Has the Return on Australian Public Investment in Agricultural Research Declined?

Eran Binenbaum*, John Mullen**, Chang Tao Wang***

* Corresponding author; School of Economics, University of Adelaide, eran.binenbaum@adelaide.edu.au

** New South Wales Department of Primary Industries

*** School of Economics, University of Adelaide

Paper to be presented at the 2008 AARES conference in Canberra

Abstract: We examine whether there has been a decline in the returns from Australian public investment in research on broadacre agriculture. Complementing a forthcoming paper by Mullen, we use alternative specifications for the regression equation, which employs the log of total factor productivity (TFP) as the dependent variable. The rate of return is computed on an annual basis rather than by using multi-year averages. In contrast to Mullen's earlier preliminary analysis, we have now found some evidence of a decline in the rate of return on public R&D investment, lending some support to recently voiced concerns on this matter.

Keywords: R&D, rate of return, agricultural productivity, public investment

1. Introduction

The purpose of this paper is to examine whether there has been a decline in the returns from Australian public investment in research on broadacre agriculture. Mullen (forthcoming) discusses concerns that such a decline may have occurred (p. 2) and provides a preliminary answer ("No") based on regression analysis (section 4). The present paper presents more extensive analyses of the same data set.

Our starting point is Table 1, which reproduces Mullen's (forthcoming) Table 5. Section 2 discusses some of the merits and demerits of the linear and quadratic versions of this model. Section 3 introduces a novel method for computing annual rates of return on public R&D investments, and applies this method to the linear and quadratic model versions. Section 4 discusses three kinds of alternative regression models: fractional exponents, a time interaction term, and linear spline models. The last type of model is found to be superior to all others, and the annual return method is applied to it. Section 5 concludes.

2. The linear and quadratic models

Mullen and Cox (1995), Mullen (forthcoming), and Wang (2006), regressed total factor productivity (TFP) against a knowledge stock variable, a weather index, farmers' terms of trade and farmers' education, where all variables were in logs and the models were linear. Mullen (forthcoming) additionally included a regression with a quadratic logged knowledge stock term. The results are reproduced in Table 1.

Table 1: Econometric results and IRRs from the 35 and 16 year lag models (identical to Table 5 in Mullen forthcoming)

35 year models										
Period	**1953-1988**		**1953-2003**		**1953-2003**		**1969-2003**		**1969-2003**	
	Coeff.	t-stat	Coeff.	t-stat	Coeff.	t-stat	Coeff.	t-stat	Coeff.	t-stat
Knowledge stock:										
Linear	0.16	3.44	0.14	2.58	-1.9	-5.53	-3.05	-1.68	0.25	3.93
quadratic					0.08	5.96	0.13	1.82		
Weather	0.04	5.19	0.3	5.48	0.24	5.67	0.3	3.76	0.31	3.84
Education	2.22	3.05	2.35	2.72	3.7	5.36	3.33	2.87	4.42	4.31
Terms of Trade	-0.27	-2.52	-0.49	-3.58	-0.29	-2.71	-0.27	-2.09	-0.30	-2.25
R^2		0.95		0.95		0.97		0.94		0.93
D-W		2.02		1.13		1.96		1.92		1.74
Reset		na		38.1		3.72		1.93		5.24
IRR%		17		10		15		16		13
16 year models										
Period	**1968-1988**		**1953-2003**		**1953-2003**		**1969-2003**		**1969-2003**	
	Coeff.	t-stat	Coeff.	t-stat	Coeff.	t-stat	Coeff.	t-stat	Coeff.	t-stat
Knowledge stock:										
Linear	0.22	2.22	0.001	0.25	-4.39	-5.45	-13.8	-3.11	0.22	2.49
quadratic					0.17	5.46	0.52	3.16		
Weather	0.26	3.22	0.3	5.01	0.24	5.06	0.3	3.75	0.3	3.24
Education	2.11	0.98	2.94	3.24	6.08	6.65	6.89	6.71	6.01	5.34
Terms of Trade	-0.27	-1.92	-0.8	-7.24	-0.43	-3.88	-0.3	-2.30	-0.49	-3.72
R^2		0.83		0.95		0.97		0.94		0.92
D-W		1.73		1.23		1.69		1.86		1.53
Reset		na		32.4		11.2		5.72		10.4
IRR%		30		0		23		39		21

In these regressions, the knowledge stock variable is constructed as a weighted average of past R&D investments, based on 35-year and 16-year trapezoidal lag profiles respectively (see Mullen and Cox 1995, Mullen forthcoming, and Wang 2006 for details).[1] As to the relative performance of these models, Mullen (forthcoming) notes: "The RESET test provides some guidance as to whether quadratic or interaction terms are missing from the model. Adding a quadratic knowledge stock term led to a marked improvement in the properties of both models[2] as can be seen in Table [1]… The econometric properties of both the 16 and 35 year quadratic models are strong. All coefficients are precisely estimated and have the expected sign (expectations about the signs on the knowledge stock variables are discussed further below). For the 35 year model, the D-W and RESET statistics and the plot of the CUSUM values all suggest few problems with the specification of this model. These same specification statistics for the 16 year model suggested that specification remains a problem. Non-nested testing of these two models provided clear evidence in favour of the 35 year model and supported concerns about the specification of the 16 year model…" Additionally, the 16-year model featured a decreasing knowledge stock term for some years as opposed to a uniformly increasing knowledge stock in the 35-year model (Figure 5 in Mullen forthcoming). The more natural assumption, of course, would be to expect the term to be uniformly increasing. Furthermore, distinguished experts such as Huffman and Evenson (2006) and Alston (personal communication) favour the longer lag specification over the shorter one. All these considerations point towards the superiority of the 35-year model. We will thus in this paper focus on it and ignore the 16-year model.

One potential point of concern relates to the possibility of spurious regression due to some of the variables exhibiting unit roots (integrated of order one). A complicating factor here is that the other variables do not exhibit unit roots (integrated of order zero). As not all variables are integrated of the same order, our set of variables cannot be said to be cointegrated, and some of the cointegration tests – such as the Johansen (1988) algorithm – that might allow one to proceed with the regression analysis, are unavailable. Fortunately, the Engle-Granger (1987) test is still available. Whilst this test in the absence of unit roots in some of the variables no longer indicates cointegration, it can still be used to allay concerns of spurious regression (Hamilton

[1] To obtain a full knowledge stock time series beginning in 1953, a backcasting procedure is used to fill in the missing past investment figures.
[2] As suggested by Garry Griffith.

1994). In this two-step procedure, a regression is run on a vector of time series variables, generating a time series of the residuals. This latter time series is then subjected to a standard Dickey-Fuller unit root test. Engle-Granger tests on the linear and quadratic models indicated an absence of a unit root in the residuals of each regression, thus validating our usage of these regressions.

We now mention the issue of greatest concern with the quadratic model. It is actually not straightforward whether the quadratic 35-year model is preferable to the linear one, for the quadratic model does suffer from a major flaw: for the first seven years, the marginal effect of central importance, namely β_t, defined as the elasticity of TFP_t with respect to K_t,

$$\beta_t \equiv \frac{\partial \ln TFP_t}{\partial \ln K_t} = \beta_L + 2\beta_Q \ln K_t \ ,$$

(where β_L and β_Q denote the linear and quadratic regression coefficients respectively), is negative: see Table 2.

Table 2: Marginal effect of ln K_t on ln TFP_t in quadratic regression model

t	Beta(t)	t	beta(t)	t	beta(t)	t	beta(t)
1953	-0.088	1966	0.092	1979	0.268	1992	0.384
1954	-0.074	1967	0.107	1980	0.280	1993	0.388
1955	-0.059	1968	0.121	1981	0.292	1994	0.391
1956	-0.045	1969	0.135	1982	0.303	1995	0.394
1957	-0.031	1970	0.149	1983	0.314	1996	0.397
1958	-0.016	1971	0.163	1984	0.324	1997	0.399
1959	-0.003	1972	0.176	1985	0.334	1998	0.401
1960	0.011	1973	0.190	1986	0.343	1999	0.403
1961	0.024	1974	0.204	1987	0.351	2000	0.404
1962	0.037	1975	0.218	1988	0.359	2001	0.406
1963	0.051	1976	0.231	1989	0.366	2002	0.406
1964	0.065	1977	0.244	1990	0.373	2003	0.406
1965	0.078	1978	0.256	1991	0.384		

Clearly, an increase in the knowledge stock ought to have a positive effect on productivity. The fact that this marginal effect is negative for the early years is also suggestive of a downward bias in the (positive) estimates of this effect for subsequent years, at least well into the 1960s. This suggests model misspecification: imposition of a quadratic curvature yields a downward-sloping fitted curve for (early and therefore) low values of K_t.

3. A procedure for estimating annual rates of return

Mullen (forthcoming), following Mullen and Cox (1995), derives the rate of return on public investment in R&D through a three-step procedure which is fairly standard in the literature (see references cited in these two papers):

1. Construct the knowledge stock variable as a weighted sum of past R&D investments, e.g. (the formula used in Mullen, Mullen and Cox, and in the present paper) $\ln K_t = \sum_{j=1}^{L_R} r_j \ln R_{t-j}$, where L_R is the maximum lag length (in our case, 35), r_j are a set of weights that sum to one and comprise the lag profile (in our case, a trapezoidal one), and R_t are annual research investments.
2. Regress ln TFP_t on $\ln K_t$ and other explanatory variables, which in the simplest (linear) case yields a constant coefficient β but which more generally yields a time-variant marginal effect β_t.
3. Compute the internal rate of return (IRR) i by solving the equation $$1000 \equiv TVMP = \sum_{j=0}^{L_R} \frac{\beta r_j \mathrm{GM}(TFP)\mathrm{GM}(P)}{\mathrm{GM}(R)(1+i)^j},$$ where TMVP (set equal to 1000 due to units used on the right-hand side) stands for total value of marginal product, GM( ) stands for "the geometric mean of…" and P is a price index.

The logic by which the formula of step 3 is derived is explained in Mullen and Cox (1995). A critical part of this logic is that the purpose of the exercise is to obtain a unique rate-of-return estimate, and hence the time-variant terms TFP_t, P_t and R_t are replaced by their geometric means.

One problem with this approach is that it tends to overestimate *R* relative to *TFP* – step 1 indicates that *R* has a distributed-lag effect on *TFP*, and as both *R* and *TFP* tend to grow over time, the relevant value of *R* used here ought to be lower than its geometric mean. This bias can be corrected by using the geometric mean of an earlier time series of *R*.

A more fundamental problem, and one central to the present paper, is that *i* itself is not time-invariant and has inherently no unique solution. Note also that the problem is compounded for regression models that are not strictly linear, as some averaging procedure has to be devised for β as well. It is more meaningful to compute an IRR for each year of investment. The IRR is then based on a stream of future benefits for the investment in that year, using the relevant values of β_t, TFP_t, P_t and R_t which are weighted by the weights r_j:

$$1000 \equiv TMVP_t = \sum_{j=0}^{L_R} \frac{\beta_{t+j} r_j TFP_{t+j} P_{t+j}}{R_t \left(1+i_t\right)^j}$$

This procedure is much more consistent with steps 1 and 2. If the analyst is interested in a unique IRR value characterising the entire data set, it can be obtained by first computing annual IRRs and then averaging these geometrically. The time series of i_t generated by our revised time-variant IRR formula serves our present primary interest, which is to extract evidence on a possible change in IRR. Figure 1 and Table 4 show the annual IRR time series thus obtained.

Figure 1: IRR time series for linear, quadratic, and spline regression models

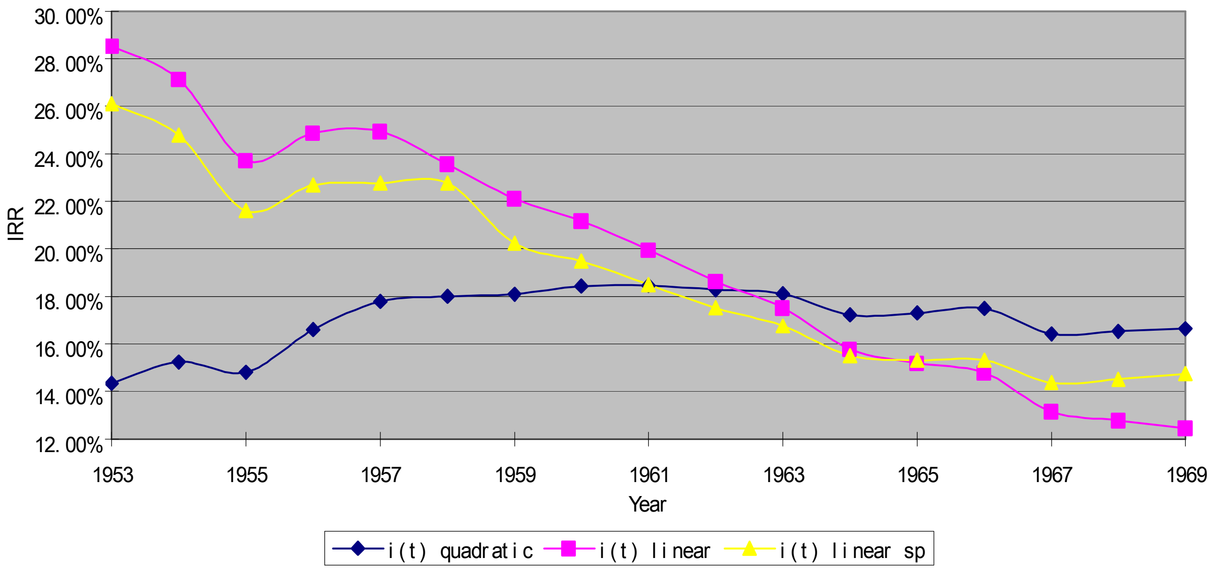

This procedure relates R&D investment in any given year to the actual values of future *TFP* and *P* properly weighted in accordance with the lag profile used. Variations based on future *K* are also affecting the result via the time-variant β_t term. For example, the investment in 1953 contributes – assuming that the lag profile used is correct – via the future knowledge stock to future *TFP* through 1987. Thus, the last investment year for which full information on the future flow of benefits is available is 1969. We thus obtain time series of 17 years of IRR for each model, which is not quite long enough for formal hypothesis testing but which does allow the drawing of preliminary conclusions.

Figure 1 shows a striking contrast between the linear and quadratic regression models. Whilst the linear model indicates a fairly dramatic decline in the IRR from over 28% for 1953 investments to less than 13% for 1969 investments, the IRR time series based on the quadratic model is generally flat, never falling below 14% and never exceeding 19%. However, we know that the linear and quadratic models are both flawed. The linear model completely misses the increasing slope that is highly significant in the quadratic model – it is constrained to produce a single β which is likely an overestimate (underestimate) of β_t for the early (later) years. The quadratic model, on the other hand, evidently underestimates β_t for the earlier years, as discussed in the previous section. Let us, then, look for alternative model specifications.

4. Alternative regression models

In this section we consider, in turn, fractional exponents, a time interaction term, and linear spline models.

The fact that the quadratic model outperforms the linear one in most respects may be taken to indicate that there is a significant nonlinear component in the model. The first alternative to look at naturally would be the polynomial approach, i.e., having found the quadratic term highly significant when that is the highest order, to also try including a cubic term in addition to the lower-order terms, etc., and to stop when the highest-order term is insignificant. However, inclusion of a cubic ln K term in the regression yields insignificant quadratic and cubic terms. This result implies that the polynomial approach stops at the quadratic level.
Another approach is to include a single nonlinear term along the linear one. Why should this nonlinear component be quadratic? One way to examine this is to allow for a capital stock term with variable exponent E. There is no a priori reason why, in a regression with a nonlinear term with multiplicative parameter β_N,

$$\ln TFP = \alpha + \beta_L \ln K + \beta_N \left(\ln K\right)^E + \text{..........},$$

E should equal 2. We attempted to estimate this regression using STATA's 'nonlinear[3] least squares' (Davidson and McKinnon 1993) routine, but this turned out to be too sensitive to the specification of the initial parameter value for E. As an alternative, we repeatedly estimated ordinary least squares increasing the exponent parameter in the above equation by one decimal point at a time, that is, using E = 1.1, 1.2, 1.3, …… The purpose of this exercise was to try and remedy the problem of negative and low positive values of the marginal effect β_t for the earlier years in our data set in the quadratic OLS regression. That regression was, of course, replicated by setting E = 2. Estimates for β_t (for t = 1953, 1954, …, 1963) as a function of E

[3] The usage of the word *nonlinear* in the term "nonlinear least squares" refers to the regression being nonlinear in the parameter E and should not be confused with the usage of the word in the rest of this paper.

are shown in Figure 2. Values of E are depicted on the horizontal axis at β_t = 0. Estimates of β_t for E = 2 correspond to those in Table 2.

Figure 2: Marginal effect as a function of exponent in nonlinear regression term

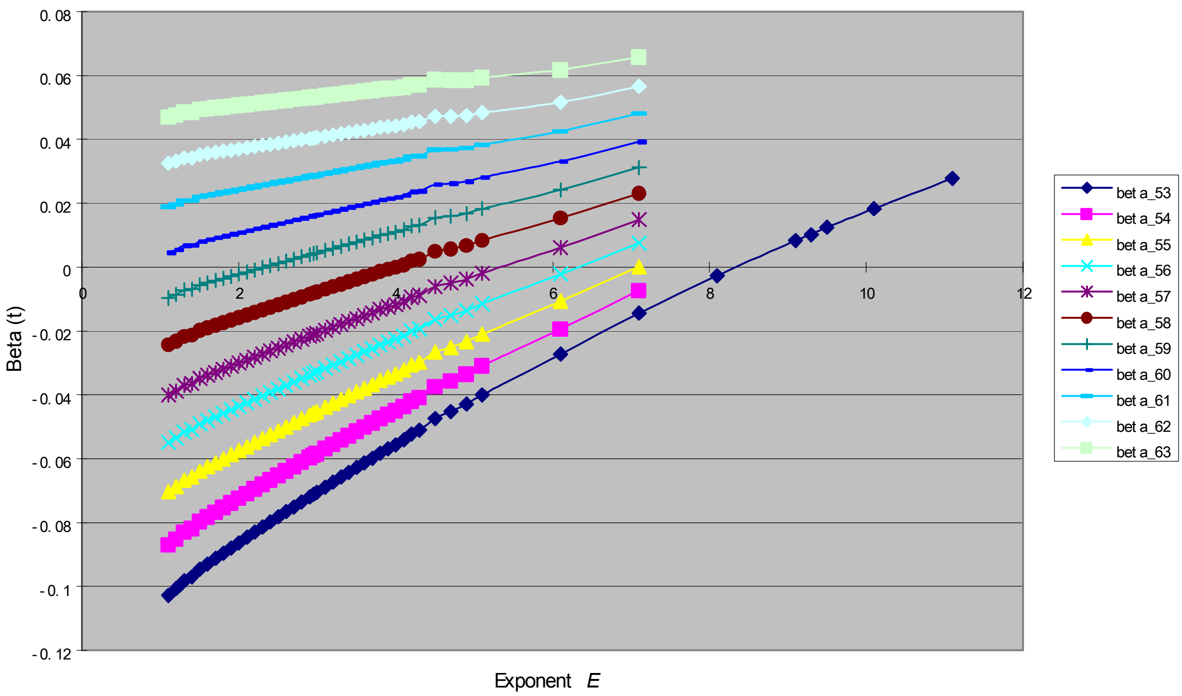

These results show that the issue of a negative or small positive marginal effect for the earlier years in our data set cannot be resolved by substituting an exponent value different from 2. Even for $E = 7$, there are still a few negative marginal effect statistics left, whilst the marginal effects for the first eleven years (all the ones depicted in the graph) are still arguably too low. Moreover, for the higher values of E shown in the graph, β_L becomes insignificant. In summary, the variable-exponent approach is not satisfactory as an alternative to the linear and quadratic regression models.

Next, consider the possibility of a time interaction term. This is based on a simple idea: in contrast to the linear model, we would like the marginal effect of ln K on ln *TFP* to be time-variant. Then why not have it directly be a function of time? This translates into the equation

$$\ln TFP = \alpha + \left(\beta_L + \beta_{\text{int}} t\right) \ln K + \text{ ....}$$

which implies that β_{int}, rather than a separate time trend coefficient (tried unsuccessfully before – see Mullen forthcoming and Wang 2006), is a time interaction coefficient. This variation of the regression model yields a significant β_{int} and an insignificant but negative β_L. Using the 'best estimate' value for β_L instead of zero, which in spite of its insignificance is the statistically preferable approach, we obtained values for the marginal effect,

$$\beta_t = \beta_L + \beta_{\text{int}} t ,$$

that were negative for the first 23 years. This clearly indicates that the time interaction approach has to be rejected too.

Finally, we considered a linear spline approach (Greene 2007). This approach allows for structural change to occur in the course of the data set. It includes multiple linear segments ('bins') with different slopes, separated by kink points ('knots'). The regression is conducted under the constraint that at the knots the left-hand and right-hand limits of the fitted line are equal, i.e. that the function be continuous. We attempted a three-bin spline regression with approximately 30 possible combinations of the two knots, but found that consistently the slope parameter in the first bin was insignificant. Furthermore, using three bins for the small number of observations may be asking too much of the data. We thus settled on two bins instead. We considered as candidate years for the knot the years. Results are reported in Table 3.

Table 3: Results of two-bin spline regressions

Knot	beta1	beta2	R-sq	RESET	AIC	BIC
75	0.0847954	0.3491876	0.972	6.77	-124.3266	-112.7356
76	0.0924833	0.3564951	0.9718	7.02	-124.0772	-112.4862
77	0.1007951	0.3704909	0.9721	6.82	-124.4601	-112.8692
78	0.1107759	0.3868867	0.9718	7.17	-124.0525	-112.4616
79	0.1222749	0.4050378	0.9711	8.23	-122.7413	-111.1503
80	0.1331629	0.4314032	0.9709	8.7	-122.3016	-110.7107
81	0.1415096	0.4683473	0.9717	7.35	-123.7859	-112.1949
82	0.1477874	0.5077437	0.9724	6.21	-125.0773	-113.4863
83	0.1533036	0.5539901	0.9731	5.12	-126.343	-114.7521
84	0.1576375	0.5931646	0.9726	5.57	-125.3746	-113.7837
85	0.1619398	0.643485	0.972	6.09	-124.4154	-112.8244
86	0.1670265	0.7148216	0.972	5.97	-124.4353	-112.8443

The RESET results were satisfactory for all the knots considered, but these are essentially t-statistics that are not appropriate for knot selection. Instead, we consider

R-squared, the Akaike Information Criterion (AIC), and the Bayesian Information Criterion (BIC) (Greene 2007) for knot selection. These consistently point to 1983 as the knot with the best performing spline. Moreover, 1983 is the middle year of the three consecutive years 1982-1984 that comprise the top-three results in terms of each of these criteria. Thus we selected 1983 as the preferred knot in the linear spline approach. This yielded a time series of IRRs for the years 1953-1969 that is displayed in Figure 1. Values are provided in Table 4.

Table 4: IRR time series and means for linear, quadratic and linear spline regressions

Year	Quadratic	Linear	Linear spline
1953	14.33%	28.54%	26.11%
1954	15.24%	27.12%	24.78%
1955	14.81%	23.72%	21.60%
1956	16.60%	24.87%	22.68%
1957	17.79%	24.94%	22.76%
1958	17.99%	23.57%	22.76%
1959	18.10%	22.11%	20.23%
1960	18.43%	21.17%	19.47%
1961	18.45%	19.94%	18.50%
1962	18.28%	18.63%	17.51%
1963	18.10%	17.52%	16.75%
1964	17.23%	15.77%	15.50%
1965	17.31%	15.18%	15.30%
1966	17.48%	14.78%	15.31%
1967	16.42%	13.16%	14.36%
1968	16.52%	12.79%	14.50%
1969	16.63%	12.45%	14.74%
GEO MEAN	16.99%	19.11%	18.63%

The linear spline model avoids the linear model's upward bias as well as the quadratic model's downward bias in estimating the marginal effect β_t for the earlier years. For the last several years in Figure 1 and Table 4, the linear and quadratic models' biases are likely reversed, and the spline model may be avoiding these as well. These features are reflected in Figure 1, with the spline-based IRR values being below the linear-based ones and above the quadratic-based ones, with these inequalities being reversed for the last several years for which IRR estimates are available. The spline-based IRR series does show a general decline in IRR, which however is less pronounced than for the linear-based IRR series. Furthermore, there is a subtle but potentially important difference between the spline-based series and the linear series for the last three years 1967-1969, where the former exhibits a levelling off or even a slight recovery in IRRs whereas the latter exhibits a continued decline.

5. Conclusions

Among the models examined here, all exhibit serious pitfalls except the linear spline model, which is therefore the most likely to produce reliable results. Using the linear spline model, there is some evidence of a decline in the rate of return on public investment in R&D for investments made in the years from 1953 trough 1969; however, there appears to be levelling off or a slight recovery for investments made in the last three years of this period. As we are using a 35-year model with observations through 2003, 1969 is the last year for which full information about benefit streams is available. Given the small number of observations and some strong assumptions made, particularly with regard to the construction of the knowledge capital stock, these results should be treated with caution.

References

Davidson, R. and McKinnon, J.G. (1993). *Estimation and Inference in Econometrics.* New York: Oxford University Press.

Engle, R.F. and Granger, C.W.J. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica,* 55, 251-276.

Hamilton, J.D. (1994). *Time Series Analysis.* Princeton: Princeton University Press.

Huffman, W.E. and Evenson, R.E. (2006). Do formula or competitive grant funds have greater impacts on state agricultural productivity? *American Journal of Agricultural Economics,* 88, 783-798.

Greene, W.H, (2007). *Econometric Analysis,* 6th edition. Upper Saddle River, NJ, USA: Prentice Hall.

Johansen, S. (1988). Statistical analysis of cointegration factors. *Journal of Economic Dynamics and Control,* 12, 231-254.

Mullen, J.D. (forthcoming). Productivity growth and the returns from public investment in R&D in Australian broadacre agriculture. *Australian Journal of Agricultural and Resource Economics.*

Mullen, J.D. and Cox, T.L. (1995). The returns from research in Australian broadacre agriculture. *Australian Journal of Agricultural Economics,* 39**,** 105-128.

Wang, Chang Tao, (2006). The returns from research in Australian broadacre agriculture revisited. Unpublished Honours thesis, School of Economics, University of Adelaide.