

NBER WORKING PAPER SERIES

IS THERE A LAG?

Bronwyn H. Hall

Zvi Griliches

Jerry A. Hausman

Working Paper No. 1454

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1984

The research reported here is part of the NBER's research program in Productivity and project in Productivity (World Economy). Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

Is There a Lag?

ABSTRACT

This paper extends earlier work on the R&D to patents relationship (Pakes-Griliches 1980, and Hausman, Hall, and Griliches, 1984) to a larger but shorter panel of firms. The focus of the paper is on solving a number of econometric problems associated with the discreteness of the dependent variable and the shortness of the panel in the time dimension. We compare weighted nonlinear least squares as well as Poisson-type models as solutions to the former problem. In attempting to estimate a lag structure on R&D in the absence of a sufficient history of the variable, we take two approaches: first, we use the conditional version of the negative binomial model, and second, we estimate the R&D variable itself as a low order stochastic process and use this information to control for unobserved R&D. R&D itself turns out to be fairly well approximated by a random walk. Neither approach yields strong evidence of a long lag. The available sample, though numerically large, turns out not to be particularly informative on this question. It does reconfirm, however, a significant effect of R&D on patenting (with most of it occurring in the first year) and the presence of rather wide and semi-permanent differences among firms in their patenting policies.

Bronwyn H. Hall
National Bureau of
Economic Research
204 Junipero Serra Blvd.
Palo Alto, CA 94305

Zvi Griliches
National Bureau of
Economic Research
1050 Massachusetts Ave.
Cambridge, MA 02138

Jerry A. Hausman
Department of Economic
M.I.T.
E52-271A
Cambridge, MA 02139

Bronwyn H. Hall, Zvi Griliches, and Jerry A. Hausman

This paper analyzes the relationship between patenting and research and development activity at the firm level by the U. S. manufacturing sector during the 1970's. Previous work by Pakes and Griliches (1980), which looked at a subset of the firms considered here, was the first attempt to use the patenting and R&D behavior of firms over time both to control for individual firm effects and to try to learn something about lags in the productivity of R&D. The present study extends their sample to 1979 and covers almost all of the firms doing appreciable amounts of R&D in the manufacturing sector. In attempting to characterize the lag structure of the patents-R&D relationship, a number of econometric problems arising from the panel nature of the data and from the measurement of the dependent variable have to be solved or at least considered in interpreting the results.

The basic model underlying this analysis has been described elsewhere (Pakes and Griliches 1980, Hausman, Hall, and Griliches 1984) and will only be sketched here. The annual research and development expenditures of a firm are considered to be investments which add to a firm's stock of knowledge. This stock of knowledge is depreciating over time so that the contribution of older R&D investment becomes less valuable as time passes. The aim of the study is to use patent applications in any given year as an indicator of the "value" of the underlying stock of knowledge, and to infer from the lag distribution on past R&D something about gestation lags in

knowledge production. It is important to note at the outset that there are difficulties with this approach: the first is that patents are not the only output of R&D - they measure only a fraction of this output, and the fraction may vary considerably over industry and possibly also over time. Controlling for differences in the firms' propensity to patent (conditional estimation) as well as including an overall effect for each year are partial answers to this. A second and related problem with the existing data is that most of the information on the question will come from relative changes in the two variables over time within the firm; if these changes are contaminated by measurement error or they are very small, the lag structure will be extremely difficult to discern.

A third problem has to do with the economic value of the patents themselves. Researchers such as Mansfield (1977) and Taylor and Silberston (1973) have suggested that the existence of the patent system may be a relatively unimportant factor in the research and development strategy of some firms. There is also a growing body of evidence (Grabowski 1983, Pakes 1984, Schankerman and Pakes 1984) showing that a large fraction of patents granted are "worthless" or become worthless in a short period of time. This paper has very little to say on this range of topics; we observe the fact that firms do take out patents which are related to the output of their research and development laboratories (and other activities of the firm in the R&D area) and that therefore, patents can be used as an indicator of this activity in the aggregate even though the information conveyed by an individual patent may be very small.

In discussing an earlier version of this paper, Stoneman (1983) argued strongly that patents are an input to the R&D process rather than an output. That is, the patent application occurs at an early point in the development process and most of the expenditures that would be associated

with it occur after the application is made. With this data we are in a position to investigate this question of timing, and find that the evidence for it is relatively weak, at least in aggregate firm behavior. The strongest thing one can say is that R&D and patents appear to be dominated by a contemporaneous relationship, rather than leads or lags.

The earlier work in this area also found a strong contemporaneous effect of R&D on patents but was inconclusive as to whether there was a significant lagged effect. Pakes and Griliches, using the standard fixed effects model, found evidence of a lag truncation effect in the distributed lag of patents on R&D. That is, when they controlled for permanent differences across firms in the propensity to patent, the estimated coefficient on the last lag of R&D which they considered (R&D expenditures of four years prior) was significantly higher than the coefficients of more recent R&D. Hausman, Hall, and Griliches used a different functional form (which took the discreteness of the patent data explicitly into account) and found similar results for the random (uncorrelated) effects model but not in their conditional fixed effects version. When they conditioned their estimates on the total number of patents received during the whole period, no coefficients except for the contemporaneous R&D variable were statistically significant either in the Poisson or negative binomial version.

These previous studies both used samples of about 120 firms with seven to eight years of patent data and twelve to thirteen years of R&D data. For the current study, although we have about fourteen years of patent data from 1966 to 1979, we have only eight years of R&D data for about 650 firms, with an additional two years (1970 and 1971) for half the firms. This fact constrains our ability to look for very long lag effects,

especially since we cannot distinguish easily between permanent differences across firms in the propensity to patent and effects due to the unobserved past R&D history. We discuss this issue at greater length in the body of the paper.

The other problem we have to deal with is the specification of the error term in our model. The difficulty arises from two somewhat related causes: the presence of a large number of zeroes in our dependent variable, the number of patents applied for by a firm in a particular year, and the large size range of the firms in our sample. A previous paper (Bound et al 1984) which analyzed a large cross section sample, including the firms under consideration here, demonstrated that estimates were quite sensitive to the specification of the distribution of the error term. Since most of the estimators used were consistent, this can be construed as informal evidence of misspecification of the underlying model, possibly due to nonlinearity in the relationship of log patents to log R&D or to heteroskedasticity which is size-related. We have taken two approaches in obtaining our estimates in this paper: the first uses a nonlinear least squares specification with additive errors; for these estimates we are able to obtain robust standard errors which are correct in the presence of arbitrary heteroskedasticity, including year-to-year correlation within firms. For this version of the model however, we are unable to obtain conditional estimates, due to its intrinsic nonlinearity and the shortness of our panel. The second approach uses an explicit stochastic specification for the patents variable, that it follows a Poisson or negative binomial distribution, which enables us to obtain conditional estimates of the slope parameters, but at the price of a distributional assumption which may not hold.

The plan of this paper is the following: first we discuss the

derivation of our dataset and look at the properties of our independent variable, R&D expenditures. Then we present some estimates of the basic patents-R&D relationship, followed by a discussion of the biases which may be present in the cross section estimation of this relationship. Finally we present conditional estimates of our model in an attempt to control for some of these biases and we conclude with a brief discussion of what we can learn from these data.

I. Data

The data we use are an extract from a larger and longer panel of firms in U. S. manufacturing drawn from the Compustat (Standard and Poor 1980). This dataset was assembled and combined with patent data from the Office of Technology Assessment and Forecasting at the NBER and is described in Bound et al (1984) and Cummins, Hall, and Laderman (1982). The original universe from which our sample comes consisted of approximately 2700 firms in the manufacturing sector in 1976, and included almost all of the firms which report R&D expenditures to the Bureau of Census-NSF R&D survey.

Our sample of firms was chosen from this universe by requiring that data on sales, gross capital, market value (value of common stock), and R & D be available for all years from 1972 through 1979 with no large jumps during that period. A jump is defined as an increase in capital stock or employment of more than 100 per cent or a decrease of more than 50 per cent. This test was not applied unless the change in employment was greater than 500 employees or the change in capital stock was greater than two million dollars. We also removed six firms which had abnormally small R & D values (less than \$10,000) in one of the years. The number of firms remaining in the sample after these cuts was 642, with a size distribution heavily tilted toward the larger firms in our original universe. Table 1 shows the selectivity of this sample with respect to size and indicates that although we have only a quarter of our original sample of firms, most of those lost were either smaller or were not R & D doing (and reporting) firms. Our coverage of the larger R & D firms is almost complete, and our sample includes 90 per cent of the R&D dollars expended by the manufacturing sector in 1976.

Table 2 exhibits the characteristics of our remaining sample of firms,

both the 642 firms with R & D between 72 and 79 and a subset of firms with a longer R & D history back to 1970. Quantiles are shown in order to give some indication of the skewness of the data: for example, median sales for this sample in 1976 were 182 million dollars, while mean sales were 1.06 billion dollars. The subset of firms with a longer R&D history consists of somewhat larger firms and is more heavily tilted toward the scientific sector. Even for this sample of relatively R&D-intensive firms, we find that over 20 per cent of the firms did not apply for patents in 1976 and that more than half applied for less than five. This confirms our impression that the patents variable in these data must be treated in a way which correctly reflects its relative imprecision at small values. Previous experience with estimation of the patents equation in the cross section (Bound et al 1984) has shown us that slope coefficient estimates may not be robust to changes in the way in which we specify the error in the equation (and the weighting which is implied by such specification).

In the later sections of this paper we look at this question again in an effort both to draw some robust conclusions from the data and to understand the reasons for the unstable coefficients. However, first we take a closer look at the process generating our independent variable, R&D, since the way it evolves has important implications for our ability to identify the true coefficients in the presence of lag truncation and firm effects.

To study the stochastic process for R&D, we use a procedure due to MaCurdy (1983) for computing the sample autocorrelation and partial autocorrelation functions. Basically, this method treats each firm as an independent draw on a time series process, so that we have 642 observations on each autocorrelation and partial autocorrelation. Since these can be estimated for each year, it is not necessary to impose covariance

stationarity. This model allows each firm to have its own mean, but assumes that the within firm variance is the same across firms, which is not an unreasonable assumption for the logarithmic form of our data. We can test for stationarity of the variances across time, and in our data we find $F(7,5128) = 6.3$ which is significant at the conventional five per cent level, but insignificant using the large sample critical values due to Leamer (1978). Accordingly, we impose stationarity in order to compute the autocorrelation and partial autocorrelation functions.

The results are shown in Table 3: the autocorrelations are all above .95, and show a very small decrease at longer lags, while the partial autocorrelations are essentially zero after the second lag, with the second lag equal to .048(.034). This is strong evidence for a low order AR process; in fact, it is difficult to reject a random walk, although there is a hint of a small positive coefficient on the second lag and a first lag coefficient of slightly less than one. In order to check this result, we compute the AR regression itself and display the results in the second part of Table 3. The standard errors shown are heteroskedastic-consistent estimates, although they are in fact almost the same as conventional estimates, which is evidence that the assumption of constant within-firm variance is not a bad one. The basic result is that the AR2 specification can be accepted at conventional significance levels, and that the process is in fact very close to a random walk.

The last two columns of this table provides a partial answer to the question of whether patents can be viewed as an input to the R&D process in this data. We use a simple version of a Granger causality test: with two lags of R&D used to predict the current level of R&D, we include contemporaneous and lagged log patents in the regression to see if they help in predicting R&D in the presence of its past history. The coefficient on contemporaneous log patents is significant ($t=3.1$), but

lagged patents are of no help in predicting future R&D, even if we leave contemporaneous patents out of the equation (last column). We tentatively conclude that there may be simultaneous movements in patents and R&D, but there is little evidence that past success in patenting leads to an increase in a firm's future R&D program above and beyond that implied by its R&D history.¹ We should qualify this result by noting that there is a considerably lower signal to noise ratio in the patents variable than in the R&D variable, both because of the skewness in patent values mentioned earlier and because it is intrinsically an integer variable. This has been well documented by Pakes and Griliches (1980) and Pakes (1984). Since R&D is highly correlated over time, it will be difficult to discern the independent contribution of patents to the R&D program in the presence of this noise.

II. Basic Results

In earlier work with the 1976 cross section of these firms, Bound *et al* found that estimates of the elasticity of patenting with respect to R&D at the average R&D in the sample varied from 0.35 to 2, depending on the choice of specification: log linear, Poisson, negative binomial, or nonlinear least squares. This difference was greatly attenuated when the firms were divided into two groups, those with R & D budgets larger than two million dollars and those with smaller R & D budgets. In the present paper, the problem is not as severe since our sample is more heavily weighted toward the firms in the larger group (approximately 50 per cent have R & D greater than two million, rather than 20 per cent), but it still persists and affects our estimates of the lag distribution.

In Table 4, we look at the differences in estimates of our basic model which are implied by differing specifications of the error structure. The model is

$$(1) \quad E(p_{it} | R_{it}, R_{i,t-1}, \dots, s_i, t) = \exp\left[\sum_{\tau} (\beta_{\tau} \log R_{i,t-\tau}) + \delta' s_i + \alpha_t\right]$$

where s_i are the observed firm characteristics (size, as measured by the log of gross plant in 1972, and a dummy for the scientific sector). Our sample is the 642 firms estimated over the years 1975 to 1979 so that we can include three lags on R&D, which yields a total of 3210 observations on the dependent variable, patents. The first column shows the nonlinear least squares estimates of the parameters, which are obtained by assuming an additive and homoskedastic error in equation (1). These estimates are consistent for the underlying coefficients, provided the model is correctly specified. The standard errors shown are robust to heteroskedasticity of the disturbances; they are computed using the formulas due to Eicker-White-

Chamberlain, and allow both for differing variances across firms and arbitrary serial correlation over time within firms.

The next two columns of Table 4 give the results of estimating the Poisson and negative binomial versions of our models. The advantage of these models is that they take into account the non-negativity and discreteness of our data. Moreover, in the next section of this paper we will see that the conditional versions of these models allow us to estimate a fixed effects model, something that we cannot do easily with the nonlinear least squares version of the model. On the other hand, these models require us to be explicit about the exact form of the distribution from which the disturbance is drawn, and may produce inconsistent estimates if the distribution is not correct (Gourieroux, Montfort, and Trognon 1984).

The Poisson and negative binomial models were described in detail in our earlier paper (Hausman, Hall, and Griliches 1984) and we shall summarize only their main features here. The log likelihood function for the Poisson model is given by

$$(2) \log L = \sum_{i=1}^N \sum_{t=1}^T [y_{it}! - e^{-X_{it}\beta} + y_{it}X_{it}\beta]$$

where y_{it} is the observed number of patent application for a firm in a year and the X_{it} are the independent variables, R&D and firm characteristics. Estimates obtained with this model differ from the nonlinear least squares estimates primarily by the weighting scheme used. The NLS estimates are unweighted, implicitly weighting the numerically larger deviations of the larger firms more than those of the small firms. The Poisson model assumes that the variance of the disturbances is proportional to the expected value

of the patents and weights the observations accordingly. The negative binomial model generalizes the Poisson model by allowing for an additional source of variance above that due to pure sampling error. The logarithm of the likelihood for this model is

$$(3) \quad \log L = \sum_{i=1}^N \sum_{t=1}^T \log [(\lambda_{it} + y_{it})^{-\lambda_{it} - y_{it}}] - \log [(\lambda_{it})^{-\lambda_{it}}] \\ - \log [(y_{it} + 1)^{y_{it}} + \lambda_{it} \log(\delta) - (\lambda_{it} + y_{it}) \log(1 + \delta)]$$

where $\lambda_{it} = \exp(X_{it}\beta)$ and δ is the variance parameter ($Vy_{it} = \exp(X_{it}\beta)(1 + \delta)/\delta$). We estimate both of these models by standard maximum likelihood techniques.

Finally, in the last column of Table 4, we show estimates computed using the quasi-generalized pseudo maximum likelihood (QGPML) method of Gonrieroux, Montfort and Trognon (1984, henceforth referred to as GMT). This estimator is based on the following idea: suppressing the t subscript for the moment, we assume that the true model for patents is

$$y_i = \exp[X_i\beta + e_i]$$

where $\exp(e_i)$ is a multiplicative disturbance drawn from an unspecified distribution. If a constant term is included as one of the X 's, we can assume $E[\exp(e_i)] = 1$ and $V[\exp(e_i)] = \eta^2$. Then the expected number of patent applications conditional on the X 's is $\exp(X_i\beta)$ and the variance is $\exp(X_i\beta) + \eta^2 \exp(2X_i\beta)$. That is, the variance equals the mean plus a parameter times the mean squared. We can obtain consistent estimates of the parameters β for this model using nonlinear least squares, use these to estimate η^2 , form a vector of GLS weights which are proportional to the variance of the dependent variable:

$$w_i = \hat{y}_i + \frac{\hat{\sigma}_i^2}{\hat{y}_i^2}$$

and use these weights to obtain more efficient estimates of the β 's. The formula for the variance of these estimates is given in GMT, and is a special case of the Eicker-White-Chamberlain formula with known weights.

Since all of these models differ only by their distributional assumptions and not by specification of the expected value, they should all yield roughly the same results unless the basic specification of the equation is wrong. In fact, it can be shown (see GMT) that both the NLS estimates and the Poisson estimates of the parameters are consistent if we have correctly specified the expectation in (1) and the true conditional distribution satisfies certain regularity conditions given in their article. Because the estimators make different assumptions about the error structure they do yield different estimates of the standard errors, even in the case of similar coefficients. In this respect, the nonlinear least squares estimates, weighted or unweighted, are the most robust, since we have computed standard errors which allow for unknown heteroskedasticity. It can be seen from the table that in return for making a relatively mild assumption about the form of the variance (that it is increasing in the mean and mean squared), we obtain a considerable increase in the precision of our estimates (compare columns 1 and 4).

In Table 4 we see that the results of using the four different estimators on this dataset are qualitatively the same, although there is a substantial increase in the coefficient on contemporaneous R&D as we move from nonlinear least squares to weighted nonlinear least squares (GMT). Since the estimators in columns 1, 2, and 4 are consistent if we have the correct model, but are estimated with different weighting schemes,² one possible explanation of the differences in coefficients, particularly the

sum, may be that the relationship between patents and R&D is not stable across the firms in our sample. An indication that this is a problem is provided by the substantial increase in standard errors when we use robust estimates.

The dimension along which the weighting schemes vary is basically related to the size of the firms in the sample. Therefore we partitioned them into roughly equal groups: those with assets (book value of net plant) less than 25 million dollars in 1972 and those with assets greater than 25 million. We then estimated the same model as column (4) of Table 4 on the two groups separately. Although the total R&D effect was the same for the two groups (.66), it was distributed differently across the lags: 0.31, .11, .14, .11 for the smaller firms and .32, .01, .02, .31 for the larger firms. This suggests that the maintained hypothesis of a roughly constant lag structure across the firms may be one reason for the apparent instability of our results. Unfortunately it is not possible with this dataset to construct a more detailed behavioral model which is capable of accounting for different lag structures across firms. We can only suggest areas for future investigation.

III. Correlated Effects or Lag Truncation Bias?

In obtaining the results shown in Table 4, there was no attempt to control for permanent differences in the propensity to patent across firms, except for the firm size variable and a dummy for the scientific sector. We expect that these differences may bias our estimate of the R&D coefficients if they are correlated with the R&D variables. All of the estimates, except possibly those for the smaller firms, exhibited evidence of a u-shaped lag structure, with the first and last coefficients being larger than those in the middle. The large coefficient on the last lag could be due to the correlation of the last R&D variable with earlier left-out R&D, but it turns out that under reasonable assumptions on the R&D process itself, it could also be caused by (correlated) fixed effects.

Assume that the log deflated R&D variable itself follows a first order autoregressive process:

$$R_t = \gamma R_{t-1} + e_t, \quad e_t \text{ a white noise process}$$

Then the autocorrelation coefficients for the R&D process are $(1, \gamma, \gamma^2, \dots)$ and the partial autocorrelation coefficients are $(\gamma, 0, 0, \dots)$. We have seen that just such a pattern is consistent with our R&D data. If we maintain the hypothesis that R&D follows an AR1 process, we can compute the bias formula for the coefficients on R&D in the presence of two types of omitted variable: 1) pre-sample R&D, which is correlated with in-sample R&D in a geometrically decaying way, and 2) a permanent fixed effect which has the same correlation with R&D in all periods. In the first case, the bias formula for $\beta = (\beta_0, \beta_1, \dots, \beta_k)$, where there are k lags in the regression, is

$$(4) \text{ plim } \beta = \beta + \frac{\beta_{k+1}}{\sigma_x^2(1-\gamma^2)} \begin{bmatrix} 1 & -\gamma & \gamma^2 & & & & \\ -\gamma & 1+\gamma^2 & -\gamma & & & & \\ & \cdot & \cdot & & & & \\ & & 0 & & & & \\ & & & & -\gamma & 1 & \gamma \\ & & & & & & \cdot \end{bmatrix} \begin{bmatrix} \gamma^k \sigma_x^2 \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \sigma_x^2 \\ \cdot \end{bmatrix}$$

$$= \beta + \beta_{k+1} \begin{bmatrix} 0 \\ \cdot \\ \cdot \\ \cdot \\ 0 \\ \gamma \end{bmatrix}$$

If only the most recent pre-sample R&D belongs in the equation, the last coefficient is biased upward by $\beta_{k+1}\gamma$ where β_{k+1} is the coefficient on the R&D one period before the sample begins. In general, the bias on the last coefficient will be equal to $\sum \beta_{k+\tau}\gamma^\tau$. In our case, since γ is close to unity, we expect the last coefficient to be roughly equal to the sum of the lag coefficients for all the earlier R&D plus its own coefficient.

On the other hand, if we assume a fixed effect μ_i for each firm has been omitted, we obtain the formula

$$(5) \text{ plim } \beta = \beta + \frac{1}{\sigma_x^2(1-\gamma^2)} \begin{bmatrix} 1 & -\gamma & \gamma^2 & & & & \\ -\gamma & 1+\gamma^2 & -\gamma & & & & \\ & \cdot & \cdot & & & & \\ & & 0 & & & & \\ & & & & -\gamma & 1 & \\ & & & & & & \cdot \end{bmatrix} \begin{bmatrix} \eta\sigma_x\sigma_\mu \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \eta\sigma_x\sigma_\mu \\ \cdot \end{bmatrix}$$

$$= \beta + \frac{\eta\sigma_\mu}{(1+\gamma)\sigma_x} \begin{bmatrix} 1 \\ 1-\gamma \\ \cdot \\ \cdot \\ 1-\gamma \\ 1 \end{bmatrix}$$

where η is the correlation of the fixed effect with R&D and σ_μ^2 is the variance of the effect. The implication is that we would see a large

positive bias in the first and last coefficient and a smaller one in the middle coefficients. If γ is close to one, as it appears to be in our data, the bias for coefficients $\beta_2, \dots, \beta_{k-1}$ would be negligible.

To explore this idea further, we use a longer sample of firms which have R&D data available back to 1970. This leaves us with 346 firms, slightly more than half of our original sample, and somewhat more heavily weighted towards larger firms. For this sample, we obtain estimates for a model with five lags on R&D (shown in Table 5). For comparison, estimates of the original model (with three lags) on this new sample are also shown. It appears from these results that some of the effect we observed in the last lag was indeed due to truncation (note how the coefficient on R_{-3} in column 1 is spread between R_{-3} , R_{-4} , and R_{-5} in column 2); if we push the idea further by estimating with seven lags on the last three years of data (1977 through 1979), the loading on the last lag seems to have disappeared. However, we have also pushed the data beyond the point where it will yield meaningful results, since significant instability in the R&D coefficients for adjacent years is now evident. The tentative conclusion is that there is unlikely to be a substantial effect of R&D more than seven years old on the patent activity of firms in the manufacturing sector.

To examine the other alternative, a fixed effects explanation of the u-shaped lag distribution, we hypothesize a differing propensity to patent for each firm which is (possibly) correlated with its R&D activity. The reasoning in this section suggests that estimates conditional on the permanent patenting propensity of the firm should reduce both the first and last lag coefficient if we have correlated effects and only the last one if the problem is lag truncation (and the lag truncation is relatively constant from year to year). This leads us to look at models which are conditional on the permanent R&D behavior of firms in the next section.

IV. Conditional Estimates

We take two different approaches to obtaining conditional estimates for our model: the first includes all observed values of R&D (for a firm) in each equation with the coefficients constrained to be equal across the different years. This is an attempt to control for fixed effects which may be correlated with our R&D variables, since we cannot simply estimate the effects due to the shortness of our panel and the nonlinearity of the model. The second approach imposes a specific distribution on the error term, namely the negative binomial, allowing us to derive an estimator which is conditional on the total number of patents applied for by the firm in all the observed years.

This second approach was described in our earlier paper (Hausman, Hall, and Griliches 1984); by conditioning on the total number of patents applied for by the firm, it essentially allows for a different intercept for each firm. Due to the multiplicative nature of the error in this model, this translates into a different variance for each firm, so that the conditional model estimates an overall variance parameter, but not the individual intercepts or variances. The log likelihood for this model is

$$(6) \quad \log L = \sum_i \sum_t \log \Gamma(\lambda_{it} + y_{it}) - \log \Gamma(\lambda_{it}) - \log \Gamma(y_{it} + 1) \\ + \log \Gamma\left(\sum_t \lambda_{it}\right) + \log \Gamma\left(\sum_t y_{it} + 1\right) - \log \Gamma\left(\sum_t \lambda_{it} + \sum_t y_{it}\right)$$

Table 6 gives the results of estimates obtained in both ways for both of our samples of firms. The first two columns are estimates of the conditional negative binomial model, while the last two are estimated using weighted nonlinear least squares on equation (1), where the firm effect s_i

includes all the R&D variables in all years, but with coefficients δ constrained to be the same across the years. These two methods of estimation are both compromises of a different sort: the negative binomial version allows for an arbitrary firm effect while making a specific distributional assumption while the GMT version controls for a firm effect correlated with R&D in a particular way (linear in the exponential function) but does not impose a distribution on the error term. It is therefore reassuring that there are not huge differences between them.

The basic result is that none of the coefficients are significant except those on current R&D, although the total effect of the lagged R&D does seem to add about .05 to the coefficient on the sum. It makes very little difference whether we look at the 642 firm sample or the sample of 346 firms which has a longer R&D history. From the fact that the coefficient on contemporaneous R&D hardly changes from the unconditional estimates, while that on the last R&D goes to zero we conclude that most of what we have removed by conditioning is the R&D prior to our longest lag. This confirms the result of Table 5 where we saw a considerable smearing of the lag coefficients when we used a longer lag in the unconditional estimates. However, the coefficients are fairly unstable and the standard errors are large, so the most we can say is that there appears to be a fairly strong contemporaneous effect, even when firm effects are controlled for. Evidence for a contribution of lagged R&D to current patenting activity is of the order of about 0.05 in the conditional estimates and possibly larger in the unconditional.

Using an idea in Pakes and Griliches (1984), we can try to estimate more lags in this equation by assuming that R&D follows a low order AR process, in this case AR1. Since this implies a correlation only between the last included R&D and the presample R&D, the estimates of all coefficients

except the last will be unbiased by the omission of earlier R&D. Accordingly, we leave the last coefficient free in each year of the equation, which allows us to estimate six lags in the 1972-1979 sample and eight lags in the 1970-1979 sample. The precision of the estimates declines with the length of the lag since we have fewer and fewer observations for the longer lags (lag six in the 1972-1979 sample is estimated only from the 1979 equation, for example). However, this constraint allows us to use all but one of the years of data on patents for each sample of firms, so that we have seven years in the 642 firms sample and nine in the 346 firms sample.

We show these results in Table 7; they are essentially the same as the conditional estimates in Table 6. We also estimated this version of the model including firm effects correlated with R&D; these turned out to be insignificant ($\chi^2(7) = 3.9$ for the first column and $\chi^2(9)=11.2$ for the second), although the model in this form is highly collinear so that it is difficult to draw firm conclusions.

The basic message of the results in this section is that permanent differences across firms in the propensity to patent do not appear to bias our estimates of the distributed lag relationship between patenting and R&D, except insofar as they are related to the presample history of R&D. The results of the previous section suggested that this bias, if it existed, would appear of equal magnitude in the first and last lag coefficients of R&D. This does not seem to be the case; what bias there is seems only to affect the last lag, and is eliminated by modelling the presample R&D process itself. This does not imply that there are no differences in the propensity to patent across firms: the size variable and the scientific sector dummy are still significant in the final version

of the model. Nor does it imply that all the differences are uncorrelated with R&D, but only that the correlation which is observed can be successfully explained by controlling for the part of the R&D history which we do not observe.

V. Conclusion

What do we conclude from this lengthy exploration of a basically simple model? First, there does seem to be a rather strong contemporaneous relationship between R&D expenditures and patenting, which does not disappear when we control for the size of the firm, its permanent patenting policy, or even the effects of its R&D history. The remaining elasticity appears to be about .3 with a fairly large standard error. Second, the contribution of the observed R&D history to the current year's patent applications is quite small, on the order of .05. Third, the contribution of the unobserved or presample R&D appears to be large, about .25, and is a possible explanation of the existence of the observed differences across the firms in the propensity to patent.

One of the most interesting results in this paper has nothing to say about patenting, although it provides one reason why we have difficulty measuring the relationship within firm over time: the characterization of the pattern of R&D investment within a firm as essentially a random walk with a relatively low error variance. In other words, R&D budgets over this short horizon (8 years) are roughly constant or growing slightly (in constant dollars) and therefore it is difficult to imagine that they are very sensitive to patenting success or vice versa.

Finally, it is difficult to give a clear cut answer to the question this paper was originally designed to answer: is there a significant longrun effect of successful R&D investment for which patents can serve as an indicator? The evidence here indicates the longrun level of R&D can be quite important, but the result is predicated on inference about the unobserved part of the R&D process. There is very little direct evidence of anything but simultaneity in the year-to-year movements of patents and

R&D. This finding suggests another way of looking at the process: in large industrial firms the fraction of R&D expenditures devoted to development rather than basic or applied research tends to be well over fifty per cent (NSF 1982). It seems reasonable to suppose that successful research leads both to a patent application and to a commitment of funds for development. A detailed investigation of this timing is beyond the scope of annual data, but the strong evidence of simultaneity in patents and R&D in our data conforms very well to this picture.

We should not close this paper on the usual note of the failure of the data to live up to our econometric expertise. Even though we have not been able to elucidate the R&D to patents lag structure better, our overall findings are quite interesting, showing a persistent significant contemporaneous relationship of R&D and patenting and rather wide and semi-permanent differences across firms in their patenting and R&D policies. The later finding provides the challenge for further research in a different style: trying to understand how and why firms differ in their responses to the technological environment they find themselves in.

Table 1

Selection of the Sample of Firms

Sales	Number in 76 Cross Section		Number in Sample	Coverage	
	All	R&D>0		All	R&D>0
less than \$1M	73	33	1	.014	.03
\$1M-10M	548	293	17	.031	.06
\$10M-100M	1102	579	224	.20	.39
\$100M-1B	669	415	259	.39	.62
\$1B-10B	204	167	131	.64	.78
more than \$10B	12	11	10	.83	.91
Total	2608	1498	642	.25	.43

1976 R & D Expenditures
in 1976 dollars

	76 Cross section	Sample	Coverage
less than \$1M	3.0	0.9	.30
\$1M-10M	65.3	4.7	.07
\$10M-\$100M	525.2	243.1	.46
\$100M-1B	2354.1	1790.7	.76
\$1B-\$10B	7830.6	7224.1	.92
more than \$10B	4593.2	4529.2	.99
Total	15,371.3	13,793.0	.90

Table 2
Key Variables in 1976

Variable	642 Firms					346 Firms
	Min	1st Q	Median	3rd Q	Max	Median
Sales (\$M)	.6	57	182	760	49,000	263
R&D (\$M)	.02	.73	2.3	11.0	1,256	3.8
Patents	0	1	3	18	831	5
Fraction with zero patents			.21			.17
Fraction in scientific sector			.37			.42

Notes to Table 2

All dollars are millions of 1976 dollars.

The scientific sector is defined as firms in the drug, computer, scientific instrument, chemical, and electric component industries.

Table 3
Time Series Analysis of Log R&D

642 Firms

Lag	Autocorrelations	Partial Autocorrelations
0	1.0	
1	.986(.056)	.998(.002)
2	.976(.056)	.048(.034)
3	.968(.056)	-.005(.034)
4	.963(.058)	.016(.034)
5	.961(.058)	-.035(.031)
6	.958(.058)	-.004(.032)
7	.954(.058)	.038(.123)

Equation	(1)	(2)	(3)	(4)	(5)
Log R ₋₁	.999(.002)	.934(.039)	.934(.039)	.920(.040)	.922(.040)
Log R ₋₂		.068(.039)	.073(.054)	.058(.040)	.059(.040)
Log R ₋₃			-.005(.034)		
Log P ₀				.028(.009)	
Log P ₋₁				.004(.011)	.017(.009)
Log P ₋₂				[-.007(.008) .004(.009)] →	

Table 4
 Estimates of the Patent Equation
 642 Firms for 1975-1979

Variable	Equation			
	(1) Nonlinear least squares	(2) Poisson	(3) Negative binomial	(4) GMT
log R ₀	.12(.30)	.28(.03)	.21(.07)	.30(.10)
log R ₋₁	.07(.21)	.03(.04)	.07(.10)	.04(.08)
log R ₋₂	-.08(.15)	-.001(.036)	.08(.10)	.06(.08)
log R ₋₃	.28(.24)	.28(.03)	.16(.07)	.25(.11)
sum log R	.39(.09)	.58	.52	.66(.05)
log book plant in 1972	.23(.07)	.21(.004)	.14(.013)	.19(.04)
Dummy (sci. sector)	.36(.23)	.30(.01)	.28(.03)	.21(.11)
δ			.051(.001)	
Log likelihood		280,034.	297,016.	

Notes:

1. All equations have a separate intercept for each year.
2. Standard errors for NLS and GMT are "robust" estimates computed by generalized Eicker-White-Chamberlain formula.

Table 5
 Estimates of the Patent Equation
 346 Firms

Time period	Equation		
	(1) 1975-1979	(2) 1975-1979	(3) 1977-1979
Log R ₀	.16(.16)	.19(.16)	.34(.23)
Log R ₋₁	-.02(.10)	-.07(.10)	-.30(.18)
Log R ₋₂	.07(.10)	.07(.10)	.09(.19)
Log R ₋₃	.36(.18)	.06(.09)	.02(.13)
Log R ₋₄		.16(.08)	.24(.17)
Log R ₋₅		.17(.12)	-.01(.16)
Log R ₋₆			.10(.12)
Log R ₋₇			.10(.12)
Sum log R	.57(.07)	.59(.07)	.57(.07)
Log book plant in 1972	.22(.05)	.20(.06)	.22(.06)
Dummy (sci. sector)	.30(.13)	.30(.13)	.30(.13)

The estimation method is GMT, with separate intercepts for each year, and robust standard errors.

Table 6

Estimates with Firm Effects

Number of Firms	Conditional Negative Binomial		GMT with Correlated Effects	
	642	346	642	346
Log R ₀	.29(.04)	.32(.07)	.23(.07)	.30(.10)
Log R ₋₁	-.01(.05)	-.08(.09)	-.02(.07)	-.10(.08)
Log R ₋₂	.08(.06)	.06(.09)	.04(.06)	.06(.06)
Log R ₋₃	.02(.04)	-.01(.01)	.03(.06)	-.0005(.06)
Log R ₋₄		.04(.07)		.06(.07)
Log R ₋₅		.01(.05)		.04(.07)
Sum log R	.38	.33	.29(.08)	.36(.12)
Log likelihood	-131,539.	-96,362.		

All equations contain time dummies.

Table 7

GMT Estimates Assuming AR1 for R&D

Number of Firms Years	642 1972-1979	346 1970-1979
Log R ₀	.33(.09)	.26(.09)
Log R ₋₁	.03(.06)	.04(.04)
Log R ₋₂	.05(.06)	.01(.04)
Log R ₋₃	-.03(.06)	-.03(.05)
Log R ₋₄	.11(.08)	.09(.06)
Log R ₋₅	-.08(.10)	-.04(.06)
Log R ₋₆	.001(.19)	.05(.08)
Log R ₋₇		.17(.12)
Log R ₋₈		-.27(.14)
Sum log R	.41(.24)	.29(.19)
Log book plant in 1972	.18(.04)	.20(.05)
Dummy (sci. sector)	.20(.10)	.27(.12)

Notes

* This is a revised version of NBER Working Paper #1227, prepared for the Conference on Quantitative Studies of R&D in Industry, Paris, September 9-10, 1983, held under the auspices of the Ecole Nationale de la Statistique et de l'Administration Economique and the National Bureau of Economic Research. We are grateful for helpful comments by the participants in that conference, especially by our discussants, Christian Gourieroux and Paul Stoneman, and by Mark Schankerman. This research was supported by the National Science Foundation (PRA 81-08635) and the National Bureau of Economic Research Program on Productivity and Technical Change. Elizabeth S. Laderman provided extremely able research assistance.

1. Ideally we would like to perform this test also in the other direction using patents on lagged patents and R&D, but there are difficulties in performing a comparable test due to the discrete nature of the dependent variable already alluded to.

2. Another way to understand what the different estimators are doing is to examine the first order conditions (again suppressing the t subscript and writing e_i for $y_i - \exp(X_i\beta)$):

$$\begin{array}{ll}
 \text{NLLS:} & \sum_i \exp(X_i\beta) \quad e_i X_i = 0 \\
 \text{Poisson:} & \sum_i 1 \quad e_i X_i = 0 \\
 \text{Negative binomial:} & \sum_i (1 + \eta^2 \exp(X_i\beta))^{-1} e_i X_i = 0 \\
 \text{GMT:} & \sum_i \frac{\exp(X_i\beta)}{\hat{y}_i + \eta^2 \hat{y}_i^2} \quad e_i X_i = 0
 \end{array}$$

Note that the first order condition which we show for the negative binomial model is conditional on the choice of η^2 . Since η^2 is being estimated simultaneously this is not the full set of first order conditions for the problem; we merely include it for illustrative purposes. Joint estimation of η^2 is precisely what makes this estimator inconsistent when the distribution is not negative binomial, although the other three estimators remain consistent in this case since they are all version of weighted least squares.

Displaying the first order conditions in this way reveals that the estimators only differ in their choice of weights, although

NLLS and GMT are minimum distance estimators, and Poisson and negative binomial are maximum likelihood estimators. They can be ranked by the weight which they give to firms with larger X's,

References

- Bhargava, A. and D. Sargan. 1983. "Estimating Dynamic Random Effects Models from Panel Data Covering Short Time Periods," Econometrica 51: 1635-1660.
- Bound, J., C. Cummins, Z. Griliches, B. H. Hall, and A. Jaffe. 1982. "Who Does R&D and Who Patents?", National Bureau of Economic Research Working Paper #908, in R & D, Patents, and Productivity, 1984, Z. Griliches (ed.): 21-54. Chicago: University of Chicago Press, 1984.
- Chamberlain, Gary. 1982. "Multivariate Regression Models for Panel Data," Journal of Econometrics (18)1: 5-46.
- _____. 1978. "Omitted Variable Bias in Panel Data: Estimating the Returns to Schooling," Annales de l'Insee, #30-31: 49-82.
- Cummins, C., B. H. Hall, and E. Laderman. 1983. "The R&D Master File: Documentation." Unpublished.
- Eicker, F. 1967, "Limit Theorems for Regressions with Unequal and Dependent Errors, in Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Vol. 1, Berkeley: University of California.
- Gourieroux, C., A. Monfort, and A. Trognon. 1984. "Pseudo Maximum Likelihood Methods: Application to Poisson Models," Econometrica 52: 701-720.
- Hausman, J. A., B. H. Hall, and Z. Griliches. 1984. "Econometric Models for Count Data with an Application to the Patents R&D Relationship," Econometrica 52: 909-938.
- Leamer, Edward E. 1978. Specification Searches: Ad Hoc Inference with Nonexperimental Data, New York: John Wiley and Sons, Inc.
- MaCurdy, T. 1981. "Time Series Models Applied to Panel Data," Mimeo, Stanford University.
- Mansfield, E. 1968. Industrial Research and Technological Innovation, New York: W. W. Norton.
- _____. 1972. Research and Innovation in the Modern Corporation, New York: W. W. Norton.
- Pakes, Ariel. 1982. "Patents, R&D, and the Stock Market Rate of Return," National Bureau of Economic Research Working Paper #786.
- _____. 1984. "The Option Value of Patents", mimeo, Hebrew University, Jerusalem.

Pakes, A. and Z. Griliches. 1980. "Patents and R&D at the Firm Level: A First Look," National Bureau of Economic Research Working Paper #561, in R & D, Patents, and Productivity, 1984, Z. Griliches (ed.): 55-72. Chicago: University of Chicago Press.

_____. 1984. "Estimating Distributed Lags in Short Panels with an Application to the Specification of Depreciation Patterns and Capital Stock Constructs," Review of Economic Studies LI(2): 243-262.

Schankerman, M. and A. Pakes. 1984. "The Rate of Obsolescence and the Distribution of Patent Values: Some Evidence from European Patent Renewals," Revue Economique. Paris: ENSAE.

Stoneman, Paul. 1983. "Patents and R&D: Searching for a Lag Structure - Comment," Proceedings of the Conference on Quantitative Studies of Research and Development in Industry. Paris: ENSAE.

Taylor, C. T. and A. Silberston. 1973. The Economic Impact of the Patent System. Cambridge: The University Press.

White, Halbert. 1982. "Instrumental Variables Regression with Independent Observations," Econometrica 50: 483-500.

_____. 1980. "Nonlinear Regression on Cross-Section Data," Econometrica 48: 721-746.