

EPTD DISCUSSION PAPER NO. 35

**DYNAMICS IN THE CREATION AND DEPRECIATION OF
KNOWLEDGE, AND THE RETURNS TO RESEARCH**

Julian Alston, Barbara Craig, and Philip Pardey

Environment and Production Technology Division

International Food Policy Research Institute

2033 K Street, N.W.

Washington, D.C. 20006 U.S.A.

August 1998

EPTD Discussion Papers contain preliminary material and research results, and are circulated prior to a full peer review in order to stimulate discussion and critical comment. It is expected that most Discussion Papers will eventually be published in some other form, and that their content may also be revised.

ABSTRACT

Econometric studies of the effects of research on productivity have typically imposed arbitrary restrictions on the length and shape of the R&D lag profile. These restrictions are likely to have biased up both the measured effects of R&D on productivity and the estimated rates of return to research. This paper argues that the useful stock of public knowledge depreciates, if at all, only gradually, and we use this notion to develop a new model, which we test using data on aggregate U.S. agriculture. We reject the conventional specification in favor of a more flexible, dynamic, alternative model, in which the impact of R&D on productivity lasts much longer than in previous studies. Consequently, the real, marginal rate of return to public agricultural R&D in the United States is estimated to be less than 10 percent per annum, much smaller than the typical rates of return reported in scores of previous studies, based on conceptually flawed and inappropriately restrictive dynamic specifications. We show that conventional approaches using the same data would have resulted in a much greater (biased) estimate of the rate of return to research.

CONTENTS

1. Introduction	1
2. Conceptual Framework	4
Research Investment and Additions to the Stock of Knowledge	5
Changes in Production from Publicly Produced Knowledge	7
Aggregate Productivity Models	11
3. Previous Empirical Work on Agricultural R&D	14
4. Econometric Models	20
The Model and Hypotheses	21
Model Selection	22
Rates of Return	32
5. Conclusion	39
References	42
Appendix A: Approximating the Internal Rate of Return	47

DYNAMICS IN THE CREATION AND DEPRECIATION OF KNOWLEDGE, AND THE RETURNS TO RESEARCH

Julian M. Alston, Barbara J. Craig, and Philip G. Pardey*

1. INTRODUCTION

For as long as most of us can remember, economists have argued that public research has yielded such high rates of return that there is clear evidence of persistent underinvestment. This underinvestment hypothesis has taken on the character of a stylized fact, to be explained.¹ In this paper we argue, provocatively, that the previous

* Julian M. Alston is a professor at the University of California, Davis, Barbara J. Craig is an associate professor at Oberlin College, and Philip G. Pardey is a research fellow at International Food Policy Research Institute. Senior authorship is not assigned. The authors wish to thank Albert Acquaye for excellent research assistance, and T. J. Wyatt for help with constructing table 1. Helpful comments on an earlier version were provided by participants at the international conference on Global Agricultural Science Policy for the Twenty-First Century held in Melbourne, Australia, August 1996. Particular thanks are due to Ron Bewley, Zvi Griliches, and Robert Myers whose suggestions led to substantial revisions in the econometric models, and to Daniel Sumner, Jenni James, Jock Anderson, Bob Evenson, and Wally Huffman for detailed comments on drafts. Partial funding support for this work was provided by the University of California's Pacific Rim Project.

¹ Boyce and Evenson (1975, pp.116-7) wrote " ... the extraordinarily high rates of return that have been measured in virtually all of the studies of agricultural research productivity must be taken to show that investment levels have been too low to represent efficient allocation of scarce resources." More recently, Alston and Pardey (1996, 199-227) reviewed the literature and speculated about likely biases in the reported rates of return but nonetheless concluded that "... it seems likely that the rate of return to agricultural R&D has been relatively high, and that there has been some underinvestment (p.219)." Fuglie et al. (1996, 28) also concluded that the returns are high stating that "Most studies that have estimated the aggregate social rate of return to research consistently found rates of return between 40 and 60 percent." In fact the range of results in the literature is much broader than that: Alston et al. (1997) reviewed a total of 141 studies and found that 80 percent of the estimates fall between 18 percent and 108 percent.

evidence has been severely biased as a result of a conceptual error, leading to misspecification of the dynamics of knowledge creation and productivity. Further, we show that when we measure the effects of research on production more appropriately, the estimated rate of return to research is closer to a normal market rate of return than to the very high rates that predominate in the literature.

Modeling and measuring the productivity consequences of R&D is a tricky business. Our primary issues concern the dynamic structure linking research spending, knowledge stocks, and productivity. Research takes a long time to affect production, and then it affects production for a long time. Since the typical study has available only a single, short time-series of annual data, however, the usual practice has been to impose arbitrary restrictions on the length and shape of the research lag profile. These assumptions about the form of the lag structure have remained largely untested, even though the use of misspecified lag structures is likely to lead to biased estimates of both the effects of R&D on productivity and the rate of return to research.

The lag structures used to represent research effects imply specific, and in many cases extreme, assumptions about the obsolescence of knowledge. Some studies have treated knowledge as if it never depreciates, having the same productive effect forever, once it has reached its maximum. This view is clearly inaccurate for technologies that we have seen come and go, such as eight-track cartridges, or are simply no longer as popular or as important as they once were, such as the scythe or the buggy whip. Other studies have treated *all* knowledge as eventually becoming obsolete, so that its effect on production falls to zero—within, say, 30 years from beginning the research that created it.

This is clearly not an accurate depiction of the contribution of knowledge embodied in the wheel, the electric light, or the internal combustion engine. Hence, in a model in which total production depends on aggregate research investments, aggregating across the complete spectrum of different types of research and knowledge, a short, finite research lag is questionable.

Like Griliches (1996), we observe that knowledge does not depreciate physically, but that some knowledge becomes obsolete. Research leads to new knowledge that adds to the existing stock of knowledge, and while knowledge itself does not deteriorate, its utilization, and hence the relevant stock of knowledge, may change as circumstances change. This view implies a significantly different specification of the lag relationship between investments in R&D and production, compared with those previous studies that have used finite lag structures.

Studies of industrial R&D often have fewer than 30 years of data, while our application to U.S. agriculture has almost 100 years of R&D data. This allows for an exceptionally long-run empirical perspective on the role of R&D in productivity growth. Our conceptual model leads to new specifications in which we nest the typical empirical model as a special case. Using the U.S. agricultural data, we reject the typical restrictions and find in favor of a model with a more complicated dynamic structure and an infinite research lag. With that model, we illustrate a substantial upward bias in the measured rates of return to R&D from mis-specification of the dynamic structure.

Our findings are qualitatively as well as quantitatively different from those in previous studies. Agricultural R&D affects productivity for much longer than prior work

suggests. Our much lower rates of return to publicly performed agricultural R&D in total do not support the position argued in almost every previous empirical study of returns to agricultural research—that there has been a gross underinvestment overall. Rather, we find respectable rates of return that support the idea that the investment has been socially profitable, but not enormously so. We argue, therefore, that better evidence must be found to support the persistent claims for expanding rural research support from the public purse. Though motivated by our interest in important questions about agricultural productivity and agricultural development, we believe our contribution is more broadly applicable. If confirmed by further work, our results are of more general importance in understanding and measuring the role of knowledge creation and utilization in productivity growth.

2. CONCEPTUAL FRAMEWORK

Estimation of the returns to research requires an understanding of the process by which research leads to innovations or additions to the stock of knowledge as well as an understanding of how the stock of knowledge itself affects production. Most studies fail to distinguish between the stock of knowledge and the flow of innovations, and this is likely to give rise to serious empirical mis-specifications. The predominant empirical practice in the literature has been to treat current and a number of years of lagged research spending as a direct input to production. This is equivalent to assuming that the currently useful knowledge stock depends *only* upon innovations arising from research undertaken

over a finite number of recent years. When the stock of knowledge affecting production contains the innovations having arisen from *all* past research, estimating the returns to research would appear to require an infinite set of parameters and infinite data. Lacking direct observations on the stock of knowledge itself or such long time series, we must find a feasible way to use available data that does the least damage to the resulting estimates. The imposition of some structure on lagged effects is probably inescapable. However, lag structures that inappropriately truncate or shape the possible pattern of returns to research will bias estimates in predictable ways. In what follows, we consider some fairly general theoretical specifications of the link between research and the stock of knowledge, and illustrate the consequences of employing these specifications to estimate returns to research.

RESEARCH INVESTMENT AND ADDITIONS TO THE STOCK OF KNOWLEDGE

Research expenditures can be thought of as investments in knowledge capital, but unlike physical capital, there is little hope of ever directly measuring the knowledge created by research. The investment may generate a process or idea that can be represented tangibly, such as a blueprint or book, or a marketable form such as a new pesticide or seed variety, but we cannot build an index of the aggregate stock of knowledge. We can usefully speculate, however, on the likely features of the relationship between the flow of research investments and the stock of knowledge. In particular, it is reasonable to expect to find long and variable lags between research investments and their eventual contributions to the stock of useful knowledge and, thereby, production. The

effects of a particular past research investment could continue indefinitely, but it is also reasonable to presume that the effects of some particular past research on production will eventually diminish in the future. The process of decline over time in the effects on production of some past line of research can be thought of in two alternative ways: one is in terms of physical depreciation of the knowledge stock; another is in terms of changes in utilization of a stock that never depreciates.² The physical depreciation perspective has dominated empirical research on agricultural R&D and is pervasive in the literature on R&D and productivity more generally.

Does the stock of socially useful knowledge capital depreciate with the march of time? One conceptual problem here is that—at least at first glance—we do not have a simple analogue in the depreciation of physical capital that we might extend to knowledge capital.³ Perhaps ideas are “lost” as historical experience exposes their shortcomings or they lose their effectiveness as pests adapt to outwit them. But it is harder to picture this than to picture the inevitable wear and tear on machinery or the gradual decay of buildings. It seems transparent that the stock of *old* knowledge decays as new ideas *replace* old ones. This means it is important to consider how obsolescence could be modeled so that we distinguish between gross and net increases in the stock of knowledge or, relatedly, the total stock versus the useful stock. But apart from the replacement of

² A third option is a combination of the depreciation and obsolescence perspectives, e.g., Evenson (1968).

³ Griliches (1996) distinguishes between depreciation in social value and in private value of knowledge stocks, which he illustrates with the example of loss of patent monopolies.

old ideas with new, contradictory or, in some other sense, superior ideas, new knowledge does not reduce the old knowledge stock and there is no other physical process that diminishes the stock of knowledge.

The depreciation view is much less applicable in situations where the reason for a change in application of a particular piece of knowledge or technology is driven by changes in circumstances that make alternative technologies more economic. It is probably more useful to think of changing resistance of pests as a change in the pests themselves rather than a depreciation in a particular pest-control technology that has not, in fact, changed at all. The “knowledge” is the same; the effectiveness and hence the utilization of the knowledge is what has changed. Changing utilization of knowledge, adoption and disadoption—driven by relative prices, the emergence of new technologies, or other changes in the physical and economic environment of production—seems likely to be the predominant source of changes in the productive effectiveness of knowledge.

CHANGES IN PRODUCTION FROM PUBLICLY PRODUCED KNOWLEDGE

Suppose output, Q_t , is taken to be a function of inputs, \mathbf{X}_t , the stock of useful knowledge, K_t , a vector of nonmarket inputs, \mathbf{Z}_t , such as weather, and purely random elements, u_t .

$$Q_t = f(\mathbf{X}_t, K_t, \mathbf{Z}_t, u_t) \quad (1)$$

Assuming the production function is separable in these variables and characterized by constant returns to scale, it may be expressed, alternatively, as a productivity function:

$$MFP_t = \frac{Q_t}{X_t} = g(K_t, Z_t, u_t) \quad (2)$$

where MFP_t is a multifactor productivity index and X_t is an input quantity index.⁴ Such a model implicitly treats the knowledge stock as causing, but not being caused by, changes in output.⁵

One approach to specifying the effect of research on knowledge is to treat all knowledge as fungible and nondepreciable, so that the aggregate stock of knowledge evolves over time according to

$$K_t = K_{t-1} + I_t, \quad (3)$$

where I_t represents current innovations or additions to the stock of knowledge. These innovations are made possible by current and past research investments, R , and chance, ϵ_t . Here, the research variable is general, without any distinction between public and private, or basic and applied, research and extension.

$$I_t = i(R_t, R_{t-1}, \dots, R_{t-L}, \epsilon_t) \quad (4)$$

⁴ Some new knowledge becomes embodied in inputs and outputs, creating conceptual and measurement problems to be addressed in deciding how to adjust measured input and output quantities for changes in their qualities and how to interpret productivity. Below we use quality-adjusted input data which means it may be a form of double counting to include the additional knowledge that caused those quality changes in the production function. See Alston, Norton, and Pardey (1995, 153-65) for some discussion of these issues.

⁵ Simultaneity between research investments and output is not likely to be an issue in modeling the relationship between public agricultural R&D and aggregate output, so we do not deal with it here. It could be a serious issue in modeling public research variation across industries, say, or in modeling impacts of private research.

Research investments affect increments to knowledge for up to L years. Implicit in the structure in (3) is the idea that increments to knowledge in the past, that were caused by research in the very distant past, are just as effective today as the most recent increments to knowledge. Knowledge grows unidirectionally, and utilization does not vary according to the vintage of the knowledge.⁶

One way to add empirical realism to this specification is to treat the problem *as if* knowledge depreciates. The imposition of a depreciation rule on the stock of knowledge would make effective utilization of more-recently developed knowledge greater. Just what sort of depreciation rule to use depends, in turn, on what sorts of innovations or ideas one is modeling (and, perhaps, who owns the results).⁷ In an aggregate context, a range of depreciation patterns may coexist. Hence, in the spirit of Jorgenson (1973), and as suggested by Griliches (1996), a proportional declining balance, or geometric depreciation rule would seem to be a simple yet natural choice to represent changes in an aggregate stock of knowledge. Using δ to denote the depreciation rate, the aggregate stock of knowledge would evolve over time according to⁸

⁶ Fulginiti and Perrin (1993) use this model.

⁷ Noting Boulding's (1966) point that knowledge does not physically deteriorate, Pakes and Shankerman (1987, 74) argue that its *value* to the firm who owns a patent does, owing to displacement by new innovations and rising appropriability problems. For further discussion on the creative destruction of knowledge stocks through private R&D, see Caballero and Jaffe (1993).

⁸ Another idea is to make the rate of depreciation of the knowledge stock a function of the rate of invention, so that $K_t = (1 - \delta(I_t))K_{t-1} + I_t$. It is sometimes suggested that the rate of decay of knowledge is higher in sectors that are more inventive. This effectively makes the rate of decay vary over time and, while it may be an appealing conceptual view, it would be difficult to make much empirical use of it, especially with

$$K_t = (1 - \delta)K_{t-1} + I_t. \quad (5)$$

The industrial R&D literature reports a range of estimates of geometric depreciation rates for R&D capital stocks. Adams (1990) estimated an annual depreciation rate for basic research of 9 to 13 percent, while Nadiri and Prucha (1993) estimated a rate of 12 percent for industrial R&D. Based on this, and other, evidence, the Bureau of Economic Analysis (BEA) (1994) used a straight-line life span that corresponds to a geometric depreciation rate of 11 percent in constructing estimates of R&D net capital stocks. This rate implies that only 10 percent of today's knowledge stock will remain in use in 20 years' time. The Bureau of Labor Statistics (BLS) (1989) used a slightly smaller rate of 10 percent as its central estimate, which also implies a rapid rundown of the stock of useful knowledge. BLS (1989, 38) also considered 0 and 20 percent depreciation and noted that "Choice of a specific rate of depreciation turns out to have important implications for the effect of R&D on productivity growth." In one earlier study, Griliches (1980) considered depreciation rates of 0, 10 and 20 percent; in another (Griliches 1986), 15 percent. More recently, Coe and Helpman (1993) suggested a depreciation rate of 5 percent for research applied to business-sector R&D capital, implying a much longer-lived effective stock. Coe and Helpman are in the distinct minority with their comparatively low depreciation rate

aggregate data on a single sector. The time-invariant rate of decay discussed and measured in this paper (and by others) can be thought of as a long-term, average rate of depreciation of the relevant knowledge stock.

(and implied long-lived effects of research on productivity).⁹

The agricultural R&D literature has not dealt with knowledge depreciation as explicitly as the industrial R&D literature. Most previous studies of returns to agricultural R&D have effectively set $\delta = 1$ so that $K_t = I_t$; hence, they have built in the idea that *all* knowledge created by research is lost. This is true of all studies that model current production, or productivity, as a function of a finite lag of research investments. In these studies, the finite lag means that if research were to cease, the knowledge stock would run down to zero over a relatively short period—and, in many models, the implication would be for zero production as a result. All models are an abstraction from reality, but we argue that this abstraction results in seriously biased estimates of the effects of research on production and the resulting rate of return.

AGGREGATE PRODUCTIVITY MODELS

Combining the elements in equations (1), (2), (4), and (5), current production or productivity may depend on the entire history, an infinite lag, of past research investments. It is plausible that the lag relationship defining *increments* to knowledge is finite, and it might also be reasonable to suppose that its length, L , is less than 30 years. But this lag structure should not be confused with the lag relationship between investments in research and production; this lag is infinite, even if the stock of knowledge can be regarded as

⁹ Hall and Mairesse (1995) recently explored the R&D-productivity relationship using 16 years of R&D data for French manufacturing firms. They formed an R&D capital stock using $K_t = (1 - \delta)K_{t-1} + R_{t-1}$ and depreciation rates, δ , of 15, 25, and (implicitly) 100 percent.

depreciating geometrically. A finite research lag structure in a model of production means that the stock of *all* knowledge depreciates in some fashion such that, if research were to cease permanently, eventually the knowledge stock would be driven to zero.

Even though a finite lag between research and its effects on production is implausible, this does not mean the econometric problem is hopelessly intractable. Suppose the knowledge stock depreciates according to the pattern defined in (5). In that case, if the effect of the knowledge stock on production (or productivity) were additive, then taking quasi-differences would yield a model from which the stock variable can be eliminated. Hence, finite lags of research investments can be used to represent the infinite effect of research on productivity. Assuming a linear form, and a single weather-related variable, Z_t ,

$$MFP_t = \alpha + \beta K_t + \gamma Z_t + u_t \quad (6)$$

Using equation (5), we can replace the current knowledge stock in equation (6) and quasi-difference the equation to get¹⁰

$$\begin{aligned} MFP_t - \alpha\delta - \beta[K_t - (1 - \delta)K_{t-1}] - \gamma[Z_t - (1 - \delta)Z_{t-1}] \\ = (1 - \delta)MFP_{t-1} + u_t - (1 - \delta)u_{t-1}. \end{aligned} \quad (7)$$

Taking innovations in equation (4) to be a linear function of a finite lag of past logarithms

¹⁰ In this case, quasi differencing involves lagging equation (6) one period, multiplying by $(1 - \delta)$, subtracting the result from equation (6), and reducing terms to get equation (7).

of research investments, the unobserved capital stock terms can be replaced to yield¹¹

$$MFP_t = \alpha \delta \left[\beta \sum_{s=0}^{L_R} b_s \ln R_{t-s} + \gamma [Z_t - (1 - \delta)Z_{t-1}] \right] + (1 - \delta)MFP_{t-1} + v_t. \quad (8)$$

This model nests the three primary alternatives: the stock of useful knowledge never depreciates, $\delta = 0$; the stock of useful knowledge vanishes in finite time, $\delta = 1$; and the intermediate case in which useful knowledge decays only gradually, $0 < \delta < 1$.

In the model most frequently used to estimate effects of research on agricultural productivity and production, only lagged research investments and sometimes a weather-related variable appear on the right-hand side. This amounts to an implicit assumption that $\delta = 1$, so that the stock of useful knowledge decays in finite time. Most models also impose additional restrictions on the lag weights (the b_s coefficients). Both types of restrictions on the model can have significant effects on the empirical assessment of research benefits. In particular, Alston and Pardey (1996, 219-27) have argued that the truncation of the lag length is likely to lead to larger rate-of-return estimates. Their reasoning is that, if the omission of longer lags of R&D investments represents the omission of relevant explanatory variables, and R&D investments are strongly positively correlated over time, the weights on the shorter lags will be biased up and so, accordingly, will be the rate of return—so long as the effect of the higher values for short lag weights

¹¹ In earlier work we tried a linear research stock (see Alston, Craig, and Pardey, 1996). We thank Zvi Griliches for suggesting the use of a logarithmic research stock, which we found to have more attractive time-series properties. The ultimate outcomes were not materially affected by this choice, however.

more than compensates for the imposition of zero restrictions on the longer lag weights, which is likely with high rates of return that heavily discount benefits in the distant future. Both the direction and magnitude of the bias in estimated rates of return to R&D from restricting the length and shape of the lag profile are empirical issues, but we suspect the bias has been upwards.

3. PREVIOUS EMPIRICAL WORK ON AGRICULTURAL R&D

A large number of previous studies have regressed a measure of agricultural production or productivity against variables representing agricultural research and extension, often with a view to estimating the rate of return to research. Many of those studies were reviewed by Alston and Pardey (1996), following Echeverría (1990) and Huffman and Evenson (1993), who extended earlier reviews such as by Evenson, Waggoner, and Ruttan (1979). A comprehensive reporting and evaluation of this and more recent literature is provided by Alston et al. (1997).

Here we are concerned in particular with studies of aggregate agricultural production or productivity using time-series data, and especially with the treatment of the research lag. We note, however, that the same issues arise more generally in models linking R&D and productivity, not just in agriculture, and the excessive restriction of R&D lags may be similarly prevalent in the nonagricultural applications. For instance, drawing on his own previous work and that by Minasian (1969), Griliches (1980, 424) discussed a knowledge stock defined as a distributed lag of a finite number of past

research investments with the summation running over the available range of data. In support, Griliches noted that Evenson (1968) found evidence of lags in agriculture of up to 16 years, and he appears to have suggested even shorter lags are appropriate for private R&D, given the predominance of applied research in industry. While the more recent literature on returns to industrial R&D has shifted towards geometric depreciation structures with, implicitly, infinite lags, the relatively large assumed rates of depreciation mean the lags are effectively still short.

In one common type of study, total factor productivity is regressed against a research stock variable, defined as a weighted sum of real research expenditures in some number of past years. Other variables, such as weather indexes, are often included as well. Sometimes the variables are expressed in logarithms.¹² The specification of the research variable concerns us most here, especially the lag length. The number of lags and the shape of the lag structure are usually chosen arbitrarily; rarely is either the lag length or the lag form tested formally. Common types of lag structures include the de Leeuw or inverted-V (e.g., Evenson 1967), polynomial (e.g., Davis 1980; Leiby and Adams 1991; Thirtle and Bottomley 1988), and trapezoidal (e.g., Huffman and Evenson 1989a, 1992, 1993; Alston, Pardey, and Carter 1994; Evenson 1996). A small number of studies have used free-form lags (e.g., Ravenscraft and Scherer 1982; Pardey and Craig 1989; Chavas and Cox 1992), but most have restricted the lag distribution to be represented by a small number of parameters, because the time span of the data set is usually not much longer

¹² This transformation is almost always used when the model is estimated as a production function, with output quantity replacing productivity as the dependent variable, in which case the explanatory variables include conventional input quantities as well.

than the assumed maximum lag length.

Until quite recently, it was common to restrict the lag length to less than 20 years. In the first studies, available time series were short and lag lengths were very short. For instance, Griliches (1964) used an average of the nominal expenditures on research and extension in two years, $t-1$ and $t-6$, with no notion of a knowledge stock. Other early studies sometimes used more lagged research variables, but only up to 10 or 12 lags (e.g., Evenson 1967; Bredahl and Peterson 1976; Norton, Coffey, and Frye 1984).

The more recent studies have tended to use more, and longer lags. For instance, Thirtle and Bottomley (1988) reported a range of results from using different productivity indices, and having tested for lag length with a maximum lag of 15 years (maintaining a second-degree polynomial with end-point restrictions). Their preferred lag lengths range from 11 to 15 years, with a small “gestation lag”—a number of years between making an investment and beginning to have an effect on productivity. This is an important point since, when they allow for some gestation lag, holding the maximum lag constant at 15 years, a lead time of five years seems statistically preferred. Increasing the lead time from zero to five years reduced the marginal internal rate of return from 79 percent per annum to 15 percent. A related point can be seen in the results of Leiby and Adams (1991), who used second-order polynomial lags of between eight and 16 years. Increasing the lag lengths from eight to 16 years resulted in much smaller estimated rates of return (to about one quarter of the value estimated with eight years). Both these effects are consistent with arguments presented by Alston and Pardey (1996).

None of the above studies used as many as 20 lagged research variables. Pardey

and Craig (1989) used a free-form lag structure to model the relationship between agricultural productivity and public-sector agricultural research, and found “strong evidence that the impact of research expenditures on agricultural output may persist for as long as thirty years” (p. 9) and that “long lags—at least thirty years—may be necessary to capture all of the impact of research on agricultural output.” (p. 18). Schimmelpfennig and Thirtle (1994) replicated Pardey and Craig's study using U.K. data, although their results on lags are less clear. Using a nonparametric approach, Chavas and Cox (1991, 590) confirmed Pardey and Craig's result, finding that “at least 30 years of lags are necessary to capture the effects of public research.” Several subsequent studies have followed this advice. However, none of these studies, including Pardey and Craig (1989) and Chavas and Cox (1991), tested how much longer than 30 years or so the lag should be, nor did they consider the problem in the context of an infinite lag structure. For example, Huffman and Evenson (1989b) used a stock of public applied research, calculated as a trapezoidal lag with a total of 35 lags: after two years of gestation lags, lag weights are assumed to rise linearly for seven years, then to stay constant for six years, and finally to decline linearly for 20 years.¹³ A number of other studies have simply adopted Huffman and Evenson's trapezoidal lag structure (e.g., Alston, Pardey, and Carter

¹³ In subsequent work, Huffman and Evenson (1992, 1993) do not mention the gestation lag, but continue to use the same shape for the rest of the lag for a total lag of 33 years. The identical structure was imposed to create a stock of public pre-technology science research. Huffman and Evenson (1989b, 1992, and 1993) do not describe how this lag structure was chosen, but a draft version of Huffman and Evenson (1993) published in 1989 states “A minimum-mean-square-error statistical criterion was used to evaluate the performance of the different distributions [of the lag weights]...” (p. 11-17). Extension was also included with three lags, assuming weights of 0.5, 0.25, 0.25.

1994).

Here we have argued for representing an infinite lag between research investments and productivity with a finite lag between research investments and *changes in* the stock of knowledge. Using the same types of arguments, Adams (1990) developed a similar model, although his dependent variable was the growth rate of productivity. Some other recent studies, beginning from an examination of the time-series structure of the data, rather than reflection about the structural relationships, have been tending in a similar direction (e.g., Akgüngör et al. 1996; Makki, Tweeten, and Thraen 1996; Myers et al. 1996). They have used time-series methods involving data transformations, such as first differences, and they have found smaller estimated rates of return as a result. An important difference is that our theory indicates a particular transformation of the data, and a specific interpretation of the resulting estimates.

Table 1 summarizes the results from past econometric studies of returns to agricultural research across countries, classified according to the length and form of the research lag, and it can be seen that the results are consistent with our expectations. Most studies have used short lags (and other restrictions on the form of the lag) and shorter lags tend to coincide with larger estimated rates of return.

Table 1 Review of rates of return to agricultural R&D (preliminary, subject to revision)

Form of Lag	Total Lag Length		Number of Estimates*	Internal Rate of Return		
	Average	Range		Average	Lowest	Highest
	<i>years</i>			<i>-- percent per annum --</i>		
Free Form	14	<15	2	111.2	104.8	117.6
	∞	>30	1	14.0	14.0	14.0
All Free Form		14- ∞	3	78.8	14.0	117.6
Polynomial [‡]	11.4	<15	91	73.4	14.3	578.5
	18.0	15-30	17	48.3	15.0	175.8
		unspecified	9	32.0	0.0	55.0
All Polynomial	12.4[§]	0-29	117	66.6	0.0	578.5
Trapezoidal	27.4	15-30	8	107.3	19.6	218.2
	35.0	>30	7	268.9	62.0	386.4
All Trapezoidal	30.9	27-35	15	178.0	19.6	386.4
Other [†]	7.0	<15	25	95.6	33.0	325.0
	28.0	15-30	46	59.4	6.0	130.0
	45.5 [§]	>30	16	65.2	33.0	260.0
		unspecified	20	59.9	25.0	94.0
All Other	21.3[§]	0-∞	107	68.8	6.0	325.0
All Forms [†]	10.1	<15	122	77.5	14.3	578.5
	25.5	15-30	71	62.1	6.0	218.2
	37.3 [§]	>30	24	119.6	14.0	386.4
		unspecified	29	51.3	0.0	94.0
All Studies	16.8[§]	0-∞	246	74.1	0.0	578.5

* Alston et al. (1997) reviewed a total of 50 econometric studies that reported a total of 240 estimates of rates of return to research. In this column, the number in each row refers to the number of estimates that were obtained using the form and length of research lag, as defined for that row, to estimate the parameters. We have tabulated separately studies that reported estimates according to the lag structure used for *computation* of benefits, which is sometimes not the same as that used for *estimation* of parameters.

† These include 11 models with inverted-V lag structures, 1 with exponential lags, and 35 models with various other, imposed lag structures.

‡ All but two of the polynomial lag models used second-order polynomials.

§ The average includes only those with a lag length less than infinity.

4. ECONOMETRIC MODELS

The foregoing discussion suggests a variety of models that may be tried. Here we consider linear models of productivity. We use annual data on U.S. aggregate agricultural productivity for 48 contiguous states over the period 1949-1991, and annual data on total agricultural R&D (and extension) expenditures by the federal and 48 state governments, for the period 1890-1991. According to Craig and Pardey's (1996) estimates, U.S. agricultural productivity grew by 1.76 percent per annum during 1949-1991, with a slightly slower rate of growth of 1.38 percent per annum during the last decade in the sample, 1980-1991, which covers years not included in most previous studies. In nominal terms, public spending on agricultural R&D totaled \$2,503 million in 1991, compared with just \$81 million in 1949, and this does not include spending on extension, which also grew, but less rapidly, from \$66 million in 1949 to \$1,365 million in 1991 (see Alston and Pardey 1996 for details).

In this section we first examine the time-series properties of the variables in the model. Then we test for lag length jointly with model structure (i.e., alternative restrictions on the parameter on the lagged dependent variable) using various forms of lag distribution. We begin with free-form lags; then we impose an Almon polynomial lag structure and evaluate the imposition of endpoint restrictions; and finally we try a trapezoidal lag structure.

THE MODEL AND HYPOTHESES

The model is based on equation (8). We include only one variable to represent weather (and other) effects on productivity, a national pasture and rangeland condition index, measured in September for each year, PRC_t . MFP_t is a Törnqvist-Theil divisia index of multifactor productivity in U.S. agriculture, as defined by Craig and Pardey (1996), and based on their state-level, quality-adjusted, indexes of output and inputs.¹⁴ R_t is the real value of U.S. state and federal government spending on agricultural R&D. This variable is found by adding the total R&D spending (Alston and Pardey 1996, table 2-A3) to total extension spending (Alston and Pardey 1996, table 2-A6) and dividing by a research deflator, which is based upon average salaries of university professors.

To rule out the possibility that the empirical relationships found between MFP and research spending are spurious owing to unit roots in both the underlying series, unit-root tests were employed. Using either Dickey-Fuller tests or the augmented Dickey-Fuller tests, we reject the hypothesis of a unit root in all three series in favor of a stationary trend. We cannot reject a unit root in the research spending variable when it is measured in levels, but the 99-year series of logarithms of research spending is judged to be stationary about a linear trend. For this reason we prefer the use of the logarithms.

¹⁴ Craig and Pardey (1996) measured significant changes in the quality of land, labor, and capital over time, using state-specific hedonic prices. Some of these changes reflect public and private investment in infrastructure, education, and R&D; but most of the changes are really not attributable to public agricultural R&D, and it is therefore appropriate to adjust the productivity data and remove these effects before attempting to measure the effects of public agricultural R&D investments on productivity. These adjustments will also mitigate the consequences of leaving private research and education out of the productivity model.

MODEL SELECTION

The first step in the empirical work emphasizes the identification of lag length (following Greene 1993). Testing for lag length is tricky. The tests that are available are low-powered and tend to lead to overfitting. Both Akaike's Information Criterion (AIC) and Amemiya's Probability Criterion (APC) were applied to test restrictions on the lag length. The model was estimated with and without the restriction that $\delta = 1$, allowing up to 30 free-form lags of R_t and gestation lags of up to 10 years.¹⁵

The results of these estimations, reported in table 2, can be summarized as follows. First, if we impose $\delta = 1$ and hence exclude both the lagged dependent variable and the lagged weather variable, we reject models that restrict the maximum lag length below 30 years—like Pardey and Craig (1989) and Chavas and Cox (1991), who could not restrict below the maximum lag length they tried. This result is consistent with the theory that suggests an infinite lag is appropriate. Also, a model with no gestation lag (the third model in table 2) is preferred according to AIC, while a model with a 2-year gestation lag (the fourth model in table 2) is preferred according to APC.

Second, in the same model, when δ is not restricted and the lagged dependent variable is included in the model, OLS estimates are inconsistent. The transformation used to eliminate the unobserved knowledge stock induces problems of serial correlation in (8), even when there is no serial correlation in the underlying linear productivity function. Consequently, the instrumental variables (IV) technique was used to estimate equation (8)

¹⁵ With our 42-year series of *MFP*, 35 lags on research are feasible to estimate, but this leaves us with only two degrees of freedom in the most general case. Testing for longer lags is possible when we impose restrictions on the form of the lag.

with a linear trend as the instrument for lagged *MFP*. Based on the IV estimates, both the APC and AIC criteria support a specification with only 4-18 lags of research when the lagged dependent variable is also included. The coefficient on lagged *MFP* is 0.79, and is statistically significantly different from zero. This coefficient is, however, rather imprecisely estimated since it would be consistent with annual depreciation rates as large as 30 percent, and as small as zero.

In short, when we impose the restriction $\delta = 1$, the results are not consistent with a finite lag of less than 30 years and the results are consistent with the theory that suggests infinite lags of research investments affect current production. Furthermore, the results indicate that the restriction that $\delta = 1$ should not be imposed. With free-form lags, when we do not impose $\delta = 1$ and, accordingly, the lagged dependent variable is included in the model, a finite lag of much less than 30 years may be used. Of course, such a model effectively includes an infinite lag linking productivity to R&D spending.

One of the disadvantages of free-form lags is that the individual coefficients may not make sense—for example, when many of the individual lag weights are negative numbers. And because of multicollinearity problems, neither the individual lag coefficients nor the sum of these coefficients is very precisely estimated. While this is not so much of a problem when testing hypotheses about lag length, we cannot place

Table 2 U.S. agricultural productivity models with free-form research lags

Explanatory Variables	Research Lags in Years				
	0-0	0-0	0-30	2-30	4-18
<i>CONSTANT</i>	183.5 (61.3)	7.0 (15.1)	-1546.7 (143.7)	-1547.7 (130.4)	-351.01 (312.4)
PRC_t	-0.3938 (0.8256)	0.5565 (0.1433)	0.6306 (0.2902)	0.6384 (0.2226)	0.5198 (0.1525)
PRC_{t-1}		-0.6361 (0.1412)			0.6041 (0.1746)
MFP_{t-1}		1.0121 (0.0281)			0.7905 (0.2104)
$\Sigma_s b_s$			114.8 (12.7)	114.4 (11.8)	27.9 (24.0)
d.f.	40	38	9	11	23
R ²	0.006	0.973	0.995	0.995	0.989
SSE	60,116	1,641	289	323	674
APC	1,576	47.3	57.4	51.1	42.5
AIC	7.36	3.86	3.50	3.52	3.68
$\delta = 1$		reject			reject
$b_s = 0 \forall s$			reject	reject	reject

Note: Standard errors in parentheses.

much confidence in estimates of research benefits based on these imprecise lag coefficients. For similar reasons, most studies impose restrictions on the lag profile such that the lag weights are positive and change smoothly, and can be represented by a small number of parameters. Likewise, in the next set of models reported in table 3, we use Almon polynomial distributed lags to restrict the *shape* of the lag distribution, and we tested lag length and zero endpoint restrictions.

Many modeling options were tried that could be reported, but in the interests of brevity the table includes only a representative selection of key results. Tests of endpoint restrictions were F tests, and the AIC and APC criteria were used to select the preferred order of the polynomial and the lag length.

When $\delta = 1$ is imposed, the fit of the model is dramatically worse at shorter lag lengths, as can be seen by comparing the first five columns in table 3 from left to right. Regardless of the order of the polynomial (anywhere from second to eighth) or the maximum lag length allowed (anywhere from 20 to 50 years), the restriction that the terminal coefficient be zero was always rejected for lags of fewer than 45 years. And, when fewer than 30 lags of research were used, the restriction that *either* endpoint coefficient could be restricted to equal zero was strongly rejected. Considering only the most-commonly used, second-order polynomials, both the AIC and APC criterion suggested no gestation lag, a lag length of 35 years, and no end-point constraints. When considering higher-order polynomials, both the AIC and the APC criteria supported a fourth-order polynomial with a gestation lag of four years and either 45 or 50 lags of research. Using these criteria we could not distinguish between these two terminal lag

Table 3 U.S. agricultural productivity models with polynomial distributed lags

Order of Lag	4	4	2	2	2	8	2	2	2
Lag Length	4-50	4-50	0-35	0-30	0-20	3-43	4-18	0-30	0-20
Endpoint Constraint	none	final	none	none	none	two	two	two	two
<i>CONSTANT</i>	45.78 (11.02)	51.59 (10.33)	71.16 (9.87)	66.06 (13.88)	67.54 (37.48)	57.57 (31.95)	7.586 (15.48)	11.91 (16.60)	9.20 (15.79)
PRC_t	0.5242 (0.1110)	0.4974 (0.1108)	0.3927 (0.1219)	0.3698 (0.1664)	0.0811 (0.3991)	0.5050 (0.1144)	0.5526 (0.1475)	0.5637 (0.1443)	0.5432 (0.1468)
PRC_{t-1}						-0.2065 (0.2835)	-0.6369 (0.1428)	-0.6317 (0.1393)	-0.6384 (0.1417)
MFP_{t-1}						0.1338 (0.5184)	1.0080 (0.0406)	0.9674 (0.0749)	0.9950 (0.0480)
$\Sigma_s b_s$	11.81 (1.90)	9.23 (0.39)	7.98 (0.33)	7.86 (0.46)	7.92 (1.35)	7.93 (4.63)	0.04 (0.35)	0.34 (0.54)	0.15 (0.37)
d.f.	35	36	37	37	37	31	37	37	37
R ²	0.986	0.985	0.981	0.965	0.799	0.988	0.973	0.974	0.973
SSE	833	879	1,149	2,140	12,132	747	1,633	1,549	1,606
APC	27.7	27.9	34.8	64.7	366.9	31.5	51.5	48.8	50.6
AIC	3.32	3.33	3.55	4.17	5.90	3.50	3.99	3.94	3.98
$\delta = 1$						not reject	reject	reject	reject
$b_s = 0 \forall s$	reject	reject	reject	reject	reject	reject	not reject	not reject	not reject

Note: Standard errors in parentheses.

lengths, and the restriction that the lag coefficients terminate at zero could not be rejected.

The last four columns in table 3 refer to models including the lagged dependent variable. When the lagged dependent variable was included, and only 4-18 lags of research (as suggested by the free-form lag tests) were fitted with a second-order polynomial, we could not reject zero endpoint constraints.¹⁶ Moreover, the coefficient on the lagged dependent variable is significantly different from zero and more precisely estimated once the research lag coefficients are restricted to follow the polynomial form. The implied estimate of δ is consistent with a depreciation rate of zero, or with a positive value but not as large as 10 percent. In this restricted model, however, the research coefficients are jointly not significant. The preferred polynomial lag structure in the model that includes the lagged dependent variable is an eighth-order polynomial with 3-43 lags (column 6, table 3). The most important difference between the polynomial lag models with 4-18 versus 3-43 lags is that, with the inclusion of the additional lags of research, the lagged dependent variable does not make a statistically significant contribution to the regression.

Models with the lagged dependent variable and any order polynomial reject the hypothesis of $\delta = 1$ when 30 or fewer lags of research are considered. These results are more consistent with an infinite lag between research and productivity.¹⁷ If we were to

¹⁶ We also could not reject endpoint constraints when longer lags of research were allowed in the model estimated with lagged dependent variables.

¹⁷ In fact, in this reduced form model it is difficult to distinguish depreciation effects from other elements of the lag structure. The 50-year lag may be an approximation of an infinite lag between research and production, reflecting a much shorter lag between research and increments to knowledge and some depreciation of the knowledge stock.


impose $\delta = 1$ (i.e., drop both the lagged dependent variable and lagged *PRC* from the model), an even longer maximum research lag is implied (perhaps 4-50 lags and a fourth-order polynomial, as in the first two columns of table 3). In sum, both the AIC and APC criteria support a fourth-order Almon polynomial model with four gestation lags and at least 45 research lags when the lagged dependent variable is omitted.

Taking 50 lags of research as given and imposing $\delta = 1$, allowing different orders of polynomial lags did not imply materially different patterns in the lag profile. Figure 1 plots the implied lag weights against years for Almon polynomial lags from second- to eighth-order, with a four-year gestation lag and a maximum lag of 50 years, and terminal endpoint restrictions imposed. The overall shape is strikingly similar across the models. Two notable features are the early spike, at around five years (which becomes taller with higher-order, more flexible, polynomials), and a second peak, at around 30 years. The darkest line, for a fourth-order polynomial (column 2 in table 3), is in the middle of the range and represents the statistically preferred model in this set. At the longest lags, for the most flexible models, a third peak appears but we cannot reject the restriction that the individual coefficients after 45 lags are jointly zero, nor an endpoint restriction that makes the last coefficient zero. Thus, all of these models have the satisfactory characteristic of implying plausible lag distributions, with all weights positive and changing smoothly. The bimodal feature could be rationalized as a reflection of the aggregate R&D flow including more-basic research, having longer lags, and more-applied research and extension having

Thus we cannot say with this model that knowledge does not depreciate. More generally, interpreting an estimate of the parameter on the lagged dependent variable as one minus the knowledge depreciation rate is probably overinterpreting.

its main effect within a few years. But drawing such specific inferences may be overinterpreting the results. The polynomial structure does, after all, impose a shape.

Figure 1 Effects of order of the polynomial on estimated lag weights



In order to allow comparisons with previous literature on agricultural R&D, a third set of models was estimated in which a trapezoidal research lag structure, as assumed by Huffman and Evenson (1989b), was used to construct a research variable ($RESVAR_t$), with MFP_t as the dependent variable, with and without the lagged dependent variable. The

results in table 4 have several aspects in common with those from the corresponding free-form and polynomial lag models. When the lagged dependent variable is included, its coefficient is not significantly different from zero; owing to multicollinearity between the lagged dependent variable and the constructed research variable, no individual coefficients are significant in a model that nevertheless explains a significant portion of the variance in *MFP*. Interestingly, a model with the lagged dependent variable but excluding the research stock (shown in table 2, column 2) actually fits better than a model including the research stock, constructed using the trapezoidal lag structure, but not the lagged dependent variable. A reasonable interpretation of this may be that the trapezoidal research stock does a worse job of representing the true distributed lag than the simplest (and most restrictive) of all alternatives, the lagged dependent variable. Both of these models represent all past research investments with only one free parameter.

Across the columns in table 4, the models differ only with respect to the depreciation rate of the knowledge stock. In the first column, as in Huffman and Evenson (1989b), the depreciation rate is 100 percent; in the second column, δ is free but the point estimate is not statistically significantly different from one. In the last two columns, where the depreciation rate is imposed as 5 percent or 10 percent, the models fit better

Table 4 U.S. agricultural productivity models with trapezoidal lags

Lag length	3-33	3-33	3-33	3-33
Depreciation	$\delta = 1$	δ free	$\delta = 0.05$	$\delta = 0.10$
<i>CONSTANT</i>	-1,513.0 (53.24)	-1,464.3 (2,392.7)	-95.41 (42.58)	-171.6 (41.69)
<i>PRC_t</i>	0.3368 (0.1585)	0.3750 (0.3338)		
<i>PRC_t-(1-δ)PRC_{t-1}</i>			0.5922 (0.1059)	0.5905 (0.1062)
<i>PRC_{t-1}</i>		-0.2581 (0.6346)		
<i>MFP_{t-1}</i>		0.0145 (1.621)	0.95 (0.00)	0.90 (0.00)
<i>RESVAR_t</i>	6.4811 (0.1985)	6.3447 (10.32)	0.4094 (0.1680)	0.7314 (0.1644)
$\Sigma_s b_s$	119.90 (3.67)	117.38 (190.8)	7.58 (3.11)	13.53 (3.04)
d.f.	39	37	39	39
R ²	0.965	0.968	0.974	0.975
SSE	2,123	1,951	1,551	1,487
APC	58.3	59.0	42.6	40.8
AIC	4.07	4.08	3.75	3.71
<i>b_s = 0 \forall s</i>	reject	reject	reject	reject

Notes: Standard errors in parentheses.


according to all criteria, lending support to the idea that long, perhaps infinite, lags are appropriate. Even still, the SSEs from the models with trapezoidal lags are twice the size of those for the polynomial lag models in table 3 with the preferred longer lags (3-43 or 4-50 years).

In figure 2, the implied lag weights from the trapezoidal research lag model with 3-33 lags are plotted against the corresponding weights from the fourth-order polynomial with 4-50 lags. The dotted lines around each lag distribution represent 95 percent confidence bands. The contrast is striking. The trapezoidal model imposes a very different shape to the lag profile than the data support with a more flexible model and, as a result, much greater weight is given especially to the shorter lags—precisely as anticipated by Alston and Pardey (1996). Surprisingly, perhaps, the sum of the lag coefficients across all the lags, not just the early ones, is much greater from the trapezoidal structure. The implications for estimated rates of return are obvious and are quantified below.

RATES OF RETURN

To evaluate the proposition that the inappropriate truncation of research lags, and other restrictions, have contributed to overstated estimates of rates of return to research, we computed the internal rates of return to research implied by the alternative models reported above. Many studies that report internal rates of return do not state clearly how the rates were computed, and in some cases the results and their interpretation may be significantly affected (e.g., comparing marginal and average rates of return, for instance).

Figure 2 Polynomial versus trapezoidal lag weights



The real marginal internal rate of return to a change in research investment in year t is given by computing the value of r_t that solves¹⁸

$$0 = \sum_{s=0}^{\infty} \frac{V_{t+s}}{MFP_{t+s}} b_s (1 - r_t)^s - VR_t. \quad (9)$$

where MFP_{t+s} is multifactor productivity in year $t+s$, V_{t+s} is the real value of production in year $t+s$, b_s is the research lag weight for the effect of research in year t on productivity in year $t+s$, and VR_t is the real cost (1995 dollars) of research in year t . This rate of return can be computed for every year, t , after the first year in which the annual net benefits are positive. This requires projecting the stream of multifactor productivity, MFP_{t+s} , and real value of production, V_{t+s} , over the indefinite future. Recognizing that the aggregate index

¹⁸ See Appendix A for the derivation of these formulas.

of input quantities has been essentially constant (see Craig and Pardey 1996), the growth rates of the aggregate output quantity and productivity have been approximately equal, and we can use the approximation:

$$0 = \frac{V_t}{MFP_t} \sum_{s=0}^{\infty} b_s (1 - r_t)^s - VR_t. \quad (10)$$

A similar approximation was used by Huffman and Evenson (1989b) and Evenson (1996).

This formula can be used to compute the rate of return at any particular data point, at the mean of the data, or using the mean values of V , MFP , and VR . We computed all these possibilities but chose to report the rates of return computed using the mean values. The lagged dependent variable models imply an infinite stream of benefits arising from the change in research spending; for these we truncated the stream at 100 years.

Table 5 shows the estimates of the marginal internal rate of return implied by alternative models both with and without the lagged dependent variable, and using three alternative research lag structures, free-form, polynomial, and trapezoidal, computed using equation (10). In each case we report the results using the “preferred” lag length for the model in question, except for the trapezoidal model, for which we use the lag structure that was imposed in the estimation. To illustrate the effects of truncation of the lags and the choice of order of the polynomial, we also report results from polynomial lag models with a range of maximum lags and orders.

The first column of rates of return in table 5 shows the results from models imposing $\delta = 1$, so that the lagged dependent variable was not included, and only a finite lag of R&D expenditures would affect productivity. We first consider the effects of

different lag shapes and lengths within this set of models that include finite lags. The model with 0-30 years of free-form lags yielded an internal rate of return of 33.9 percent per annum, but this is a statistically unfavored model. The next set of results were obtained from models in which polynomial shapes were imposed to save degrees of freedom, allow longer lags, and assure smoothness in the finite research lag coefficients. First, we tried a typically short lag (0-20 years) in a second-order polynomial with no end-point restrictions, and no lagged dependent variable—albeit a longer lag than used by Davis (1980), Thirtle and Bottomley (1988), and Leiby and Adams (1991) in a second-order model. The estimated rate of return here is 22.5 percent, but the data reject that model in favor of models with longer lags included either explicitly, or implicitly by including the lagged dependent variable. Increasing the maximum lag length from 20 to 30 and then 35 years resulted in substantial reductions in the estimated rate of return, to 8.1 percent per annum, and 6.5 percent per annum, respectively. Clearly lag length matters. These models were all rejected in favor of models that have longer lags, and even lower rates of return to research.

Table 5 Real marginal internal rates of return to U.S. agricultural R&D

Lag Specification			Number of Lag Parameters	Rates of Return	
Form	Order	Length		$\delta = 1$	δ free
		<i>years</i>		<i>-- percent per annum --</i>	
Free Form		0-30	31	33.9	
		4-18	14		147.0
Polynomial	2	0-20	3	22.5	3.9
	2	0-30	3	8.1	4.1
	2	0-35	3	6.5	
	2	4-50*	2	4.5	
	4	4-50*	4	6.8	
	6	4-50*	6	8.6	
	8	4-50*	8	9.3	
Trapezoidal	8	3-43 [†]	7		6.8
		3-33 [†]	1	47.8	47.6 [‡]

NOTES: δ is the knowledge depreciation rate.

* Terminal endpoint constraint imposed.

[†] Initial and terminal endpoint constraints imposed.

[‡] The point estimate of δ was almost 1.0, here. When δ was restricted to be 0.05 or 0.10 in the econometric estimation, the resulting rates of return were computed as 24.3 and 29.2 percent per annum, respectively.

The next four estimates in the same column show the effect of changing the order of the polynomial, holding the lag length constant at 4-50 years, with a zero terminal endpoint restriction imposed. The estimated rate of return increases monotonically with increases in the order of the polynomial, from 4.5 percent per annum with a second-order polynomial to 9.3 percent per annum with an eighth-order polynomial lag. Among these models with 4-50 research lags, the fourth-order polynomial is statistically preferred, and it implies a rate of return of 6.8 percent per annum. The last row shows the rates of return

estimated when the trapezoidal lag was imposed. This model yields a rate of return of 47.8 percent per annum, very close to the rate of return reported by Huffman and Evenson (1989b).¹⁹

Next, in the second column of rates of return in table 5, we show rates of return from models including the lagged dependent variable (δ free); models that imply an infinite lag effect of research on productivity. First, consider the model with free-form lags. The statistically preferred version of this model includes 4-18 years of research lags, and the estimated rate of return is a whopping 147 percent per annum. This is a very high rate of return, but not one that we can view with much confidence given the imprecise estimate of all the parameters in this model.²⁰ On the other hand, with the short second-order polynomial lags (0-20 or 0-30 years), when the lagged dependent variable is included, the estimated internal rate of return is a modest 3.7 or 4.1 percent per annum. Among the polynomial lag models when the lagged dependent variable is included, the statistically

¹⁹ Huffman and Evenson (1989b) used different data, covering a different (earlier) time period, and their model included other factors, such as education and private research, as explanatory variables in their *MFP* model, whereas we dealt with these aspects implicitly in developing *MFP* estimates. Although the two sets of estimates cannot be compared directly, for these reasons, it is notable that when we use the same lag structure as used by Huffman and Evenson, our estimates are of a similar magnitude to theirs, and much greater than we obtain with our statistically preferred models with longer and more flexible lag structures

²⁰ We also tried a second-order polynomial lag model with 4-18 lags (as preferred for the free-form lag model with the lagged dependent variable), and the estimated rate of return was essentially zero percent per annum. We tried this lag length in the polynomial form just for purposes of comparison, and the results show that the imposition of the smooth second-order polynomial form on the free-form lag model has dramatic effects on the estimated rate of return. Our data do not favor either the free form or polynomial models with these lag lengths.

preferred version is an eighth-order polynomial lag model, with 3-43 research lags, and initial and terminal endpoint restrictions imposed. The resulting rate of return is 6.8 percent per annum, the same as from the preferred fourth-order polynomial lag model when the lagged dependent variable is not included.

Finally, when the lagged dependent variable is included, the trapezoidal lag model still generates a very high rate of return, of 48 percent per annum. The coefficient on the lagged dependent variable in this case is not statistically significant (and neither is the coefficient on the trapezoidal lag, research stock variable), but the implied point estimate of δ is rather large (almost 1.0). We also estimated trapezoidal models with δ imposed as either 0.10 or 0.05, which are much more plausible rates of depreciation of the knowledge stock. The data do not reject these restrictions, and the resulting rate-of-return estimates are much smaller than with $\delta = 1$ or δ free, 29.2 and 24.3 percent per annum, respectively.

To sum up, the weight of the econometric evidence points to a marginal internal rate-of-return of 6.8 percent per annum. The same estimate is obtained with a fourth-order polynomial (4-50 lags and not including the lagged dependent variable) and an eighth-order polynomial (3-43 lags and including the lagged dependent variable).²¹

The estimated rate of return depends on choices in how to compute internal rates of return, including whether it is approximated or simulated, whether it is marginal or average, or which year of investment it applies to. We computed the rate of return for every year using data observations from 1948-1991, and the estimated rates of return from

²¹ Even smaller estimates were obtained when we followed the advice of Fox (1985) and made allowance for a 20 percent excess marginal burden of taxation, but in our application the rates were only slightly smaller.

equation (10) declined progressively. This is to be expected given the growth in real expenditure on agricultural R&D (see appendix A). For comparison, we also computed marginal internal rates of return by simulating the stream of benefits through 1991 from a one percent increase in R&D spending in 1948. The resulting rates of return for the simulation are, as expected, much greater than the returns found using the approximation in equation (10) at the data means. When the approximation uses 1948 data, the computed rates of return are very similar for the simulated and approximation approaches (about 30 percent per annum for our preferred polynomial lag models, and almost 100 percent per annum with the trapezoidal lags). This suggests that equation (10) may be providing an adequate approximation. The rate of return also depends on the size of the change in research spending. To illustrate this we computed “average” rates of return by simulating a 100 percent decrease in R&D spending in 1948. Because of diminishing returns, the estimated “average” internal rate of return to R&D spending in 1948 was substantially greater than the corresponding marginal rate.

5. CONCLUSION

We have developed arguments about how to structure research lags meaningfully, and examined the implications for estimates of returns to research. The ideas developed here have been virtually absent from the agricultural R&D literature and, although some of the same conceptions have been raised in some of the more recent industrial R&D literature, they have not often been carried through systematically to the empirical parts of

that literature (and when the ideas have been carried through, the empirical consequences have usually been imposed).

Our theoretical arguments led to a view that a given research investment could affect production for much longer than most studies have allowed. Our econometric results consistently indicated that we should prefer a model with infinite lags, to a more conventional model with finite lags, unless we can include many more lags than previous studies have done.

We speculated that the truncation of the lags in previous studies has biased the rates of return upward, and our results confirm that view. Our results also suggest that many previous studies have unduly restricted the shape of the research lag profile. Although free-form lags are likely to be precluded by data constraints and a preference for smooth and precisely estimated lags, many studies have effectively estimated only one parameter to represent the effects of the past 20, 30, 40, or even 50 years of research. Our data support either a model with four free parameters, a fourth-order Almon polynomial with one endpoint restriction and 4-50 lags of research without the lagged dependent variable, or a model with seven parameters in an eighth-order polynomial with initial and terminal endpoint restrictions, 3-43 years of research, and the lagged dependent variable included.

The preferred models indicate a real, marginal rate of return to public agricultural research in the United States of 6.8 percent per annum. This is well below most previously reported results. A part of the story must be that many of the previous studies have unduly truncated the lags, and truncating the lags can have profound effects on the

computed rate of return. Also, a number of prominent and influential studies have used a trapezoidal lag shape, which results in rates of return of almost 50 percent per annum in our data set, where the statistically preferred models yield rates of return of less than 10 percent per annum. Different approaches used to compute rates of return may also have had important effects.

Our results make clear that the biases from restricting the lags can be very large. Perhaps our estimates have been biased for other reasons, such as the omission of technology spillovers from non-agricultural industries or from overseas agriculture, along with private research, from the productivity model—although the input data used to construct our multifactor productivity measure were adjusted explicitly for quality changes in ways that would mitigate some consequences of omitting private R&D and certain types of technology spillovers. By the same token, our benefit measures exclude the spillovers to non-agricultural uses in the United States and overseas agriculture. We hope to explore these and other possibilities using the state-level production data for U.S. agriculture developed by Craig and Pardey (1996). State-level (panel) data may allow greater potential for identifying the research lag structures, but will also introduce an additional set of problems—how to handle interstate research spillovers. Even so, state-level U.S. data, or data from other countries, will provide additional opportunities to further evaluate the issues we have raised and illustrated with the national U.S. time series.

REFERENCES

- Adams, J.D. 1990. Fundamental stocks of knowledge and productivity growth. *Journal of Political Economy* 98 (4): 673-702.
- Akgüngör, S., D.W. Makanda, J.F. Oehmke, R.J. Myers, and Y.C. Choe. 1996. Dynamic analysis of Kenyan wheat research and rate of return. Contributed paper proceedings from the Conference on Global Agricultural Science Policy for the Twenty-First Century, August 26-28, Melbourne, Australia, pp. 1-32.
- Alston, J.M. and P.G. Pardey. 1996. *Making science pay: Economics of agricultural R&D policy*. Washington D.C.: American Enterprise Institute for Public Policy.
- Alston, J.M., B.J. Craig, and P.G. Pardey. 1996. Research lags and research returns. Invited paper proceedings from the Conference on Global Agricultural Science Policy for the Twenty-First Century, August 26-28, Melbourne, Australia, pp. 333-366.
- Alston, J.M., G.W. Norton, and P.G. Pardey. 1995. *Science under scarcity: Principles and practice for agricultural research evaluation and priority setting*. Ithaca: Cornell University Press.
- Alston, J.M., P.G. Pardey and H.O. Carter. 1994. Valuing California's agricultural research and extension. Publication No. VR-1 (March). Davis, Ca.: University of California Agricultural Issues Center.
- Alston, J.M., M.C. Marra, P.G. Pardey, and T J Wyatt. 1997. Research returns redux: A meta-analysis of agricultural R&D evaluations. IFPRI Research Report. Washington D.C.: International Food Policy Research Institute (draft).
- Boulding, K.E. 1966. The economics of knowledge and the knowledge of economics. *American Economic Review* 56 (Proceedings 1966): 1-13.
- Boyce, J.K. and R.E. Evenson. 1975. *Agricultural research and extension programs*. New York: Agricultural Development Council.
- Bredahl, M. and W. Peterson. 1976. The productivity and allocation of research: U.S. agricultural experiment stations. *American Journal of Agricultural Economics* 58: 684-692.
- Bureau of Economic Analysis. 1994. A satellite account for research and development. In *Survey of Current Business*, pp. 37-71 (November). Washington, D.C.: Department of Commerce.

- Bureau of Labor Statistics (BLS). 1989. The impact of research and development on productivity growth. BLS Bulletin No. 2331 (September). Washington, D.C.: Department of Labor.
- Caballero, R.J. and A.B. Jaffe. 1993. How high are the giants' shoulders: An empirical assessment of knowledge spillovers and creative destruction in a model of economic growth. *Macroeconomics Annual 1993*, pp. 15-74. Cambridge: National Bureau Economic Research.
- Chavas, J.P. and T.L. Cox. 1992. A nonparametric analysis of the effects of research on agricultural productivity. *American Journal of Agricultural Economics* 74 (August): 583-591.
- Coe, D.T. and E. Helpman. 1993. International R&D spillovers. IMF Working Paper No. WP/93/84 (November). Washington, D.C.: International Monetary Fund.
- Craig, B.J. and P.G. Pardey. 1996. Input, output, and productivity developments in U.S. agriculture. Contributed paper proceedings from the Conference on Global Agricultural Science Policy for the Twenty-First Century, August 26-28, Melbourne, Australia, pp. 193-216.
- Davis, J.S. 1980. A note on the use of alternative lag structures for research expenditure in aggregate production function models. *Canadian Journal of Agricultural Economics* 28: 72-76.
- Echeverría, R.G. 1990. Assessing the impact of agricultural research. In *Methods for diagnosing research system constraints and assessing the impact of agricultural research, Vol.II: Assessing the impact of agricultural research*, ed. R.G. Echeverría. The Hague: ISNAR.
- Evenson, R.E. 1967. The contribution of agricultural research to production. *Journal of Farm Economics* 49 (December): 1415-1425.
- Evenson, R.E. 1968. The contribution of agricultural research and extension to agricultural production." Unpublished Ph.D. diss. University of Chicago, Illinois.
- Evenson, R.E. 1996. Two blades of grass: Research for U.S. agriculture. Chapter 11 in *The economics of agriculture, volume 2: Papers in honor of D. Gale Johnson*, ed. J.M. Antle and D.A. Sumner, pp. 171-203. Chicago: University of Chicago Press.
- Evenson, R.E., P.E. Waggoner, and V.W. Ruttan. 1979. Economic benefits from research: An example from agriculture. *Science* 205: 1101-1107.
- Fox, G. 1985. Is the United States really underinvesting in agricultural research?

- American Journal of Agricultural Economics* 67 (November): 806-812.
- Fuglie, K., N. Ballenger, K. Day, C. Klotz, M. Ollinger, J. Reilly, U. Vasavada, and J. Yee. 1996. Agricultural research and development: Public and private investments under alternative markets and institutions. Agricultural Economic Report No. 735, Economic Research Service, USDA, May. Washington D.C.: U.S. Department of Agriculture.
- Fulginiti, L.E. and R.K. Perrin. 1993. Prices and productivity in agriculture. *Review of Economics and Statistics* 75 (August): 471-482.
- Greene, W.H. 1993. *Econometric analysis*. 2d ed. New York: Macmillan Publishing Company.
- Griliches, Z. 1964. Research expenditures, education, and the aggregate agricultural production function. *American Economic Review* 54 (6) (December): 961-974.
- Griliches, Z. 1979. Issues in assessing the contribution of research and development to productivity growth. *Bell Journal of Economics* 10 (Spring): 92-116.
- Griliches, Z. 1980. Returns to research and development expenditures in the private sector. Chapter 8 in *New developments in productivity measurement and analysis*, ed. J. Kendrick and B. Vaccara. Chicago: University of Chicago Press (for the National Bureau of Economic Research).
- Griliches, Z. 1986. Productivity, R&D, and basic research at the firm level in the 1970s. *American Economic Review* 76 (March): 141-154.
- Griliches, Z. 1996. R&D and productivity: The unfinished business. Invited paper proceedings from the Conference on Global Agricultural Science Policy for the Twenty-First Century, August 26-28, Melbourne, Australia, pp. 1-20.
- Hall, B.H. and J. Mairesse. 1995. Exploring the relationship between R&D and productivity in French manufacturing firms. *Journal of Econometrics* 65: 263-293.
- Huffman, W.E. and R.E. Evenson. 1989a. The development of U.S. agricultural research and education: An economic perspective, Part IV. Department of Economics Staff Paper No. 174, Iowa State University.
- Huffman, W.E. and R.E. Evenson. 1989b. Supply and demand functions for multiproduct U.S. cash grain farms: Biases caused by research and other policies. *American Journal of Agricultural Economics* 71 (August): 761-773.

- Huffman, W.E. and R.E. Evenson. 1992. Contributions of public and private science and technology to U.S. agricultural productivity. *American Journal of Agricultural Economics* 74 (August): 752-756.
- Huffman, W.E. and R.E. Evenson. 1993. *Science for agriculture: A long-term perspective*. Ames: Iowa State University Press.
- Jorgenson, D.W. 1973. The economic theory of replacement and depreciation. In *Econometrics and Economic Theory*, ed. W. Sellekates, pp. 189-281. New York: Macmillan.
- Leiby, J.D. and G.D. Adams. 1991. The returns to agricultural research in Maine: The case of a small northeastern experiment station. *Northeastern Journal of Agricultural and Resource Economics* 20: 1-14.
- Makki, S.S., L.G. Tweeten, and C.S. Thraen. 1996. Returns to agricultural research: Are we assessing right? Contributed paper proceedings from the Conference on Global Agricultural Science Policy for the Twenty-First Century, August 26-28, Melbourne, Australia, pp. 89-114.
- Martin, W. and J.M. Alston. 1997. Producer surplus without apology? Evaluating investments in R&D. *The Economic Record* 73 (221) (June): 146-158.
- Minasian, J.R. 1969. Research and development, production functions, and rates of return. *American Economic Review* 59: 80-85.
- Myers, R.J. and T. Jayne. 1996. Regime shifts and technology diffusion in crop yield growth paths: An application to maize yields in Zimbabwe. Contributed paper proceedings from the Conference on Global Agricultural Science Policy for the Twenty-First Century, August 26-28, Melbourne, Australia, pp. 63-88.
- Nadiri, M.I. and I.R. Prucha. 1993. Estimation of the depreciation rate of physical and R&D capital in the U.S. total manufacturing sector. National Bureau of Economic Research Working Paper No. 4591 (December). Washington, D.C.: NBER.
- Norton, G.W., J.D. Coffey, and E.B. Frye. 1984. Estimating returns to agricultural research, extension, and teaching at the state level. *Southern Journal of Agricultural Economics*: 121-128.
- Pakes, A. and M. Shankerman. 1987. The rate of obsolescence of patents, research gestation lags, and the private rate of return to research resources. Chapter 4 in *R&D, patents, and productivity*, ed. Zvi Griliches. Report of the National Bureau of Economic Research. Chicago: University of Chicago Press.

Pardey, P.G. and B. Craig. 1989. Causal relationships between public sector agricultural research expenditures and output. *American Journal of Agricultural Economics* 71 (February): 9-19.

Ravenscraft, D. and F.M. Scherer. 1982. The lag structure of returns to research and development. *Applied Economics* 14: 603-620.

Schimmelpfennig, D. and C. Thirtle. 1994. Cointegration, and causality: Exploring the relationship between agricultural R&D and productivity. *Journal of Agricultural Economics* 45: 220-231.

Thirtle, C.G. and P. Bottomley. 1988. Is publicly funded agricultural research excessive? *Journal of Agricultural Economics* 31: 99-111.

Appendix A: Approximating the Internal Rate of Return

Assume a linear productivity function with no shifters:

$$MFP_t = \beta_0 + \beta_1 K_t \quad (A.1)$$

With an infinite lag of logarithms of research used to construct the knowledge stock:

$$MFP_t = \beta_0 + \sum_{s=0}^{\infty} b_s \ln R_{t-s}, \quad (A.2)$$

The effects of a one-shot increase in research investments in time t on changes in the stream of MFP s are given by taking the total derivative of equation (A.2), for any future year $t+n$,

$$dMFP_{t+n} = b_n d\ln R_t \quad (A.3)$$

Benefits in year $t+n$ equal the change in the value of output, V , from the change in the MFP .

Using $dV = V(d\ln V) = V(d\ln MFP) = V(dMFP/MFP)$,

$$B_{t+n} = dV_{t+n} = V_{t+n} \frac{dMFP_{t+n}}{MFP_{t+n}} = V_{t+n} \frac{b_n d\ln R_t}{MFP_{t+n}}. \quad (A.4)$$

The costs of achieving the stream of benefits given by (A.4) are $dVR_t = VR_t d\ln R_t$, where VR_t is the cost of research in year t , the quantity of research, R_t multiplied by its price index, and, with a fixed price, the relative change in expenditure equals the relative change in quantity.

The marginal internal rate of return is the value of r_t that solves:

$$\begin{aligned}
 0 & \sum_{s=0}^{\infty} B_{t,s} (1 - r_t)^s \quad dVR_t \\
 & \sum_{s=0}^{\infty} \frac{V_{t,s}}{MFP_{t,s}} dMFP_{t,s} (1 - r_t)^s \quad VR_t d\ln R_t \\
 & \sum_{s=0}^{\infty} \frac{V_{t,s}}{MFP_{t,s}} b_s d\ln R_t (1 - r_t)^s \quad VR_t d\ln R_t \\
 & \sum_{s=0}^{\infty} \frac{V_{t,s}}{MFP_{t,s}} b_s (1 - r_t)^s \quad VR_t.
 \end{aligned} \tag{A.5}$$

Benefits equal a discounted sum of lag weights, multiplied by the real value of production divided by the value of *MFP* in each future year. If we were to assume a steady state as an approximation, and project the values of these variables in year *t* forward indefinitely,

$$0 \quad \frac{V_t}{MFP_t} \sum_{s=0}^{\infty} b_s (1 - r_t)^s \quad VR_t. \tag{A.6}$$

If the research intensity (*VR/V*) is high, or if the current productivity (*MFP*) is high, a smaller value of *r_t* will be implied. Thus, computing *r₁₉₄₈* (at *t* = 1948, the year with the lowest value of research expenditure) will yield a larger value than *r₁₉₉₁* (at *t* = 1991, with the highest value).