AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

ELSEVIER

Agricultural Economics 25 (2001) 141–152

AGRICULTURAL
ECONOMICS

www.elsevier.com/locate/agecon

Attribution and other problems in assessing the returns to agricultural R&D

Julian M. Alston^{a,*}, Philip G. Pardey^b

^a Department of Agricultural and Resource Economics, University of California, Davis, CA 95616, USA

^b International Food Policy Research Institute, Washington, DC 20006, USA

Abstract

Estimated rates of return to research are distorted by problems of attributing the credit for particular research results, or for particular research-induced productivity increases, among research expenditures undertaken at different times, in different places, and by different agencies. A comprehensive assessment of the evidence from past economic evaluations of the returns to agricultural R&D indicates that studies generally report high rates of return, with enormous variation among studies, but that much of this evidence has been tainted by inadequate attention to attribution problems. This paper raises these concerns in a general way and illustrates their importance with reference to two particular types of attribution problem. First, we consider the problem of accounting for locational spillovers in attributing varietal-improvement technology among research performers, using US wheat varieties as an example. Second, we consider the temporal aspects of the attribution problem using the specification of research lags in econometric models to illustrate the problem of attributing aggregate productivity gains to research expenditures made at different times. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Rates of return; Attribution; Agricultural R&D; Research spillovers

1. Introduction

It is widely believed that public-sector agricultural R&D has paid handsome dividends for society as a whole; contrary views are the exception, e.g. Pasour and Johnson (1982), and Kealey (1996). But even those who believe that agricultural R&D is (or has been) a good investment for society may be sceptical about some of the very high reported estimates of rates of return to research. An interest in the outcome might lead to biased estimates in some cases — rate of return estimates are often intended to be used to justify past investments and shore up support for future investments, and both implausibly high and unfavourable results are less likely to be acceptable for

this purpose. Rates of return are also likely to involve errors even when the analyst is disinterested because it is inherently difficult to identify which research investment was responsible for a particular productivity improvement (or, conversely, which parts of the productivity benefits are attributable to a particular research investment).

Consider an ex post analysis of the contribution of agricultural R&D by the California Agricultural Experiment Station (CAES) to current productivity in California. For such an analysis we want to be able to meaningfully measure productivity growth and then attribute it among those investments by the CAES, other public R&D investments by the California state government and by other states and the US Department of Agriculture (USDA), international R&D, and private R&D investments. Moreover, we have to attribute the productivity growth not just between the

* Corresponding author.

E-mail address: julian@primal.ucdavis.edu (J.M. Alston).

CAES research and the other elements at a point in time, but among these elements over time, including the distant as well as the recent past. We want to be able to say which research, conducted (or paid for) by whom, and, in particular, when, was responsible for a particular productivity improvement.

This *attribution* problem is difficult; it relates to the *appropriability* problem that underpins the in-principle argument for government involvement in research. Spillover effects of research, where research conducted by one firm (or state or country) yields benefits for free-riders, account for private-sector under-investment and the possibility of high social rates of return. If it were easy to attribute benefits to particular investments, it should be possible to devise institutions to make the benefits appropriable. Thus, the characteristics of research that give rise to the potential for high rates of return also give rise to measurement problems.

Griliches (1979) laid out an agenda for work on understanding the link between R&D and productivity change, which he reiterated in Griliches (2001). Taking inspiration from these and other contributions from Griliches, in this article we reassess the evidence on rates of return to agricultural research with an emphasis on the nature of the attribution problem and the likely implications of conventional evaluation methods. We suggest that the effects of attribution problems have not been neutral; on the whole, rate of return estimates are likely to have been biased up.

2. Overview of the literature

The literature on returns to agricultural research itself has been the subject of several studies. For instance, partial periodic tabulations and narrative reviews can be found in Evenson et al. (1979), Echeverría (1990), Alston and Pardey (1996), and Fuglie et al. (1996). Alston et al. (2000a,b) provide a comprehensive compilation, synthesis, and quantitative meta-analysis of rate of return estimates that reveals interesting and useful patterns. A total of 292 benefit–cost studies of agricultural R&D (including extension) were compiled, and these studies provide 1886 separate estimates of rates of return to agricultural R&D that range from small negative numbers to more than 700,000% per annum. This large range

reflects variation within groups (such as applied versus basic research, or research on natural resources versus commodities) more than among groups, and such large within-group variation makes it difficult to discern differences among groups. The estimated annual rates of return averaged 99.6% for research only, 47.6% for research and extension combined, and 84.6% for extension only. Moreover, the distributions are generally positively skewed, with a significant number of exceptionally high rates of return. Table 1 shows the ranges of rates of return and the mean, standard deviation, mode, and median rates of return for a sub-sample of the data, according to the nature and commodity orientation of the research and the geographic location of the research performer. The preponderance of studies reported the returns to all research (mainly returns to aggregate investments in agricultural R&D), while just over half the observations pertained to field crops research and research performed in developed countries.

The estimates in Table 1 predominantly refer to real (i.e. inflation adjusted), marginal (i.e. for incremental research expenditures), ex post (i.e. for past investments), internal rates of return (IRRs). The implication when reporting an IRR is that the benefits from the research are being evaluated as though they can be reinvested, along with the original investment, at the same rate of return. Since the benefits often accrue to farmers and consumers who typically do not have opportunities to invest at such high rates, it is worth dwelling briefly on what is implied by high IRRs. If the investment of US\$ 1.21 billion (bUS\$) in 1980 in US public agricultural R&D had earned an IRR of 48% per annum, the mean for the studies of US agriculture in aggregate, the accumulated stream of benefits would have been worth 3000 bUS\$ (1980 dollars) by the year 2000, 4.5 months worth of total US GDP, and more than 20 years worth of US agricultural GDP. This is the implied benefit from the investment in 1980 alone. Such calculations might give rise to doubts about whether the estimated rates of return really represent IRRs, or for that matter the true returns to research.

Of course, even though there is a unique true rate of return to any particular set of past investments, there is no such thing as *the* rate of return to agricultural research. In a typical agricultural research portfolio, some (perhaps most) investments yield no benefits whatsoever, while others yield very high returns,

Table 1
Rates of return to agricultural R&D^a

	Number of estimates (<i>count</i>)	Rates of return (% per year)					
		Mean	Standard deviation	Mode	Median	Minimum	Maximum
Nature of research							
Basic	30	79.2	88.7	75.0	61.3	−1.3	457
Applied	192	163.5	557.0	38.0	46.0	6.0	5645
All research	904	88.4	148.6	46.0	49.0	−7.4	1720
Research & extension	643	46.8	43.4	28.0	36.0	−100	430
Commodity orientation							
Multicommodity	436	80.3	110.7	58.0	47.1	−1.0	1219
Field crops	916	74.3	139.4	40.0	43.6	−100.0	1720
Livestock	233	120.7	481.1	14.0	53.0	2.5	5645
Tree crops	108	87.6	216.4	20.0	33.3	1.4	1736
Natural resources	78	37.6	65.0	7.0	16.5	0.0	457
Geographic location							
Developed countries	990	98.2	278.1	19.0	46.0	−14.9	5645
Developing countries	683	60.1	84.1	46.0	43.0	−100.0	1490
Multiregional	74	58.8	98.3	32.0	34.0	−47.5	677
IARC	62	77.8	188.6	26.0	40.0	9.9	1490
All studies	1772	81.2	216.1	46.0	44.0	−100.0	5645

^a Source: Adapted from Alston et al. (2000a; Tables 15 and 17). Sample excludes two extreme outliers and includes only returns to research and combined research and extension, so that the overall sample size is 1772.

sufficient to make the portfolio as a whole profitable. Even though very high rates of return are not implausible in every context, they are much less plausible for the more aggregated investments represented by extensive portfolios.

3. Measurement issues and problems

Problems with data, measurement error or misconceptions can result in an estimated rate of return that is higher or lower than the true value. One important problem is to define the relevant counterfactual alternative. In particular, to define what the world might be like in the absence of the particular research investment being evaluated, we have to take account of other things that might also be caused to change. Holding the right things constant is necessary to derive a stream of benefits that properly matches the stream of expenditures being evaluated.

Alston and Pardey (1996, Chapter 6) suggest that the estimated rates of return to R&D in the literature have tended to be over-optimistic, relative to the corresponding true values, because the commonly used

procedures understate the costs, overstate the benefits, and often predetermine the research lag structure (that relates changes in productivity to past investments in research) in ways that lead to higher estimated rates of return. Some other common practices might lead to understatement of benefits, so that a particular estimated rate of return may be too high or too low. In particular, the conventional estimates may exclude benefits from ‘maintenance’ research, benefits from disease prevention, food safety R&D, or social science research related to agriculture (some of which may not show up clearly in commodity markets and some of which are not captured in conventional productivity measures), and the spillover benefits from agricultural R&D into non-agricultural applications. On balance, however, we suspect that the tendency to overestimate has predominated.

3.1. Productivity measurement

The ex post evaluation of public agricultural R&D investments often begins with a consideration of agricultural productivity. At a minimum, we want to avoid measurement problems associated with inappro-

priate aggregation or indexing procedures. Index number problems can account for some errors in the measurement of productivity growth attributable to research, and aggregate productivity measures can be statistically sensitive to aggregation procedures (e.g. Acquaye et al., 2000).

As pointed out by Schultz (1956), growth in the use of conventional inputs does not account for much of the growth in agricultural output. A part of the attribution problem is to remove the effects of various other (non-research) factors before attempting to attribute residual productivity growth to particular research investments. Understanding the sources of the growth not attributable to conventional inputs is the first step in measuring the benefits from public R&D investments. Other factors, beyond conventional inputs, include changes in input quality, output quality, improvements in infrastructure, economies of size and scale, and improvements in technology. In addition, conventional productivity measures do not account for the consumption of unpriced or underpriced natural resource stocks, such as irrigation water, in the process of production. Rate of return studies that use conventional productivity indexes will tend to overstate the social value of technological changes that involve a faster rate of consumption of natural resource stocks, and will understate the benefits from technologies that involve greater environmental amenities or resource stock savings (e.g. Alston et al., 1995; Perrin and Fulginiti, 1996).

Schultz (1956) and Griliches (1963) demonstrated the important role of changes in input quality in accounting for measured productivity growth in agriculture. Yet many subsequent studies of returns to public-sector R&D have measured aggregate input quantities using index numbers that were not adjusted appropriately to account for changes in input quality. Such analysis overstates the productivity growth attributable to public-sector R&D by giving it credit for effects attributable either to schooling (from private or public investments in education unrelated to R&D) or to private R&D (in the case of embodied technological change). It is tricky to isolate the effects of schooling from the benefits of training in the context of research programs, a benefit that should be attributed to R&D. Following Griliches (1964), some studies have included additional explanatory variables to represent the effects of factors such as education, infrastructure,

or private R&D in models of productivity. Clearly, the appropriate adjustments of the dependent variable can be different, depending upon the explanatory variables other than public R&D that are to be included in the model to account for the effects of input and output quality, and so on. Craig and Pardey (1996, 2001) and Acquaye et al. (2000), among others, have shown that correcting for changes in input quality can have major implications for understanding changes in input use and productivity in US agriculture. Making adjustments for input quality change is likely to lead to a lower estimated rate of return to public-sector R&D and a better appreciation of the different roles played by private- and public-sector R&D (in agriculture and elsewhere), and education. Less is known about the quantitative effects of accounting for research-induced changes in output quality.

3.2. Matching benefits and costs: attribution among groups

Multifactor productivity is the measurable stream of output not accounted for by measured inputs. We can translate the productivity measures into measures of streams of research benefits using conventional procedures. The attribution of these benefits to particular inputs can be thought of as a two-step process. Having accounted for the contribution of factors other than R&D in the first step, a second step involves discerning the share of these residual benefits most appropriately attributed to research by a particular individual, program, state, nation, or other aggregate. This attribution problem can be thought of as matching streams of research benefits to corresponding streams of costs.

Understatement of public research costs arises in a number of ways. As pointed out by Fox (1985), a common source of understatement is not allowing for the full social cost of using government revenues for R&D. General taxation involves a social cost of more than one dollar per dollar raised, an excess burden (Ballard and Fullerton, 1992). Most studies have not adjusted for the effects of the excess burden of taxation on costs, an omission that will lead to a systematic understatement of the social costs and an overstatement of the social rate of return.

Occasionally studies of particular research programs fail to attribute an appropriate portion of R&D overhead (including the costs of associated basic

research and institutional overheads) to the particular projects being evaluated, or they omit components of the effort involved in the development and extension phases of a project. It is not easy to estimate costs attributable to total research (let alone research on a particular set of issues), but there seems to be a tendency to understate the cost of individual research programs, and research overall, through the tendency to omit or underestimate overhead costs.

Agricultural research consists of a continuum of activities, from basic science through to field extension work, that interact with and complement one another. To properly measure the contribution of one element of the whole, it is important to control for the effects of all of the others. Many previous studies have failed to take proper account of other elements and, as a result, they have tended to overestimate the gains in productivity attributable to a particular element of total expenditures on R&D. Equivalently, many studies have underestimated the total expenditure (that includes foreign and domestic, private and public, and basic and applied work and extension) required to achieve a particular productivity gain.

Overstatement of benefits sometimes arises from not counting the effects of private-sector R&D or spillovers of technology from elsewhere (states, countries, competing institutions, or other industries) and, instead, attributing all of the gains in productivity to only a part of the total relevant R&D spending. Griliches (1992) discussed the problems of accounting for R&D spillovers, and Griliches (1974) explored the role of private-sector R&D.

Private-sector research is often omitted from the analysis, or its effects are considered but not properly taken into account. This is a problem in econometric studies, in particular, where the omission of relevant explanatory variables can lead to biased estimates of the effects of variables included in the analysis. Private R&D expenditures (R^P) are likely to be positively correlated with public R&D expenditures (R^G), and, as a result, the omission of R^P from a productivity model can be expected to lead to an upwards bias in estimates of the coefficient on R^G . The confounding of effects extends beyond overstating the rate of return to R^G when we go beyond the consequences of statistical correlation and consider causal connections between the two types of expenditure and, perhaps, complementary or substitution interactions between

R^P and R^G in affecting productivity. The omission of private-sector R&D may also give rise to biased estimates in synthetic (benefit–cost) approaches (i.e. where productivity gains are deduced or assumed rather than statistically estimated), depending on how the growth in productivity attributable to public-sector R&D is estimated. Similar concerns arise in relation to the treatment of extension, spillovers from private- or public-sector research conducted elsewhere (e.g. overseas or in sectors other than agriculture), basic (or pretechnology) research that may underpin the applied research being assessed, and development work without which the commercial adoption of research results would not be possible.

R&D spillovers appear pervasive and confound the attribution of research benefits. Using firm-level data from the chemical industry, Mansfield (1977) reported that the returns to innovators (private rates of return) were significantly smaller than ‘social’ rates of return. More recently, Jaffe (1986) developed a patent-based metric of R&D ‘spillover pools’ to investigate firm-to-firm spillover effects. He found indirect but convincing econometric evidence of the existence of R&D spillovers, demonstrating that, on average, firms had higher returns to their own R&D (in terms of accounting profits or market value) if this R&D was conducted in areas where other firms do much research. Analogous firm-to-firm spillover effects are no doubt a feature of private research related to agriculture.

Agricultural economists also have been giving attention to economies of size, scale, and scope in agricultural R&D, and the related questions of the spatial spillovers of public agricultural research benefits (and costs), especially in recent years (e.g. Johnson and Evenson, 1999; Byerlee and Traxler, 2001). Econometric efforts to measure the spatial spillovers of agricultural research have used knowledge stocks computed as spatial aggregations of R&D based on geopolitical boundaries and geographic proximity rather than agroecological similarity (e.g. Huffman and Evenson, 1993). However, the pattern of geographical spillovers is largely conditioned by agroecological factors, although economic and policy factors also play important roles. For example, Binenbaum et al. (2000) analyse the jurisdictional pattern of intellectual property rights that affect international flows of germplasm and related biotechnologies.

In our own work, still in progress, in which we have used measures of agroecological similarity to parameterise technological spillover potential, we found very substantial spillover effects among US states. An implication is that typical studies that do not allow for interstate or international spillovers will overstate own-state research responsibility for state-level productivity growth, and, thus, state-specific rates of return to research will be overstated. At the same time, the global returns to a state's research will be understated in studies that consider only within-state effects. Some pre-aggregation of state-specific research investments into knowledge stocks is unavoidable in any attempt to capture interstate spillover effects in econometric models. Errors in this pre-aggregation could distort the evidence, just as ignoring spillovers altogether does, even when the pre-aggregation has been done with care and attention to the likely underlying determinants of spillover effects.

Some other choices in an analysis may have important implications for the estimated rate of return, but often we cannot generalise about the size and direction of the resulting biases. For instance, most studies have not attempted to correct for the effects of commodity programs or other distortions, an omission which Alston et al. (1988) showed might lead to over- or under-statement of the benefits and the rate of return; exceptions include Oehmke (1988), Zachariah et al. (1989), and Huang and Sexton (1996). Similarly, selection bias can be a problem — projects may have been selected for analysis because they are known to be winners, without regard for the high proportion of unsuccessful projects, which could be regarded as contributing to overhead costs to be borne by the successful projects. On the other hand, this should not be a problem with studies based on the analysis of aggregate data, and such studies do report lower rates of return (Alston et al., 2000a,b).

3.3. Research and adoption lags: attribution over time

In some respects, investing in research is like investing in physical capital: current productivity does not simply depend on the current rate of investment, but rather on the flow from the stock of usable knowledge derived from the history of past investments. Hence, investment decisions taken in one period have

consequences that last into the future. Indeed, lags and dynamics in agricultural R&D are of greater duration and importance than for most other types of capital investment. There are lags of several years, typically, between when an expenditure is made on research and when the resulting innovation or increment to knowledge begins to be adopted and to affect production and productivity.

The effects of a particular investment today can persist over many future production periods, perhaps forever. The effects of other R&D investments may be short-lived or non-existent. Estimating the parameters that characterise this overall dynamic research–development–adoption–disadoption process is the most challenging empirical problem in evaluating R&D. In the evaluation of individual process innovations (e.g. Griliches, 1957; Schmitz and Seckler, 1970) it is sometimes possible to obtain good information on the timing of events. More often (and inevitably in the case of aggregative analysis across programs and commodities), however, the information is not directly accessible and must be either estimated as a part of the analysis, or imposed on it.

Even the more data-rich studies of aggregate national research systems typically use only 40–50 years of annual observations on research (and, perhaps, extension) expenditures to attempt to explain 20–30 years of variation in production or productivity. Such data are not sufficient to estimate the research lag profile accurately. Indeed, to obtain estimates at all, it has been found necessary to impose a great deal of structure on the lag relationship — including assumptions about its length, smoothness, and general shape. These generally untested (or inadequately tested) restrictions have an impact on the resulting answers. These assumptions are often devised arbitrarily, with a view to convenience of estimation as much as anything, rather than empirically. For example, studies have typically imposed a finite lag structure linking R&D spending to changes in productivity over less than 20 years. But some types of research have effects that persist indefinitely (e.g. we still use electricity), while others have effects that are finite, as the innovation loses effect (e.g. pest resistance is eroded) or is replaced by other innovations and becomes obsolete (e.g. new and better agricultural chemicals), and yet others are very short-lived (e.g. specific computer chips). Hence, a flexible, infinite lag

with some allowance for research obsolescence may be appropriate for econometric work; especially work that aims to estimate the returns to aggregate R&D.

In principle, given sufficient data, a flexible infinite lag model could be implemented using modern time-series econometric approaches. In practice, given data (and other) constraints, an infinite lag structure might be better approximated by the use of a longer finite lag structure than most studies have used (although the potential for bias remains). The few studies that have attempted to estimate lag lengths for aggregate agricultural R&D in the US and the UK econometrically have found that lag lengths of at least 30 years may be necessary (e.g. Pardey and Craig, 1989; Chavas and Cox, 1992; Huffman and Evenson, 1992, 1993; Schimmelpfennig and Thirtle, 1994; Alston et al., 1998). This suggests that the typical study has used a truncated lag structure that is too short.

In a synthetic study, where the research-induced shifts are given, the truncation of the lag amounts to leaving out benefits, which would, *ceteris paribus*, bias the rate of return down. In an econometric study, however, truncation of the lag amounts to omitting relevant explanatory variables. This will lead to biased parameter estimates, with too much econometric weight (yielding larger values for the parameters) on the more recent lags. By itself, the omission of long lags here, as with the synthetic approach, amounts to understating total benefits: but unlike the synthetic studies the present value of the benefits associated with the shorter lags is now greater. In a discounting context, given typically high rates of return, the latter effect is likely to dominate (since the benefits associated with the long-past research expenditures are heavily discounted), so that truncation of the lag will tend to bias rates of return up. This view is supported by the meta-analysis of Alston et al. (2000a,b) and by the econometric analysis of Alston et al. (1998).

Various other specification issues arise in econometric studies of research returns. As well as getting the lags right, in some settings it might be necessary to allow for leads. The typical study assumes unidirectional causation from research to productivity, but Pardey and Craig (1989) provide some evidence, albeit weak, in support of bidirectional causality (see also Schimmelpfennig and Thirtle, 1994). Simultaneous equations bias has an indeterminate effect on single-equation estimates of the impact of research on

productivity. Difficulties in uncovering the lead–lag relationships may be confounded with other problems if these relationships have undergone structural change, as is likely over the relatively long estimation periods that are increasingly being used in conjunction with very long lags.

4. Illustrative examples of attribution problems

To illustrate the nature and the importance of the attribution problems underlying the estimates of rates of return to research, we consider two examples. First is an assessment of the US benefits from wheat variety improvement R&D conducted by the Consultative Group on International Agricultural Research (CGIAR). Second, we consider evidence on the effects of different treatments of the research lag structure on the evaluation of rates of return to agricultural R&D.

4.1. Attribution among investors: US benefits from the CGIAR

Pardey et al. (1996) investigated the impacts in the US of varietal-improvement research performed at the international agricultural research centres funded by the CGIAR. This investigation focused on two cases: the wheat-breeding work carried out at the International Wheat and Maize Improvement Center (CIMMYT) in Mexico (and its antecedent agencies), and the rice-breeding program of the International Rice Research Institute (IRRI) in the Philippines. Both of these programs are very well known: they have been at the centre of efforts to develop the high-yielding grain varieties whose use in developing countries has contributed to large increases in world-wide food supplies — increases commonly referred to as the Green Revolution.

A review of these cases by Pardey et al. (1996) shows that substantial attribution problems can arise even when the details of the technology and the timing of events are well documented and understood. Consider the case of wheat. Pardey et al. (1996) obtained detailed data on experimental yields of the many wheat varieties at multiple locations in each of the different US wheat-growing states. Comparison of experimental-plot yields of new varieties with those in production in 1970 indicates that in the absence

of the new varieties, overall wheat yields would have been 33% lower in 1993. The authors estimated that over 1970–1993, such yield gains generated economic benefits with a present value in 1993 of about 43 bUS\$ (1993 dollars). In other words, approximately 1/9th of the total value of wheat production over the period is attributed to increases in yields resulting from the introduction of new varieties. These are the gross benefits to producers and consumers as a result of the US adoption of the new varieties.

One important aspect of the attribution problem involves determining who deserves the credit for these gains. In particular, what is the fraction of the total benefit that can be attributed to the work done at CIMMYT? Pardey et al. (1996) had complete information on the genetic (and breeding) history of each important wheat variety grown in the US, for each wheat-growing state, along with an extensive dataset on experimental yields by variety for multiple experimental sites (within states). Unfortunately, even such uncommonly detailed information is not enough to solve the attribution problem; genotype does not translate simply into yield gains or other phenotypic characteristics (such as seed size, colour and protein and fibre content) that translate into tangible economic value. How much of the credit for the improvement in US wheat yields associated with semi-dwarfing should go to Norman Borlaug (who led the effort at CIMMYT, and earlier at the Rockefeller Foundation-sponsored research program in Mexico that began in 1943), compared with the breeders at Washington State University (who previously made the first US cross with the *Norin 10* variety from Japan)? How much credit for the excellence of today's variety should go to the breeder who bred it, and how much should go to the breeders and farmers who bred or selected its parents, grandparents and so on? It is not easy to identify the separate marginal product of any particular breeder in the chain. Consequently, economists studying this type of issue have ended up using mechanistic rules to apportion the total benefits across steps in the history of the development of a new variety.

To compute and attribute the benefits from wheat-breeding, Pardey et al. (1996) examined the effects of using a variety of rules to accommodate differing perceptions of the relative importance of earlier and later breeding steps. In general they found that US benefits from the CIMMYT wheat-breeding

program were very large. Even using their most conservative attribution rule (giving the greatest credit to more recent, US-based innovations, and the least credit to the earlier CIMMYT-based innovations), the additional wheat produced in the US as a consequence of the CIMMYT program was worth 3.6 bUS\$ from 1970 to 1993. US government support of the wheat-breeding program at CIMMYT since 1960 was about US\$ 68 million (in present value terms as of the end of 1993). Counting only the benefits from the yield gains in the US, the benefit–cost ratio of US support was greater than 49:1. This is the most conservative estimate. Using alternative attribution rules, the benefit–cost ratio is as high as 199:1.

Recall, this is the benefit from US *adoption* of varieties containing CIMMYT-derived germplasm, which is a gross rather than net measure of the benefits to US from CIMMYT's wheat variety improvement program. It does not account for the costs to the US as a net exporter, which arise when the rest of the world adopts new CIMMYT-based wheat varieties and this leads to a reduction in the demand for and price of US wheat. Evaluating this effect is a much larger undertaking; it involves measuring the effect of CIMMYT's wheat-breeding program on the entire world. This is yet another form of attribution problem, one which generally has not been recognised in previous studies of the country-specific benefits from international agricultural research (one exception is Brennan and Bantilan, 1999).

4.2. Attribution over time: specifying and estimating lag relationships

In empirical work on models of the effects of research on aggregate agricultural productivity, the number of lags and the shape of the lag structure are usually chosen arbitrarily; rarely is either the lag length or structure tested formally. Common types of lag structures include 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, 1989, 1992, 1993). 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 dataset is usually not much longer 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. More recent studies have tended to use more flexible, and longer lags. 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 30 years” (p. 9) and that “long lags — at least 30 years — may be necessary to capture all of the impact of research on agricultural output” (p. 18). Using a non-parametric approach, Chavas and Cox (1992, p. 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 the two just cited, tested how much longer than 30 years the lag length should be.

In contrast, Alston et al. (1998) 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. Alston et al. (1998) laid out a model in which current aggregate production depends on the utilisation of the stock of useful knowledge, which is itself a function of the entire history of relevant investments in R&D. What results is potentially an infinite lag between past investments in research on the one hand and production on the other. While a short, finite lag may reasonably depict the link between investments in research and increments to the stock of useful knowledge, it would be a significant conceptual error to use the same lag to represent the relation between investments in research and production, since production depends on flows from the entire stock of useful knowledge, and not just on the latest increment to this stock. Other recent studies, based on an examination of the time-series nature of the data, rather than reflection about structural relationships, have been tending in a similar direction (e.g. Akgüngör et al., 1996; Makki et al., 1996). Using time-series methods involving data transformations such as first differences, they have found that smaller estimated rates of return result.

Table 2
Lag structure and estimated rates of return to research from econometric models^a

Lag structure	Mean lag (years)	Number of estimates (count)	Rate of return (% per year)				
			Mean	Mode	Median	Minimum	Maximum
Form							
Polynomial	13.2	285	79.9	58.0	58.0	4.5	729.7
Trapezoidal	32.7	55	97.7	95.0	67.0	11.0	384.4
Free-form	28.0 ^b	6	26.5	6.0	30.0	6.0	45.0
Inverted-V	12.0	33	134.5	30.0	72.0	23.0	562.0
Other	13.3 ^b	304	75.6	46.0	48.0	−1.0	1219.0
No structure	26.6	79	45.8	54.0	51.0	0.3	185.0
No lag	0	36	48.0	46.0	44.4	20.9	111.0
All forms	16.3 ^b	762	77.9	58.0	53.0	−1.0	1219.0
Length (years)							
0	0	36	48.0	46.0	44.4	20.9	111.0
1–4	9.9	408	95.2	58.0	60.7	0.0	1219.0
0–15	22.3	174	58.1	46.0	49.9	4.5	260.0
>30	38.0 ^b	144	60.1	40.0	41.6	−1.0	384.4
Unspecified	Unspecified	100	60.0	27.0	41.2	8.9	337.0

^a The figures in this table encompass studies reporting econometrically estimated rates of return to agricultural research only, and to research & extension, reported in Alston et al. (2000a; Table 16).

^b Represents the mean length of the R&D lags for rate of return estimates based on finite lag structures. One of the 6 *free-form* estimates is based on an infinite lag structure, as are 43 of the 304 *other* estimates, 44 of the 762 *all forms* estimates, and 44 of the 144 >30 years estimates.

Table 2 summarises the results from past econometric studies of returns to agricultural research across countries, classified according to the length and form of the research lag. It can be seen that the results are consistent with 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.

To illustrate their ideas and implement the arguments, Alston et al. (1998) developed a model of agricultural productivity that can be used to evaluate typical assumptions about the shape of the research lag, as well as the implicit assumptions about knowledge depreciation associated with explicit assumptions about the research lag length. They applied this model to data on US aggregate agricultural productivity for the period 1949–1991, making use of annual data on total agricultural R&D (including extension) expenditures by the federal government and 48 state governments, for the period 1890–1991. The agricultural input data were adjusted for quality change over time, which will account for certain types of private R&D expenditures and human capital improvements, and so on, but omitted-variables bias may still result from the exclusion of private R&D and spillover effects (the details of the data and estimation procedures can be found in Alston et al. (1998)). The primary conclusion reinforces the view that agricultural research affects productivity for much longer than most previous studies have allowed, possibly forever. A model consistent with infinite lags was statistically preferred over a more conventional model with finite lags. The implications for reported rates of return were quite dramatic. The statistically preferred model indicated a much lower real, marginal IRR to public agricultural research in the US than was implied by a model using shorter lags.

5. Conclusion

Taken at face value, studies of returns to agricultural research indicate that the investment has been enormously socially profitable. Some research investments no doubt have yielded extraordinarily high returns and to some extent the overall picture has been distorted by sampling bias. As well, however, many

of the estimates are likely to have been biased up as a result of attribution problems. We have presented arguments and evidence concerning the nature of these attribution problems and the resulting bias in reported rates of return to research.

For a start, many of the estimates at the upper end of the range are simply implausible. In particular, some very large estimates of IRRs to aggregate R&D investments, if taken literally as IRRs over lengthy time periods, imply unbelievable impacts of agricultural research. Part of the problem here is that the internal rate of return measure may not reasonably represent the relationship between research and returns. In particular, it is questionable whether the primary beneficiaries — the farmers and consumers to whom the benefits accrue — could reinvest research returns at a rate anything like the typical estimated rate of return to research. Hence, even if the rate of return as calculated corresponds arithmetically to the stream of benefits and costs, other summary measures may be more meaningful.

In addition, however, there are issues surrounding the measures of benefits and costs and their interpretation. Significant problems arise in attempting to determine what the pattern of productivity growth would have been in the absence of a particular research investment. Some of these problems concern the use of appropriate index number theory and making appropriate corrections for changes in quality and other characteristics of inputs and outputs. Others arise in defining the relevant counterfactual alternative in order to meaningfully match streams of research benefits and costs.

Further problems can arise in the specification of the research–returns relationship, especially in econometric applications. The typical approaches understate the period over which research affects productivity and, in econometric studies using time-series data, this means they overstate the shorter-term impacts, leading to overstated rates of return. Most studies also fail to fully account for the effects of work done by others in the research–development–extension continuum, and this gives too much credit to the particular investor being evaluated. Work remains to be done to establish the empirical importance of bias due to incomplete correction for the locational spillovers of research results in estimated rates of return to research.

Acknowledgements

Partial support for this work was provided by US-AID, Global Bureau, the Giannini Foundation of Agricultural Economics, and a grant from the UC Pacific Rim research project. The authors thank Albert Acquaye and Connie Chan-Kang for valuable research assistance, and John Brennan, David Colman, Ruben Echeverría, Will Masters and Michel Petit for comments on an earlier version of the work.

References

- Acquaye, A.K.A., Alston, J.M., Pardey P.G., 2000. A disaggregated perspective on post-war productivity growth in US agriculture: isn't that spatial? In: Proceedings of the NC-208 Conference on Agricultural Productivity: Data, Methods, and Measures, Waugh Auditorium, USDA-ERS, Washington, DC, 9–10 March 2000.
- Akgüngör, S., Makanda, D.W., Oehmke, J.F., Myers, R.J., Choe, Y.C., 1996. Dynamic analysis of Kenyan wheat research and rate of return. In: Contributed Paper Proceedings from the Conference on Global Agricultural Science Policy for the 21st Century, Melbourne, Australia, 26–28 August 1996, pp. 1–32.
- Alston, J.M., Pardey, P.G., 1996. Making Science Pay: the Economics of Agricultural R&D Policy. Washington, DC.
- Alston, J.M., Edwards, G.W., Freebairn, J.W., 1988. Market distortions and the benefits from research. *Am. J. Agric. Econ.* 70, 281–288.
- Alston, J.M., Anderson, J.R., Pardey, P.G., 1995. Perceived productivity, foregone future farm fruitfulness, and rural research resource rationalization. In: Peters, G.H., Hedley, D.D. (Eds.), Proceedings of the 22nd International Conference of Agricultural Economists on Agricultural Competitiveness: Market Forces and Policy Choice, Aldershot, Dartmouth, pp. 639–651.
- Alston, J.M., Craig, B.J., Pardey, P.G., 1998. Dynamics in the creation and depreciation of knowledge, and the returns to research. EPTD Discussion Paper no. 35, August 1998, International Food Policy Research Institute, Washington, DC.
- Alston, J.M., Chan-Kang, C., Marra, M.C., Pardey, P.G., Wyatt, T.J., 2000a. A meta-analysis of rates of return to agricultural R&D: ex pede herculem? Research Report no. 113. International Food Policy Research Institute, Washington DC.
- Alston, J.M., Marra, M.C., Pardey, P.G., Wyatt, T.J., 2000b. Research returns redux: a meta-analysis of the returns to agricultural R&D. *Aust. J. Agric. Resour. Econ.* 44 (2), 185–215.
- Ballard, C.L., Fullerton, D., 1992. Distortionary taxes and the provision of public goods. *J. Econ. Perspect.* 6, 117–131.
- Binenbaum, E., Nottenburg, C., Pardey, P.G., Wright, B.D., Zambrano, P., 2000. South-North trade, intellectual property jurisdictions, and freedom to operate in agricultural research on staple crops. EPTD Discussion Paper no. 70, December 2000, International Food Policy Research Institute, Washington, DC.
- Brennan, J.P., Bantilan, M.C.S., 1999. Impact of ICRISAT research on Australian agriculture. Report prepared for Australian Centre for International Agricultural Research, Economic Research Report no. 1. NSW Agriculture, Wagga Wagga.
- Byerlee, D., Traxler, G., 2001. The role of technology spillovers and economies of size in the efficient design of agricultural research systems. In: Alston, J.M., Pardey, P.G., Taylor, M.J. (Eds.), *Agricultural Science Policy: Changing Global Agendas*. Johns Hopkins University Press, Baltimore, MD (Chap. 9).
- Chavas, J.-P., Cox, T.L., 1992. A nonparametric analysis of the effects of research on agricultural productivity. *Am. J. Agric. Econ.* 74, 583–591.
- Craig, B.J., Pardey, P.G., 1996. Productivity in the presence of quality change. *Am. J. Agric. Econ.* 78 (5), 1349–1354.
- Craig, B.J., Pardey, P.G., 2001. Input, output, and productivity developments in US agriculture. In: Alston, J.M., Pardey, P.G., Taylor, M.J. (Eds.), *Agricultural Science Policy: Changing Global Agendas*. Johns Hopkins University Press, Baltimore, MD.
- Davis, J.S., 1980. A note on the use of alternative lag structures for research expenditure in aggregate production function models. *Can. J. Agric. Econ.* 28, 72–76.
- Echeverría, R.G., 1990. Assessing the impact of agricultural research. In: Echeverría, R.G. (Ed.), *Methods for Diagnosing Research System Constraints and Assessing the Impact of Agricultural Research: Assessing the Impact of Agricultural Research*, Vol. II. International Service for National Agricultural Research, The Hague.
- Evenson, R.E., 1967. The contribution of agricultural research to production. *J. Farm Econ.* 49, 1415–1425.
- Evenson, R.E., Waggoner, P.E., Ruttan, V.W., 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? *Am. J. Agric. Econ.* 67, 806–812.
- Fuglie, K., Ballenger, N., Day, K., Klotz, C., Ollinger, M., Reilly, J., Vasavada, U., Yee, J., Fisher, J., Payson, S., 1996. Agricultural research and development: public and private investments under alternative markets and institutions. USDA-ERS-NRED Agricultural Economic Report 735. US Department of Agriculture, Washington, DC.
- Griliches, Z., 1957. Hybrid corn: an exploration in the economics of technological change. *Econometrica* 25 (4), 501–522.
- Griliches, Z., 1963. The sources of measured productivity growth: agriculture, 1940–1960. *J. Political Econ.* 71, 331–346.
- Griliches, Z., 1964. Research expenditures, education, and the aggregate agricultural production function. *Am. Econ. Rev.* 54 (6), 961–974.
- Griliches, Z., 1974. Errors in variables and other unobservables. *Econometrica* 42, 971–978.
- Griliches, Z., 1979. Issues in assessing the contribution of research and development to productivity growth. *Bell J. Econ.* 10, 92–116.
- Griliches, Z., 1992. The search for R&D spillovers. *Scand. J. Econ.* 94, 29–47.
- Griliches, Z., 2001. R&D and productivity: the unfinished business. In: Alston, J.M., Pardey, P.G., Taylor, M.J. (Eds.), *Agricultural*

- Science Policy: Changing Global Agendas. Johns Hopkins University Press, Baltimore, MD.
- Huang, S.-Y., Sexton, R.J., 1996. Measuring returns to an innovation in an imperfectly competitive market: application to mechanical harvesting of processing tomatoes in Taiwan. *Am. J. Agric. Econ.* 78, 558–571.
- Huffman, W.E., Evenson, R.E., 1989. Supply and demand functions for multiproduct US cash grain farms: biases caused by research and other policies. *Am. J. Agric. Econ.* 71, 761–773.
- Huffman, W.E., Evenson, R.E., 1992. Contributions of public and private science and technology to US agricultural productivity. *Am. J. Agric. Econ.* 74, 752–756.
- Huffman, W.E., Evenson, R.E., 1993. *Science for Agriculture: A Long-term Perspective*. Iowa State University Press, Ames.
- Jaffe, A.B., 1986. Technological opportunity and spillovers of R&D: evidence from firm's patents, profits, and market value. *Am. Econ. Rev.* 76 (5), 984–1001.
- Johnson, D.K.N., Evenson, R.E., 1999. R&D spillovers to agriculture: measurement and application. *Contemp. Econ. Policy* 14 (4), 432–456.
- Kealey, T., 1996. *The Economic Laws of Scientific Research*. St. Martin's Press, New York, 382 pp.
- Leiby, J.D., Adams, G.D., 1991. The returns to agricultural research in Maine: the case of a small northeastern experiment station. *Northeastern J. Agric. Resour. Econ.* 20, 1–14.
- Makki, S.S., Tweeten, L.G., Thraen, C.S., 1996. Returns to agricultural research: are we assessing right? In: *Contributed Paper Proceedings from the Conference on Global Agricultural Science Policy for the 21st Century*, Melbourne, Australia, 2–28 August 1996, pp. 89–114.
- Mansfield, E., 1977. *The Production and Application of New Industrial Technology*, Norton, New York.
- Oehmke, J.F., 1988. The calculation of returns to research in distorted markets. *Agric. Econ.* 2, 291–302.
- Pardey, P.G., Craig, B., 1989. Causal relationships between public sector agricultural research expenditures and output. *Am. J. Agric. Econ.* 71, 9–19.
- Pardey, P.G., Alston, J.M., Christian, J.E., Fan, S., 1996. *Hidden harvest: US benefits from international research aid*, IFPRI Food Policy Report, International Food Policy Research Institute, Washington D.C.
- Pasour Jr., E.C., Johnson, M.A., 1982. Bureaucratic productivity: the case of agricultural research revisited. *Public Choice* 39 (2), 301–317.
- Perrin, R., Fulginiti, L., 1996. Productivity in the presence of “poorly priced” goods. *Am. J. Agric. Econ.* 78 (5), 1355–1359.
- Ravenscraft, D., Scherer, F.M., 1982. The lag structure of returns to research and development. *Appl. Econ.* 14, 603–620.
- Schimmelpfennig, D., Thirtle, C., 1994. Cointegration, and causality: exploring the relationship between agricultural R&D and productivity. *J. Agric. Econ.* 45, 220–231.
- Schmitz, A., Seckler, D., 1970. Mechanized agriculture and social welfare: the case of the tomato harvester. *Am. J. Agric. Econ.* 52, 569–577.
- Schultz, T.W., 1956. Reflections on agricultural production, output and supply. *J. Farm Econ.* 38, 748–762.
- Thirtle, C.G., Bottomley, P., 1988. Is publicly funded agricultural research excessive? *J. Agric. Econ.* 31, 99–111.
- Zachariah, O.E.R., Fox, G., Brinkman, G.L., 1989. Product market distortions and the returns to broiler chicken research in Canada. *J. Agric. Econ.* 40, 41–51.