The Australian Journal of Agricultural and Resource Economics, 43:3, pp. 323–335

The optimal management of research portfolios

Peter Bardsley*

Risky research projects are, other things being equal, intrinsically harder to monitor than projects that are less risky. It is shown using agency theory that a standard cost benefit analysis, which ignores the agency problem, will introduce a bias towards excessively risky projects, and it will under-estimate the benefits from complementary investments in libraries, scientific equipment and other expenditures that increase the productivity of scientists. Research managers should be risk-averse in their choice of projects, and they should aim to hold a balanced portfolio of projects. The nature of this portfolio problem is, however, quite different from the portfolio management problem in financial markets.

1. Introduction

How should a public research institute choose its portfolio of projects? Should it diversify its investment in some risk-averse manner, or should it evaluate the expected value of each potential project in isolation,[1] accepting those for which the expected benefits exceed the expected costs and rejecting the others?

The focus of this article is on applied scientific research in agriculture carried out in large publicly funded institutions, though the conclusions are not restricted to this sector. There is a long tradition of publicly provided research in agriculture because of the structure of agriculture and the nature of the technology. Some research outputs can be captured in a proprietary way (chemicals, seeds) but many cannot. Alston, Pardey and Smith (1996) estimate that worldwide public research expenditure in agriculture exceeds US$9 billion per year, and that it has been increasing at 5 per cent per year. This research is undertaken in a number of institutional settings: government departments of agriculture, universities, national scientific institutes; and international scientific institutes, in particular those of the Consultative Group on International Agricultural Research (CGIAR). Support for

* Peter Bardsley, Department of Economics, the University of Melbourne, Melbourne.

[1] There are straightforward cost and benefit based reasons for not considering projects in isolation if there are common resources or project interdependencies at the research or implementation stage. Interdependencies of this type will not be considered here.

research in agriculture has been seen as an important form of international development aid.

There is a large body of economic analysis to back up this public investment in scientific research. For comprehensive reviews of this literature see Norton and Davis (1981) and Alston, Pardey and Norton (1995). The overwhelming bulk of this literature is devoted to determining how to evaluate correctly the costs and benefits of research, taking into account supply shifts, spill-overs, adoption lags and research risks. In response to this literature, and also in response to ever-greater demands for public accountability, project evaluation has become a major preoccupation of many public research institutions. However, much of the existing literature focuses on the risk-neutral evaluation of isolated research projects. While this is important to research managers, it reflects only one of their problems. They are concerned with managing their research program as a whole. It will be argued here, using some fairly simple ideas from the economic theory of agency, that it is not sufficient to manage research on a simple project-by-project basis. Research managers should be risk-averse in their choice of project, and they should aim for a balanced project portfolio.

The Arrow–Lind theorem (Arrow and Lind 1970) suggests that if projects are small relative to the size of the economy, and if the returns are uncorrelated with aggregate income fluctuations, then they should be evaluated in a risk-neutral manner. Most agricultural research projects would fall within this classification, since a great deal of project risk is research risk, and there seems to be no reason why this would be correlated with income. If this argument is accepted, then project evaluators should take risk into account only to the extent of calculating the expected value of potential projects. There should be no bias either towards risky or towards safer projects, and there should be no attempt to construct a 'balanced' portfolio of projects. Good projects should be fully funded up to the point of diminishing marginal expected returns, even if they take up most of the research budget.

In practice, public research portfolios do not display either the concentration of effort or the focus on high-value projects that one might expect to see if one accepts the Arrow–Lind argument. There may be several reasons for this. Research organisations often serve an interest group base, which is diversified by region and industry. One explanation for portfolio diversity is the need to have 'something to show' to each segment of this diverse interest base. Another reason is risk avoidance by research managers. In part this may be due to managerial risk aversion. It may also be due to the judgment that it is not appropriate, from an organisational point of view, to invest in excessively risky research or to put all of one's eggs in one basket.

A notable example of this point of view is the study by Scobie and Jacobsen (1992) and Scobie, Jacobsen and Frances (1992) of the research portfolio funded by the Australian wool industry. These authors believe that there is a strong parallel between investment in research and other forms of investment. They draw on the analogy with financial investment, arguing that risk is relevant and that the same risk management principles should apply. In particular, they pick up from finance the standard tools of mean-variance analysis from the Capital Asset Pricing Model (CAPM). An immediate implication of this approach is that risk must be evaluated and managed on a portfolio-wide basis. Project-by-project evaluation is no longer good enough. It is worthwhile to be clear why this analogy with the CAPM is not appropriate for evaluating research portfolios. At first sight there are important similarities between financial investment and investing in scientific research. In both cases a risky decision must be made which affects risk-averse individuals. In both cases there is an agency relationship, where the main stake holder (the principal) stands at arm's length from the investment. There is, however, an important difference. In standard finance theory the principal (the investor) is risk-averse, while the agent (the stock broker or portfolio manager) is risk-neutral.[2] In scientific investment, the opposite is true and the agency problem is turned on its head. The principal is the corporate research manager, standing in place of the firms in the client industry, while the agent is the research scientist.[3] According to the Arrow–Lind theorem, the principal is risk-neutral. The agent, managing a single project or a small group of projects, is risk-averse. Thus the agency problem is quite different in the two cases, and one cannot rely without further argument on the analogy with finance.

So why do research managers avoid risk and diversify their portfolios? It may be that they do not understand the force of the Arrow–Lind argument, and they would change their practice if it were explained to them. It seems more likely, however, that risk avoidance in project selection is a rational and perhaps socially beneficial reaction to the agency problem that they face. It will be argued here that this is indeed so, and that proper consideration of the agency relationship will lead to a portfolio management problem which is somewhat different from that of standard finance theory.

[2] The principal-agent problem is in this case degenerate, and it is usually suppressed from the analysis. Recent work in finance focuses more on this agency relationship.

[3] One can of course consider a hierarchical chain of risk averse managers. For clarity it will be assumed throughout this article that head office is run by a sophisticated public policy analyst, who understands the Arrow–Lind theorem and has the objective of maximising expected benefits.

2. Agency theory

It is hard to monitor what scientists actually do, especially if they are located in field stations or laboratories that are physically and culturally distant from head office, and if the work that they do is inscrutable to managers. It is difficult to know whether they are working hard, whether they are following their own scientific interests rather than pursuing corporate priorities, whether they are directing their efforts into projects with positive economic benefits, and whether they are facilitating the practical application of their ideas.

In the language of agency theory, there is a 'hidden action' or 'moral hazard' problem. This problem can be managed by monitoring and supervision, and by the creation of appropriate incentive structures, but this is only an imperfect remedy. A strong incentive structure may do an excellent job of aligning the scientist's values with those of the organization, but the role of incentives is constrained by the risk that they impose on the agent.[4] Research is a risky activity, and individual scientists have only a limited capacity to absorb these risks. Risk aversion thus puts a limit on the extent to which the organisation can influence and control the work that scientists do.

There is also a 'hidden knowledge' or 'adverse selection' problem. Scientists have private knowledge about their abilities, about the probability of success of their proposals, and about the actual progress that has been made in their ongoing but incomplete projects. For natural reasons, they may not be totally frank in revealing this information to head office project evaluators and managers. Even if they wish to do so, it may be difficult to communicate this information credibly. A scientist may truly know that a proposal has a very high expected return, but if everybody exaggerates the merits of their proposals then it may be difficult to get this accurate piece of information through to management. People who manage or evaluate research confront problems of this type continually. This article will focus on the 'hidden action' problem and its connection with project risk; 'hidden knowledge' problems will at this stage be left to one side.

The projects that are the hardest to monitor are precisely those high-risk, high-return projects that, if they were measured by the yardstick of standard cost-benefit analysis, may be the most socially valuable. The reason why risky projects are hard to monitor is straightforward. If a project has only a 10 per cent chance of success, then failure of the project provides almost no information to the monitor. It is after all what is to be expected. In this

[4] Holmstrom and Milgrom (1991, 1994) point out that in the presence of measurement bias a strong incentive scheme may have perverse effects on values.

situation, it is not possible to 'pay by results' (for example, by rewarding successful researchers with promotions or research grants) because the good researchers are unlikely to differentiate themselves from the bad in a reasonable time. On the other hand, if the success probability is 90 per cent, then a failure is quite informative. For these reasons, one might expect to see a bias towards safer projects. Another way of looking at this problem is that if agents are risk-averse (which is certainly true of the agents in this problem), then any attempt to manage by incentives will impose high costs on risk-averse agents who take on risky projects.

The implication is that, if project risk is ignored, then the project portfolio may not be managed in an optimal way. This idea is explored formally in section 3 of the article, where the main ideas are developed in the context of a single project. In this section the main instruments available to the manager are the type of project that is undertaken, and complementary investments in scientific infrastructure that may affect the probability of success and the risk of failure in the project. In section 4 a project portfolio is studied. The introduction of several projects introduces a new management instrument, namely project diversification, which is addressed in this section. The analysis here is focused on a two-project portfolio, which is sufficient to explore the relevant issues (an extension to many projects, which requires new mathematical techniques, is discussed in Bardsley (1997)). Policy implications are discussed in section 5.

3. Managing a single project

Consider first the situation where the risk-neutral research manager (the principal), asks the scientist (the agent) to work on a single project. This can be modelled as a standard principal-agent model (Holmstrom 1979; Grossman and Hart 1983).

The project may succeed, with probability p, yielding a benefit b to the principal, or it may fail resulting in a benefit of zero. The agent chooses the probability p, which requires effort $e(p)$. The agent is risk averse, with a utility function that is separable in money and effort. If the project succeeds then the principal pays the agent $X = x(U)$, yielding monetary utility U. If it fails, then the principal pays the agent $x = x(u)$, yielding monetary utility u. The agent will participate provided that the expected utility is non-negative. It will be assumed that the agent is risk averse (hence $x'(v) > 0$, $x''(v) > 0$) and that $e'(0) = 0$, $e'(p) \geq 0$, $e''(p) > 0$, and $e'(p) \to \infty$ as $p \to 1$. In broad terms these assumptions about the effort function mean that the agent can increase the probability of success by working harder, but that no matter how hard he or she works, success will never be certain.

The participation and incentive compatibility constraints are as follows.

$$pU + (1-p)u = e(p) \tag{1}$$

$$U - u = e'(p) \tag{2}$$

The principal's objective is to maximise

$$p(b - x(U)) - (1-p)x(u)$$

subject to these constraints. The agent's problem is convex provided that $e''(p) > 0$, and the first-order approach is sufficient to solve the agency problem. The Lagrangean is

$$\mathrm{L} = p(b - x(U)) - (1-p)x(u) + \lambda(pU + (1-p)u - e(p)) + \mu(U - u - e'(p)).$$

This leads to the following first order conditions in addition to equations 1 and 2.

$$x'(U) = \lambda + \frac{\mu}{p} \tag{3}$$

$$x'(u) = \lambda - \frac{\mu}{(1-p)} \tag{4}$$

$$b = (x(U) - x(u)) + \mu e''(p) \tag{5}$$

After a slight rearrangement this gives

$$e(p) = u + pe'(p) \tag{6}$$

$$\lambda = px'(U) + (1-p)x'(u) \tag{7}$$

It can be shown (see Appendix) that $\mu > 0$, $\lambda > 0$, and $0 < p < 1$.

In contrast, if the principal can observe and control the agent's effort then the incentive compatibility constraint does not bind; the Lagrangean is

$$\mathrm{L}' = p'(b - x(U)) - (1-p')x(u') + \lambda'(p'U' + (1-p')u' - e(p')),$$

where the prime indicates values that would be chosen in a full information environment. The first order conditions are

$$u' = U' = e(p')$$

$$x'(u') = x'(U') = \lambda'$$

$$b = e'(p').$$

Figure 1 shows the effort function $e(p)$, the reward structure (u, U), and the probability of success p. The chord between u and U is tangent to the effort function; this geometrical fact summarises the agent's participation and incentive compatibility constraints. In this figure there is also shown the contract (u', U'), and the success probability p' that would be chosen if effort were observable and there were no agency problem. It can be shown (see

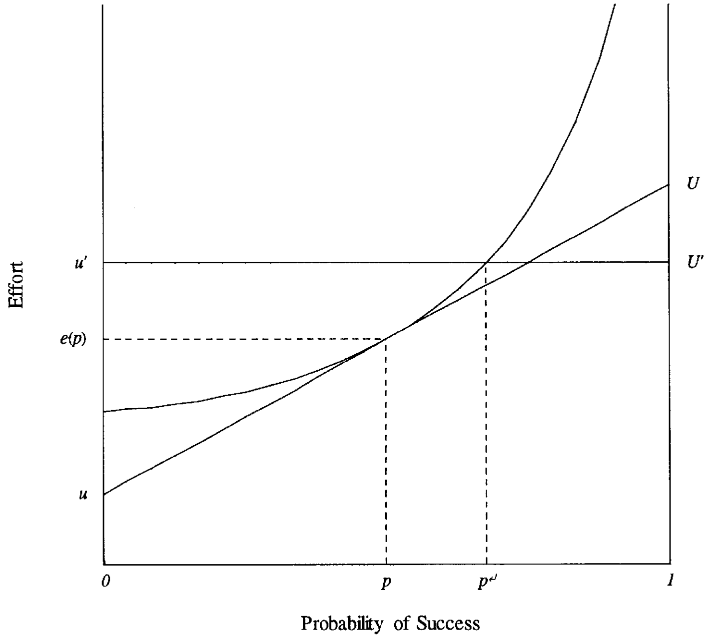

Figure 1 Probability of success

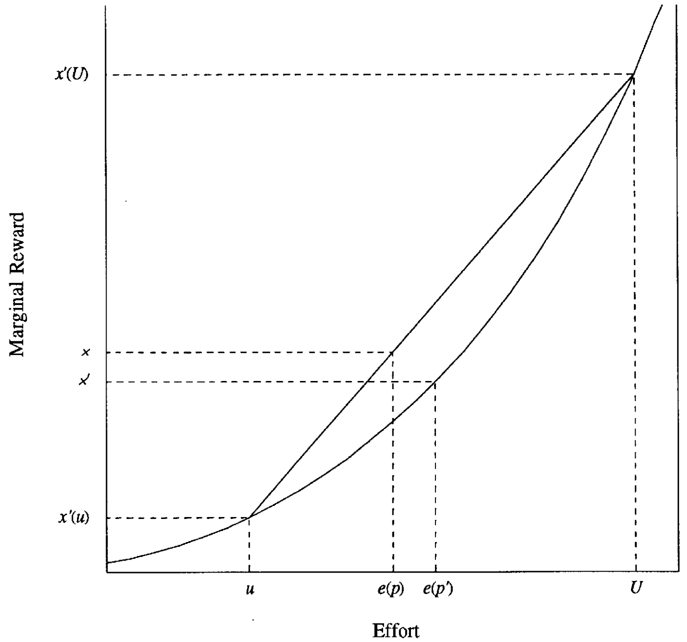

Figure 2 Effort and marginal reward

Appendix) that $p' > p$; the optimal effort is less, as is the probability of success, when effort cannot be monitored.

Figure 2 shows the relationship (equations 6 and 7) between the optimal effort $e(p)$ and the Lagrange multiplier λ. The Lagrange multiplier is an indication of how strongly the agent's participation constraint binds. If λ is large, then the agent's willingness to participate is an important constraint on the principal's ability to meet his objective. It can be seen that λ may be either greater or smaller than λ'.

3.1 Management implications

Now consider the choice of project. To do this, assume that the effort required to achieve a success rate p depends on the choice of project. Thus the effort function $e(p)$ should be replaced in the model above by $e(p, b)$. It will be assumed that $e_b(p, b) > 0$ and $e_{pb}(p, b) > 0$. More valuable projects require greater effort both in total and at the margin. The optimal value of b can be found by applying the envelope theorem to the Lagrangean:

$$p = \lambda e_b(p, b) + \mu e_{pb}(p, b). \tag{8}$$

The left-hand side of this equation represents the marginal expected benefit from moving to a more valuable project. The first term on the right represents the marginal increase in the cost of the project, assuming that there is no change in the agent's behaviour. This is precisely the calculation that would be made in a standard cost-benefit analysis, such as might be undertaken by a head office team of project evaluators, using standard techniques and ignoring the agency problem (the Lagrange multiplier λ converts marginal utility to marginal cost). The third term represents the effect on the agent's incentives. The change in the project increases the marginal effort required to achieve the success rate p. This deterioration in the agent's incentive must be corrected by a steeper reward schedule; since the agent is risk averse this increases the expected cost to the principal. This is the third term of the equation. Since μ is positive, so is this last term. This leads to the first conclusion from this analysis. A standard cost-benefit analysis, which ignores the agency problem, will choose a project that is too risky.

One can also consider the value of investing in complementary capital items or other expenses: scientific equipment, libraries, conference expenses, and so on. Let x be the complementary expenditure. The effort function $e(p)$ should be now be replaced by $e(p, x)$. It will be assumed that $e_x(p, x) > 0$ and $e_{px}(p, x) > 0$. The ancillary investment raises both total and marginal productivity. Of course x must also be subtracted from the Lagrangean as an up-front expense. Using the envelope theorem, the optimal level of complementary investment is given by

$$1 = \lambda e_x(p, x) + \mu e_{px}(p, x). \tag{9}$$

The conclusion is as before. A standard cost-benefit analysis considers only the direct reduction in the cost of the research effort, and ignores the effect on incentives. The complementary resources increase the agent's marginal productivity, leading to a higher level of effort for a given incentive structure. Thus the standard analysis will result in an under-investment in related resources that increase the productivity of scientists.

4. Managing a project portfolio

Now consider the case where there are several projects. Bergemann (1992) has analysed a general multi-task principal-agent problem that can be applied in this context. Assume for simplicity that there are two projects and that the projects are independent. The benefits and the costs are additive, and the outcomes are statistically independent. Project i may succeed with probability p_i, yielding a benefit b_i to the principal, or it may fail resulting in a benefit of zero. To achieve this probability of success the agent must exert an effort costing in utility terms $e_i(p_i)$. The principal offers a contract which yields the agent utility u^{11} if both projects succeed, u^{10} if only project one succeeds, u^{01} if only project two succeeds, and u^{00} if both projects fail. This requires a monetary transfer of $x(u^{ij})$, where the reward function $x(u)$ is the inverse of the utility function. The contract must satisfy participation and incentive compatibility constraints. The principal's problem can be summarised by the Lagrangean

$$\begin{aligned}
\mathrm{L} = {} & p_1 p_2(b_1 + b_2 - x(u^{11})) + p_1(1 - p_2)(b_1 - x(u^{10})) \\
& + (1 - p_1)p_2(b_2 - x(u^{01})) + (1 - p_1)(1 - p_2)(-x(u^{00})) \\
& + \lambda(p_1 p_2 u^{11} + p_1(1 - p_2)u^{10} + (1 - p_1)p_2 u^{01} \\
& \quad + (1 - p_1)(1 - p_2)u^{00} - e_1(p_1) - e_2(p_2)) \\
& + \mu_1(p_2(u^{11} - u^{01}) + (1 - p_2)(u^{10} - u^{00}) - e_1'(p_1)) \\
& + \mu_2(p_1(u^{11} - u^{10}) + (1 - p_1)(u^{01} - u^{00}) - e_2'(p_2)).
\end{aligned}$$

The key first-order conditions are

$$x'(u^{11}) = \lambda + \frac{\mu_1}{p_1} + \frac{\mu_2}{p_2}, \tag{10}$$

$$x'(u^{01}) = \lambda - \frac{\mu_1}{(1-p_1)} + \frac{\mu_2}{p_2}, \tag{11}$$

$$x'(u^{10}) = \lambda + \frac{\mu_1}{p_1} - \frac{\mu_2}{(1-p_2)}, \tag{12}$$

$$x'(u^{00}) = \lambda - \frac{\mu_1}{(1-p_1)} - \frac{\mu_2}{(1-p_2)}. \tag{13}$$

Just as in the single project case, one can show that $\lambda > 0$, $\mu_1 > 0$ and $\mu_2 > 0$.[5] The main implications of this model then follow: there will be a bias towards projects that are too risky, and there will be under-investment in complementary resources.

What does this mean for research management? The implication is that the research portfolio should be *well diversified*. This is clearly contrary to the full information policy under which only the project, or group of projects, promising the maximal expected return would be funded.

Bergemann explains clearly the reasoning for this type of result. The principal wants to impose sharp incentives on the agent, in order to get a high level of effort. However, strong incentives impose an unacceptable degree of risk on the agent. Anything that reduces the agent's exposure to risk is good for the principal, as it allows a more favourable incentive structure to be imposed. One way to do this is to allow the agent to hold a diversified project portfolio. The striking implication is that diversification is required at the branches of the organisation, not at the roots. Each laboratory manager or field station manager should hold a diversified portfolio, but there is no reason to hold a diversified portfolio of laboratories or field stations.

5. Conclusion

Public research is undertaken by large, complex and decentralised organisations. While the managers of these research institutions pay attention to the cost-benefit calculations produced by their project analysts, this is only a small part of their task. Much of their work is concerned with the management of people and of groups of people. This article suggests that this concern is intrinsic to the nature of the task, and that the proper economic analysis of research projects should take into account structure and incentives in the organisations that carry out research. Several aspects of the

[5] See Bardsley (1997).

agency relationship warrant investigation; this article focuses on the 'moral hazard' or 'hidden action' problem.

The main conclusions are as follows. Despite the Arrow–Lind theorem, risk does matter. Research managers should be risk-averse in choosing projects. A risk-neutral expected return calculation ignores the effect of risk on the incentives of scientists and project managers. Not all risk is relevant. Research managers should be risk-averse with regard to research risk (the risk that projects might fail), but risk-neutral with regard to market risk or the risk that discoveries might not be adopted (for these the Arrow–Lind theorem applies). A risk-neutral expected return calculation will under-estimate the benefit from complementary expenditure (on libraries, scientific equipment, or travel to conferences) that raises the marginal productivity of scientists. Such a calculation ignores the beneficial effect of complementary investment on incentives. Research institutions should hold a diversified portfolio of projects, but the diversification should occur at the local or branch level of the organisation. The reasons for diversification are quite different from those in financial markets.

In a recent report, the World Bank (1996) highlights some of the policy issues in agricultural research that are currently facing the world agricultural system. It recommends extensive adoption of the formal research evaluation paradigm exemplified by the Alston, Norton and Pardey (1995) monograph. It recommends, in particular, that expected project return should not be discounted for risk, and it argues this case on the basis of the Arrow–Lind proposition. The optimal portfolio of projects is just the portfolio of highest-ranking projects, each fully funded[6] in descending order of rank until the budget runs out. One implication of this position is that research portfolios would be likely to be very concentrated, with the lion's share of the funds going to one or two favoured projects. Another implication is that, if project evaluation can be made routine, with common approaches and standardised software, then funds can be allocated by a decentralised bidding process. There is no need for any management overview of the portfolio composition. Australia is a good example of a country where this approach is being applied more frequently. Funds under the control of industry research corporations are allocated in precisely this way, and government departments are adopting a similar purchaser-provider model.

I would argue here to moderate this rather simplistic approach. There is no doubt that the evaluation driven model of research fund allocation is an important part of the answer to the search for efficiency. But research

[6] That is, until decreasing returns reduce the value of the project to equal the value of the next best project.

managers do have a legitimate concern about the structure of their research program as a whole, and this concern goes beyond merely finding the projects with the highest expected rate of return. They are also concerned with evaluating scientists, and providing scientists with incentives to work well. A diversified portfolio with a variety of projects, some easy, some hard, providing a range of milestones and signals of success, makes this easier to do. When these incentive effects are taken properly into account, as well as the factors emphasised by Alston, Norton and Pardey, one might expect to see a slightly different project portfolio. One might also expect to see research management institutions that are responsive to, but are not completely driven by, project evaluation techniques.

Appendix

The basic model is a particular case of the general agency model analysed by Grossman and Hart (1983), in particular, Section 4 of that paper. Most of the results used here can be found in some form amongst their results, but is easier to derive them from scratch. The important results are

$$\lambda > 0, \tag{i}$$

$$\mu > 0, \tag{ii}$$

$$0 < p < 1, \tag{iii}$$

$$p < p'. \tag{iv}$$

Consider these in turn. The first follows immediately from equation 7, since the marginal payment function is positive.

The second is a version of Holmstrom's Theorem. Assume that $\mu \leq 0$; by equations 3 and 4, and the monotonicity of the reward function, $x(U) \geq x(u)$. By equation 5 this implies that $b \leq 0$, which is a contradiction.

To show (iii), assume that $p = 0$, By equation (3) $\mu = 0$, so $u = U$. Then by equation 5 this implies that $b = 0$, which is a contradiction. Thus $p > 0$. It is clear that $p < 1$ from the assumption that the effort function becomes infinitely steep before p reaches 1.

Now consider (v). For a given p, let $(u(p), U(p))$ be the contract which induces effort p. Notice that $\dfrac{du(p)}{dp} = -pe''(p)$ and $\dfrac{dU(p)}{dp} = (1-p)e''(p)$. The principal's problem is to choose p to maximise

$$V(p) = pb - px(U(p)) - (1-p)x(u(p)).$$

Without loss of generality, utility can be scaled so that $u(p') = 0$ and $U(p') = 1$ (recall that p' is the success rate chosen under full information).

Let $x(v, \alpha) = \alpha x(v) + (1-\alpha)(x(0) + x'(0)v) = x(0) + x'(0)v + \alpha\theta(v)$. As α varies between 0 and 1, this perturbs the inverse utility or reward function of the agent; if $\alpha = 0$ then the agent is risk neutral, while if $\alpha = 1$ we return to the original preferences. Note that $\theta'(v) \geq 0$ and $\theta''(v) \geq 0$ for $0 \leq v \leq 1$ (this is a consequence of risk aversion). Let $p(\alpha)$ be the effort that the principal would induce from an agent of type α; this maximises the objective

$$V(p, \alpha) = pb - px(U(p), \alpha) - (1 - p)x(u(p), \alpha).$$

Notice that $p(0) = p'$ (since a risk-neutral agent will implement the full information effort p), while $p(1) = p$. To show that $p < p'$ it is sufficient to show that $p'(v) \leq 0$ for $0 \leq v \leq 1$. By a standard envelope argument, the sign of $p'(v)$ is the same as the sign of the cross-derivative

$$V_{p\alpha} = (\theta(u(p(\alpha))) - \theta(U(p(\alpha)))) + p(1 - p)e''(p)(\theta(u'(p(\alpha))) - \theta(U'(p(\alpha)))).$$

By (ii) above $U > u$, so this quantity is negative as required.

References

Alston, J.M., Norton, G.W. and Pardey, P.G. 1995, *Science under Scarcity: Principles and Practice for Agricultural Research Evaluation and Priority Setting*, Cornell University Press, Ithaca, New York.

Anderson, J.R., Pardey, P.D. and Roseboom, J. 1994, 'Sustaining growth in agriculture: a quantitative review of agricultural research investments', *Agricultural Economics*, vol. 10, no. 2, pp. 107–23.

Arrow, K.J. and Lind, R.C. 1970, 'Uncertainty and the evaluation of public investment decisions', *American Economic Review,* vol. 60, no. 3, pp. 364–78.

Bardsley, P. 1997, *Multiple Action Agency: An Application to the Management of Scientific Research*, the University of Melbourne Department of Economics Research Paper No. 550, Melbourne.

Bergemann, D. 1992, 'Multiple tasks in the principal agent model: a generalized portfolio problem', *Caress Working Paper* 92–33.

Grossman, S.J. and Hart, O.D. 1983, 'An analysis of the principal-agent problem', *Econometrica,* vol. 51, no. 1, pp. 7–45.

Holmstrom, B. 1979, 'Moral hazard and observability', *Bell Journal of Economics,* vol. 10, no. 1, pp. 7–45.

Holstrom, B. and Milgrom, P. 1991, 'Multi-task principal-agent analyses: incentive contracts, asset ownership, and job design', *Journal of Law, Economics and Organization,* vol. 7, special issue, pp. 24–52.

Holmstrom, B. and Milgrom, P. 1994, 'The firm as an incentive system', *Yale Working Paper series D*, no. 64.

Norton, G.W. and Davis, J.S. 1981, 'Evaluating returns to agricultural research: a review', *American Journal of Agricultural Economics*, vol. 63, no. 4, pp. 685–99.

Scobie, G.M. and Jacobsen, V. 1992, *Allocation of R&D funds in the Australian Wool Industry*, Department of Economics, University of Waikato, report prepared for the Australian Wool Research and Development Corporation.

Scobie, G.M., Jacobsen, V. and Frances, J. 1992, 'Investing in R&D in the wool industry: a portfolio approach', contributed paper, 36th Annual Conference of the Australian Agricultural Economics Society, Canberra.

World Bank 1996, *Achievements and Problems in Development of National Agricultural Research Systems*, OED, World Bank Report No. 15828, World Bank, Washington, DC.