude in the net capital stock and R&D spending which will be used to produce output in three years; these increases are transitory in the sense that they work rather quickly. The puzzling fact that the technology shock appears to produce small but permanent declines in the growth of net capital and R&D spending can perhaps be understood by recalling that the shock in question is only that part of technology news which is not correlated with current output. Most of such news may indeed be output related and so we are not identifying it separately.

How do our results compare to those of Lach and Schankerman (1988)? Like them, we have transitory factors idiosyncratic to q and investment (growth in net capital stock). The difference is that we require two factors to explain the correlations between q , net capital growth, and R&D investment growth; one factor is permanent like theirs and the second one is transitory (and larger than the first). Although the response of R&D to the permanent factor (y) is similar in both our estimates, they have investment responding more slowly but positively, while we have net capital stock declining very slowly. Because we have chosen a fundamentally different decomposition of the variables, further comparisons become meaningless. Nevertheless, the interesting fact remains that our model, like theirs, accounts for about 5 percent of the total variance in year-to-year movements in the market value of the firms.

A second and distinct finding is that the upper bound on the gestation lag of physical capital appears to be closer to two and one half years than to one and a half year. Here, however, we must note a

difficulty with our model's treatment of the gestation lag: the assumption of a very precise integer gestation lag leads to clear identification of the length of that lag from the covariances of $s_t - k_{t-r}$ and q_t under the assumption of constant returns. This identification comes from the fact that the current capital-output ratio must contain only the demand news that arrived since that capital was chosen (once the permanent firm effect has been removed). If there is any smearing across years in the purchase and installation of desired capital, this sharp identification will break down, and sk will appear to contain more lags of the demand shock x , no matter what r we choose. This is what happens in our data, and it cannot be accommodated in the existing model, which makes estimation of r very difficult in our framework.

TABLE 1

THE IMPORTANCE OF R&D INVESTMENT IN THE U.S. MANUFACTURING SECTOR

Years	Total Manufacturing Sector					Total Sample		
	P&E Inv.	R&D Inv.	Priv. R&D	R/I	P.R/I	Inv.	R&D	R/I
1951-55	12.02	4.11	2.33	.34	.19			
1956-60	15.39	8.57	3.74	.56	.24			
1961-65	19.09	12.54	5.46	.66	.29			
1966-70	33.84	17.15	8.85	.51	.26			
1971-75	43.50	21.24	13.16	.49	.30	63.0	14.5	.23
1976-80	82.57	34.57	23.07	.42	.28	124.0	23.5	.20
1981-85	130.93	63.87	43.44	.49	.33	173.0	42.0	.24

Variables:

P&E Inv.: investment in plant and equipment by the manufacturing sector, from the Economic Report of the President, 1986, Table B-51.

R&D Inv. and Private R&D Inv.: Research and development expenditures by the industrial sector, total and privately funded, from Science Indicators, 1985.

The sample is all publicly traded manufacturing firms and is described in Hall et al. (1985). For the sample, Inv. is the total capital expenditures (including structures), and R&D is the total R&D expenditure for those firms that report.

All variables are five-year moving averages and are shown in current \$.

TABLE 2

SIMPLE STATISTICS FOR THE VARIABLES

1974-1981: 496 Firms

Statistic	g	Δs	Δk	Δr	Δi	sk	rk
Mean*	0.18	.108	.131	.119	.114	3.48	.045
Minimum	-0.87	-0.93	-0.92	-3.01	-2.51	0.36	.001
Maximum	2.83	0.96	0.82	2.88	2.82	107.	9.16
Variance	.2703	.0205	.0199	.0833	.2894	.366	1.57
Variance ratios:							
within year/tot	.66	.88	.96	.99	.97	.99	1.00
within year-ind/tot	.65	.86	.95	.99	.97	.75	.57
within yr.-firm/tot	.59	.71	.73	.88	.92	.09	.06

1978-1985: 539 Firms

Mean	0.20	.080	.110	.117	.096	3.203	.054
Minimum	-0.76	-0.87	-0.93	-3.01	-2.45	0.28	.001
Maximum	2.94	0.96	0.92	2.71	3.93	53.0	5.62
Variance	.1825	.0251	.0249	.0776	.2991	.381	1.53
Variance ratios:							
within-year/tot	.88	.84	.94	.98	.96	.96	1.00
within year-ind/tot	.87	.83	.92	.97	.96	.74	.55
within yr.-firm/tot	.79	.71	.73	.85	.92	.10	.07

Notes:

* For sk and rk, the geometric means of the ratio variables s/k and r/k are shown. All other variables are growth rates.

The row labelled "within/total" shows the ratio of the within year and industry variance to the total variance. The row labelled "within firm/total" shows the ratio of the within year and firm variance to the total variance.

TABLE 3
COVARIANCES OF THE SAMPLE DATA

367 Firms, 1980-1985 (with lags back to 1974)

Auto Covariances

Lag	q_{it}	$s_{it} - k_{it-1}$	Δk_{it}	$r_{it} - k_{it-1}$
0	.1281 (.0070)	.227 (.009)	.0223 (.0013)	.602 (.021)
-1	.0013 (.0037)	.208 (.009)	.0044 (.0005)	.558 (.019)
-2	-.0077 (.0032)	.195 (.009)	.0013 (.0005)	.529 (.019)
-3	-.0049 (.0034)	.184 (.008)	-.0004 (.0005)	.512 (.018)
-4	-.0042 (.0037)	.180 (.008)	.0006 (.0005)	.502 (.018)
-5	-.0057 (.0034)	.179 (.008)	.0005 (.0005)	.489 (.018)
-6	.0025 (.0033)	.176 (.008)	.0010 (.0005)	.473 (.017)
First lag = 0 or const.	1	3	2	6

Notes:

Covariances are estimated by methods due to MaCurdy (1985). Standard errors that are robust to heteroskedasticity are shown in parentheses.

Tests were conducted for the hypothesis that the lagged covariances were zero (columns 1 and 3) or equal (columns 2 and 4). The first lag for which the test was accepted is shown in the table.

TABLE 3 (continued)

Cross Covariances						
Lag ⁺	q,sk	q,Δk	q,rk	sk,Δk	sk,rk	Δk,rk
-6	.0080	-.0001	-.0064	-.0021	.117	.0073 *
-5	.0055	-.0015	-.0076	-.0024	.121	.0070 *
-4	.0034	-.0010	-.0060	-.0029	.124	.0081 *
-3	.0019	-.0004	-.0059	-.0051 *	.124	.0070 *
-2	.0032	-.0053 *	-.0132 *	-.0037	.130 *	.0094 *
-1	.0208 *	.0017	-.0006	-.0021	.137 *	.0135 *
0	.0296 *	.0116 *	.0128 *	.0160 *	.153 *	.0188 *
+1	.0206 *	.0080 *	.0111 *	.0114 *	.138 *	-.0003
+2	.0173 *	.0061 *	.0100 *	.0087 *	.128 *	.0016
+3	.0081	.0029	.0017	.0054 *	.120	.0016
+4	.0086	.0006	.0094	.0046 *	.117	.0047
+5	.0063	-.0005	.0042	.0049 *	.116	.0036
+6	.0085	.0004	.0041	.0055 *	.111	.0023

Notes:

⁺ The lag shown is for the second variable in each column. That is, the first entry in the first column is the covariance of q and sk(-6).

^{*} The starred covariances are significantly different from zero (a "robust" t-statistic greater than 2). In column 5 (rk,sk), the t-statistic is for a difference from a "permanent" covariance of around 0.12.

TABLE 4: DYNAMIC FACTOR MODEL ESTIMATES

Sample*	Gestation lag(τ) = 2			Gestation lag(τ) = 3		
	78-81	82-85	80-85	78-81	82-85	80-85
σ	-.011(.310)	.286(.098)	.081(.141)	1.08(.35)	1.08(.14)	1.01(.11)
λ_0	.084(.012)	.071(.005)	.063(.004)	.044(.021)	.118(.003)	.111(.003)
λ_1	.159(.011)	.109(.004)	.099(.004)	.178(.012)	.105(.005)	.099(.005)
λ_2	--	--	--	.099(.020)	.068(.005)	.072(.005)
β_2	.027(.006)	.026(.003)	.025(.003)	--	--	--
β_3	.021(.008)	.017(.003)	.027(.003)	.012(.007)	.002(.003)	.005(.002)
η_2	.017(.005)	.019(.007)	.008(.006)	--	--	--
η_3	.017(.005)	.042(.007)	.027(.007)	.019(.005)	.040(.006)	.046(.006)
γ_0	27.(225.)	-3.31(1.07)	-9.41(12.43)	-.902(.208)	-.909(.119)	-.977(.105)
δ_0	.006(.053)	-.063(.021)	-.022(.029)	-.193(.063)	-.232(.032)	-.215(.025)
$\sigma(x)$.016(.019)	.070(.011)	.037(.010)	.025(.019)	.100(.010)	.071(.009)
$\sigma(y)$.049(.013)	.035(.010)	.022(.008)	.043(.012)	.025(.009)	.025(.008)
$\sigma(z)$.406(.012)	.288(.011)	.296(.009)	.411(.012)	.297(.011)	.303(.008)
$\sigma(\epsilon)$.184(.010)	.121(.004)	.108(.003)	.167(.010)	.126(.004)	.114(.003)
ρ^{**}	.995	.995	.988(.006)	.995	.995	.998(.002)
Trace	218.7	372.5	396.2	213.2	332.7	367.1
#params	16	16	17	15	15	16
#obs.	1984	2156	2202	1984	2156	2202

Notes to Table 4:

x and y are the two common factors. z is the idiosyncratic factor in q, and ϵ is the idiosyncratic factor in Δk . All estimates were obtained by the generalized method of moments (see Appendix A for a description).

* The samples are defined in Tables 2 and 3. The 78-81 and 82-85 samples have four lags, while the 80-85 sample has six.

** Estimates of ρ with no standard error were obtained by searching and then fixing ρ , owing to convergence difficulties in the shorter samples.

TABLE 5
 VARIANCE DECOMPOSITION
 367 Firms, 1980-1985

Var	r	Profitability Shock		Technology Shock		Measurement Error	
		Variance	%	Variance	%	Variance	%
q	2	.0014	1.5	.0005	0.5	.0876	97.9
	3	.0050	5.2	.0006	0.6	.0918	94.2
Δs	2	.0114	90.8	.0012	9.2	--	
	3	.0254	92.6	.0020	7.4	--	
Δk	2	.0014	9.8	.0012	8.1	.0117	82.0
	3	.0027	15.1	.0020	11.4	.0130	73.5
Δr	2	.0011	2.3	.0454	97.7	--	
	3	.0021	4.3	.0468	95.7	--	

The columns labelled percent show the fraction of the total variance of year-to-year growth rates which is accounted for by the factor (x, y, and the measurement errors z and ϵ in the last column).

APPENDIX A
MINIMUM DISTANCE ESTIMATION OF NONLINEAR
MODELS OF A COVARIANCE MATRIX

The MOMENTS User's Manual (Hall 1987) describes how to estimate models of the form

$$A(\theta)Y = B(\theta)X$$

where $X \sim MVN(0, \Sigma)$ and Y is an observed data matrix (T by N), by the method of maximum likelihood. If the distribution of X is truly multivariate normal, this method produces consistent, efficient, and asymptotically normal estimates. However, if the normalizing assumption fails, we might wish to use a more robust method such as minimum distance. Another reason why a different estimating method might be preferred is the following: The ML method maximizes a likelihood of the form

$$\log L \sim \log |\Omega(\theta)| + \text{tr}[M \Omega^{-1}(\theta)]$$

where $\Omega(\theta) = A^{-1}(\theta) B(\theta) \Sigma(\theta) B'(\theta) A^{-1}(\theta)$, and M is the sample covariance of Y . If N is large (greater than about 10), the determinant and inversion routines necessary to evaluate the likelihood may easily encounter numerical difficulties. Using a minimum distance method overcomes this drawback also. Finally, the minimum distance method of estimating this model may be easily implemented in TSP and, in fact, involves only SUR estimation with a single observation, once the correct covariance (weighting) matrix has been formed. This property arises

from the fact that $M = (1/T)Y'Y$ is a sufficient statistic for the data under normality and hence may be used to estimate the second moments of the model, although fourth moments are required to obtain correct standard errors. The methodology used here is due to Chamberlain (1982) and MaCurdy (1981, 1985) and is related to that of Hansen (1982).

To use minimum distance on this model, we set up a series of equations that are nonlinear in the parameters, using the theoretical model for Y as a guide:

$$\text{vec}E[Y'Y] = \text{vec}\Omega(\theta)$$

Then substitute the sample covariances of Y for the expectation (population moment):

$$m_t = \text{vec}(Y_t Y_t') - \text{vec}E[Y'Y] + \xi_t = \text{vec}\Omega(\theta) + \xi_t$$

Y_t is an N by 1 vector, and m_t , $\text{vec}\Omega$, and ξ_t are $P=N(N+1)/2$ by 1 vectors. The properties of ξ_t are that it is independent and identically distributed with mean zero and covariance matrix Ξ . The minimum distance method suggests that we minimize $m_t - \text{vec}\Omega(\theta)$ in the metric of its inverse variance, which is Ξ^{-1} . A consistent estimate of the variance of ξ_t may be obtained from the sample variance of m_t :

$$\hat{\Xi} = (1/T) \sum_{t=1}^T (m_t - \text{vec}M)(m_t - \text{vec}M)'$$

where $\hat{\Xi}$ is a P by P matrix and $E(\hat{\Xi}) = \Xi$.

Once $\hat{\Xi}$ is obtained (which may be done once for any particular dataset Y), the parameter vector θ may be estimated consistently and asymptotically efficiently by

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} (1/T) \sum_{t=1}^T (m_t - \operatorname{vec}\Omega(\theta))' \hat{\Xi}^{-1} (m_t - \operatorname{vec}\Omega(\theta))$$

The FOC for this problem is

$$(\operatorname{vec}M - \operatorname{vec}\Omega(\theta))' \hat{\Xi}^{-1} (\partial \operatorname{vec}\Omega(\theta) / \partial \theta) = 0$$

This is also the FOC for the problem

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} (\operatorname{vec}M - \operatorname{vec}\Omega(\theta))' \hat{\Xi}^{-1} (\operatorname{vec}M - \operatorname{vec}\Omega(\theta))$$

So in this case the method of moments and minimum distance will coincide. Note that the second minimization depends only on the sample covariance matrix $\operatorname{vec}M$ and its estimated covariance $\hat{\Xi}$. Depending on the order of magnitude of the problem, this can create a considerable savings in computation time. Note also that the trace criterion for the second problem

$$S = (\operatorname{vec}M - \operatorname{vec}\Omega(\theta))' \hat{\Xi}^{-1} (\operatorname{vec}M - \operatorname{vec}\Omega(\theta))$$

is precisely equal to the trace criterion for the first problem, multiplied by T , the number of observations. This implies that the standard errors resulting from minimizing S will need to be adjusted by a \sqrt{T} factor so they have the correct order of magnitude. This is best done by dividing the estimated $\hat{\Xi}$ by T before estimation.

REFERENCES

- Bernstein, Jeffrey I., and M. Ishaq Nadiri. 1986. "Research and Development and Intraindustry Spillovers: An Empirical Application of Dynamic Duality." New York University. Photocopied.
- Bureau of Labor Statistics. 1979. Capital Stock Estimates for I/O Industries: Methods and Data, Bulletin No. 2034. Washington, D.C.: Government Printing Office.
- Chirinko, Robert. 1986. "Will the Neoclassical Theory of Investment Please Rise?: The General Structure of Investment Models and Their Implications for Tax Policy." University of Chicago and National Bureau of Economic Research. Photocopied.
- Chirinko, Robert. 1987. "Investment, Tobin's Q, and Multiple Capital Inputs." University of Chicago and NBER. Photocopied.
- Duffie, Darrell. 1988. Security Markets: Stochastic Models. New York: Academic Press.
- Epstein, L. G., and M. G. S. Denny. 1983. "The Multivariate Flexible Accelerator Model: Its Empirical Restrictions and an Application to U.S. Manufacturing." Econometrica 51: 647-674.
- Fama, Eugene F., and Kenneth R. French. 1987. "Permanent and Temporary Components of Stock Prices." Chicago: Center for Research in Security Prices Working Paper Number 178 (July).
- Griliches, Zvi. 1981. "Market Value, R&D, and Patents." Economics Letters 7.
- Griliches, Zvi, Bronwyn H. Hall, and Ariel Pakes. 1986. "Is There a Second (Technological Opportunity) Factor?" NBER. Photocopied.
- Hall, Bronwyn H. 1986. MOMENTS User's Manual. Palo Alto, CA: TSP International.
- _____. 1987. "The Relationship Between Firm Size and Firm Growth in the U.S. Manufacturing Sector." Journal of Industrial Economics 35: 583-606.
- _____. 1988. "Research and Development Investment and the Evolution of the U.S. Manufacturing Sector: Econometric Studies at the Firm Level." Ph.D. diss., Stanford University.
- Hall, Bronwyn H., Clint Cummins, Elizabeth S. Laderman, and Joy A. Mundy. 1985; revised 1988. "The R&D Master File: Documentation." NBER. Photocopied.

- Hall, Bronwyn H., Zvi Griliches, and Jerry A. Hausman. 1986. "Patents and R&D: Is There a Lag?" International Economic Review 27: 77-106.
- Hall, Robert E. 1977. "Investment, Interest Rates, and the Effect of Stabilization Policies." Brookings Papers on Economic Activity.
- Hayashi, Fumio. 1982. "Tobin's Marginal q and Average q: A Neoclassical Interpretation." Econometrica 50: 556-69.
- Kydland, Finn E., and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." Econometrica 50(6): 1345-71.
- Lach, Saul, and Mark Schankerman. 1988. "The Interaction Between Capital Investment and R&D in Science-Based Firms." Journal of Political Economy, forthcoming.
- Lucas, Robert E. 1967. "Optimal Investment Policy and the Flexible Accelerator." International Economic Review 8: 78-85.
- MaCurdy, Thomas E. 1985. "A Guide to Applying Time Series Models to Panel Data." Stanford University. Photocopied.
- Mairesse, Jacques. 1987. "Le lois de la production ne sont plus ce qu'elles etaient: une introduction a l'econometric des panels." Revue economique 1(janvier 1988): 225-71.
- Mairesse, Jacques, and Alan K. Siu. 1984. "An Extended Accelerator Model of R&D and Physical Investment," in R&D, Patents, and Productivity, Zvi Griliches, ed. Chicago: University of Chicago Press.
- National Science Board. 1985. Science Indicators: The 1985 Report. Washington, D.C.: U.S. Government Printing Office.
- Pakes, Ariel. 1985. "On Patents, R&D, and the Stock Market Rate of Return." Journal of Political Economy 93(2): 390-409.
- Pakes, Ariel, and Zvi Griliches. 1984. "Estimating Distributed Lags in Short Panels with an Application to the Specification of Depreciation Patterns and Capital Stock Constructs." Review of Economic Studies Vol. LI(2), No. 165.
- Poterba, James M., and Lawrence H. Summers. 1987. "Mean Reversion in Stock Prices: Evidence and Implications." NBER Working Paper No. 2343.
- Sargent, Thomas. 1986. Dynamic Macroeconomic Theory. New York: Academic Press.
- Standard and Poor. 1985. Compustat Annual Industrial Data File. New York: Standard and Poor Corporation.
- U.S. President. 1986. Economic Report of the President. Washington, D.C.: U.S. Government Printing Office.