Scottish Journal of Political Economy, Vol. 55, No. 5, November 2008

ASSESSING THE ASSESSMENT OR, THE RAE AND THE OPTIMAL ORGANIZATION OF UNIVERSITY RESEARCH

*Manfredi M. A. La Manna**

Abstract

The UK Research Assessment Exercise (RAE) is assessed as an incentive scheme affecting the allocation of research talent of varying 'quality' across departments. The 'centres of excellence' policy implicitly pursued through the RAE is an optimal allocation strategy only if all departments in all disciplines are of the generalist variety, i.e. each pursues a research path through all its stages. Conversely, the RAE-induced research allocation minimizes efficiency if applied to specialist departments, when resources are concentrated on one specific research obstacle. It is argued that the RAE should not take the organization of University research as exogenous, but rather should encourage specialization. All results are obtained by applying to University research concepts and solutions borrowed from the mathematical theory of systems reliability.

I Introduction

What is going to have a direct cost of £10m with an additional indirect cost of $35m, and to be used to allocate nearly £8000m of research funding in the United Kingdom in 2009? The answer, of course, is the 2008 Research Assessment Exercise (RAE), the latest round of the appraisal of University research regarded by some as 'the most sophisticated exercise of its type in the world'. In 2001,[1] 69 panels of experts assessed the 'research performance' of 2560 units of assessment (departments) from 190 British higher education institutions, involving, among other things, the evaluation of 203,743 publications written by 50,672 active researchers.

The estimate of man- (and woman-) hours expended for the preparation and assessment of thousands of departmental submissions – involving highly distinguished and senior researchers – is difficult to calculate. The Higher Education Funding Council for England estimates that RAE 2008 will have a

*University of St Andrews

[1] The figures for the 1996 RAE were: 60 panels, 2896 units of assessment, 192 HE institutions, 212,553 publications, and 55,893 researchers.

direct cost of £10m, with the indirect cost sustained by universities bringing the likely cost to £45m (source: HEFCE, 2004). These costs pale into insignificance when compared with the £8000m of research funding that are going to be distributed across the British HE system on the basis of the outcome of RAE 2008. Entire departments have been closed as a direct result of poor rankings in the 1996 and 2001 RAEs. Moreover, the RAE is a critical factor in the hiring decisions (and their timing) by Universities, especially in the period preceding the RAE deadline for submissions (see Hayri, 1997). The quality of the debate on the merits and demerits of the RAE has been quite disappointing, with rivers of ink and endless Common-Room conversations devoted to the minutiae of the assessment procedure. Academic post-mortems on the RAE have tended to concentrate on the number and type of relevant 'research outputs', the revealed-preference rankings of various academic journals, etc.,[2] rather than the fundamental purpose of the exercise.

This paper adopts an altogether different view in so far as it side-steps the question of how to measure 'research performance' and focuses directly on one of the main problems that the RAE is (implicitly) assumed to be addressing, namely the allocation of research resources of varying quality to a given number of research units. Obviously, no single paper can analyse satisfactorily the complicated interplay between the process whereby new ideas are produced ('research') and the monitoring and funding of research specifically undertaken within University departments. In order to make some progress in this under-researched area, the present paper makes some drastic simplifications concerning both the research process and the management of University research. Specifically, out of the many characteristic features of research, the paper focuses on just two (albeit fundamental) aspects of research, namely:

(i) the successful completion of almost any research project involves the overcoming of more than one intellectual obstacle (the *multi-task* nature of research) and
(ii) typically there is more than one way to overcome any given research obstacle (the *multi-path* nature of research).

As far as the organization and management of University research is concerned, the paper simplifies dramatically by assuming that

(iii) each discipline pursues only one (multi-task) programme;[3]
(iv) attention can be restricted to research carried out by individuals within the same department/unit;[4]

[2] After each RAE, most disciplines undertake some post-mortem analysis of the outcome. For some examples from economics and management studies, see Ball and Butler (2004), Bessant *et al.* (2003), Ball (1997), Taylor and Izadi (1996), Johnes (1990), Johnes and Johnes (1993), and Doyle *et al.* (1996).

[3] The main results of the paper still hold if multiple programmes are allowed, and indeed would be strengthened in the presence of any inter-programme inter-task correlation.

[4] In this paper I restrict the analysis to the case where the units of assessments are individual departments and not research groupings across departments (both domestic and foreign). The reason for this restriction is twofold: firstly, I wish to analyse the RAE scheme as is currently

(v) the productivity of individual researchers is independent of the productivity of others;[5] and
(vi) interactions between research and teaching can be ignored.[6]

The justification for such obviously unrealistic assumptions is that the resulting model, apart from being manageable, yields efficiency gains of significant magnitude, thereby suggesting that its qualitative results may still apply in more realistic settings.

In the simplified scenario summarized in assumptions (iii)–(vi), the paper poses two key questions:

(a) If the organization of research within Universities is taken as exogenously given, is the incentive scheme implicit in the RAE optimal, i.e. does it promote the optimal allocation of talent within disciplines?
(b) Alternatively, if the organization itself of University research can be affected by a suitable incentive mechanism, how should the RAE be modified to organize research optimally?

The answer to the former question turns out to be ambiguous in the sense that the incentive scheme implicit in the RAE, by promoting 'centres of excellence,' is indeed optimal if and only if for *every* discipline *each* department is of the 'generalist' type, i.e., each department tackles *every* task in the given research programme. If, on the other hand, at least for some disciplines, departments specialize in tackling specific tasks (i.e., are of the 'specialist' type), then the distribution of talent induced by the RAE *minimizes* efficiency.

As to the latter question, the answer is unambiguous: whenever possible, research ought to be carried out by specialist departments.

The paper is organized as follows: in Section II, a simple example is provided to show that special care ought to be taken in allocating funds to multi-task multi-path research. Section III contrasts two polar models of research organization whereby the most productive researchers are either concentrated in elite departments ('centres of excellence') or dispersed across all departments ('diffusion of talent'). Section IV provides a brief sketch of the relevant incentive effects of the RAE on the distribution of research talent across departments. In Section V, it is argued that if the RAE incentive mechanism is changed to favour specialist departments, then substantial efficiency gains can be achieved. Section VI concludes.

organized (which, of course, is based on departments being the relevant units) and, secondly, it can be argued that in an exercise ultimately aimed at allocating resources to domestic departments, it is reasonable to take the non-domestic 'components' of the research system as exogenous.

[5] This assumption is relaxed in 'Extension to correlated productivities'.

[6] Two recent models where research and teaching interact can be found in De Fraja and Iossa (2002) and Gautier and Wauthy (2007). The current paper addresses the more focused issue of whether the RAE provides the best incentives for the optimal assignment of resources to research and in this it follows the RAE in ignoring any complementaries/synergies between teaching and research.

II A Simple Example of Multi-Task, Multi-Path Research

The simplest model of multi-task research is clearly one in which a research project yields a gross benefit π iff each of *two* tasks is completed successfully. Thus, if we let p_i and w_i be, respectively, the probability of success in, and the cost of performing, task i ($i = 1, 2$), the net expected return of the project is given by $p_1 p_2 \pi - w_1 - w_2$. Suppose next that there are two independent units, A and B, which can carry out this project deploying different, independent, and uncorrelated research strategies. For simplicity, let $p_1^A = p_2^A \equiv p^A > p^B \equiv p_1^B = p_2^B$ and $w_1^A = w_2^A \equiv w^A < w^B \equiv w_1^B = w_2^B$, i.e., unit A is both more productive (in terms of probability of success) and less costly than unit B.

For illustration, consider the case where $p^A = \frac{1}{3}$; $p^B = \frac{1}{6}$ and $w^A = £1\text{m}$; $w^B = £0.5\text{m}$; and $\pi = £18\text{m}$. Any funding council adopting the non-negative net expected return criterion for approving research proposals would be indifferent between funding unit A or not and would definitely reject the funding proposal from unit B, as their expected net returns are, respectively, 0 and $-£0.5\text{m}$. However, if the two units can share information about success in each task, the cost–benefit analysis is altogether different, as now the expected net return for this multi-task multi-path project is given by

$$\begin{aligned}\left[1 - \left(1 - p^A\right)\left(1 - p^B\right)\right]^2 \pi - 2w^A - 2w^B &= \frac{16}{81} £18\text{m} - £2\text{m} - £1\text{m} \\ &= £0.56.\end{aligned}$$

This simple example shows that in the case of multi-task, multi-path research, unless research strategies are perfectly correlated, then assessing projects/departments *individually* (rather than as a *system*) is likely to be significantly sub-optimal.

III Two Polar Models of Research Organization

If one takes departments as the units to be deployed to tackle the basic research multi-task programme of any given scientific discipline, then it may be useful to consider two polar cases of research organization:

(a) The '*specialist*' department: all researchers are assigned at tackling the *same* (broadly defined) task, exploring possible alternative solutions.
(b) The '*generalist*' department: researchers are assigned in such a way that *all* tasks required to complete a given research path successfully are tackled.

To illustrate, consider the simplest case of multi-task research, where success in the research programme requires completing two tasks, t_1 and t_2, and suppose that each task can be tackled by two uncorrelated strategies: one succeeding with probability p_H and the other with probability p_L (with $p_H > p_L$). To make the task allocation problem significant, we need a minimum of four researchers to be allocated between two departments. Thus, assuming that two researchers are of

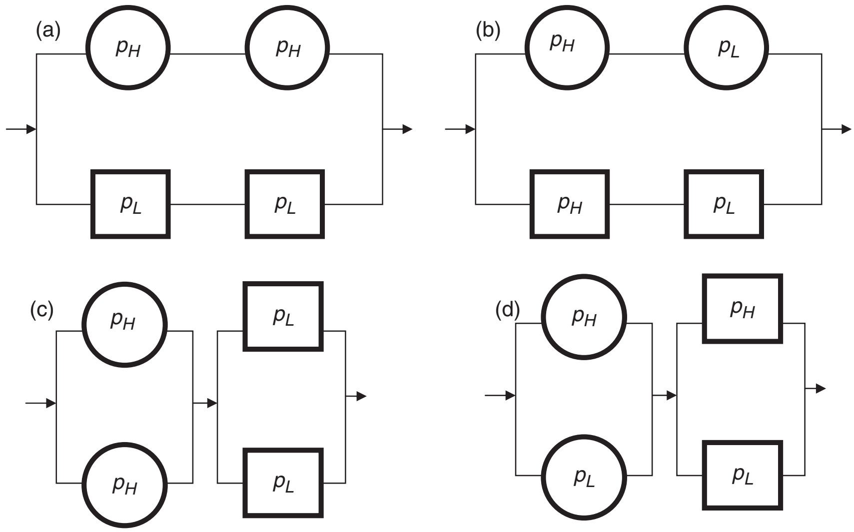

Figure 1. (a) Two generalist departments with talent concentration $\pi_1 = p_H^2 + p_L^2 - p_H^2 p_L^2$. (b) Two generalist departments with talent dispersion $\pi_2 = 2p_H p_L - p_H^2 p_L^2$. (c) Two specialist departments with talent concentration $\pi_3 = (2p_H - p_H^2)(2p_L - p_L^2)$. (d) Two specialist departments with talent dispersion $\pi_4 = (p_L + p_H - p_H p_L)^2$.

high productivity (i.e., succeed at their task with probability p_H) and the remaining two are of lower productivity (i.e., succeed at their task with probability p_L), there are four possible research structures, described in Figure 1(a)–(d), where the 'round' researchers belong to one department and the 'square' researchers to the other.

Even though this is obviously a very stylized example of research organization, some observations of a more general nature can be made:

1. The first obvious point to note is that, depending on the choice between generalist vs. specialist departments, *combined with* the choice between concentration vs. dispersion of talent, each of the four possible research structures succeeds in completing the overall research programme with a different probability denoted by π_i ($i = 1, \ldots, 4$) in Figure 1. Therefore, an effective incentive scheme ought to have at least two dimensions, affecting *both* the basic organization of research by departments (specialist/generalist) and *the* allocation of talent.
2. Indeed, even in the extreme case in which researchers have the *same* productivity and thus the issue of the distribution of talent does not arise, the choice between specialist or generalist departments does matter as it has a significant impact on the overall probability of success (see examples in Table 6 below).
3. The effect of concentrating talent on the overall efficiency of the research process is vastly different depending on the organization of research, whether by specialist or generalist departments.

Table 1
Funding of Research in England and Scotland, 2002–2003

	HEFCE		SHEFC	
RAE Rating	RAE1996	RAE2001	RAE1996	RAE2001
1	0	0	0	0
2	0	0	0	0
3b	1.00	0	1.00	0
3a (unchanged)	1.50	0.31	1.55	0
3a (improved)	1.50	0.31	1.55	1.00
4	2.25	1.00	2.40	1.55
5	3.375	1.89	3.72	2.80
5*	4.05	2.71	3.72	3.20

IV Centres of Excellence vs. Diffusion of Talent

Some basic features of the RAE in England and Scotland

The available literature produced by the UK Funding Councils is silent on the objective(s) that the whole exercise is supposed to achieve, apart from the obvious aim of ranking departments according to a common set of criteria.[7] One indisputable fact is that the revealed preference of the RAE is to foster a 'centres of excellence' strategy, whereby the most efficient resources are concentrated in a few 'star' departments. This is shown by the criterion for the allocation of research funding as a function of the departmental ranking.

Table 1 (reproduced from Hare, 2002) shows the 'funding ratios' in 1996 and 2001, i.e. the amount of research funding allocated to each member of staff, whereby, for example in 1996 a researcher in a 5*-rated department would receive 4.05 times the amount given to someone in a 3b-rated department.

It may be helpful to recall that for any given discipline (or Unit of Assessment), the amount of Research Funding RF_z accruing to the set of all z-rated departments (where $z = 1, 2, \ldots, 5^*$) is given by the formula:

$$RF_z = A \times \sigma_z \times w_z, \tag{1}$$

where A is the total amount of Research Funding for the given discipline, σ_z is the share of active researchers in all z-rated departments as a percentage of total number of active researchers and w_z is the weight assigned to active researchers in z-rated departments. To appreciate the significant change in the funding allocation formula between 1996 and 2001, it is useful to re-calibrate the 'funding ratios' of Table 1 in terms of w_z, as described in Table 2:

[7] Only the most basic feature of the application of the RAE incentive scheme in England and Scotland is covered here. For a detailed analysis, see the recent articles by Chatterji and Seaman (2006, 2007).

Table 2
Funding of Research in England and Scotland, 2002–2003 (rating weights, w_z)

	HEFCE (%)		SHEFC (%)	
RAE Rating	RAE1996	RAE2001	RAE1996	RAE2001
1	0	0	0	0
2	0	0	0	0
3b	7.3	0	7.1	0
3a (unchanged)	11	5	11.1	0
3a (improved)	11	5	11.1	11.7
4	16.4	16	17.3	18.1
5	24.7	30	26.7	32.7
5*	29.6	44	26.7	37.5

Note: Percentages rounded to add up to 100%.

Table 2 shows that whereas in 1996 departments ranked 1–3a obtained 29.3% of overall funding for England and Scotland, in 2001 the figures had dropped to 10% and 11.7%, respectively, whereas the share of 5 and 5* ranked departments had risen from 54.3% to 74% and 70.2%, respectively, with the share of four-ranked departments remaining essentially constant around 16–18%. Both the shift to the right and the increased gradient of the funding curve are evident when the figures in Table 2 are plotted on the rating weight/departmental ranking plane, as in Figure 2(a) and (b).

When is the RAE the optimal incentive scheme?

The above analysis shows fairly conclusively that as an incentive mechanism, the RAE unambiguously favours a centres-of-excellence policy, where the most productive researchers are concentrated in a few 'starred' departments.[8] The simple analytical apparatus sketched in the previous section enables us to determine when such a policy is likely to be optimal.

Consider first the case of a discipline where all departments are of the generalist type and specifically assume that there are D departments, that each department has with T researchers,[9] each of whom tackles a different research task.

The problem that the RAE is (albeit indirectly) trying to solve is how to allocate researchers of varying quality (or 'probability of task success') to a given

[8] As suggested by a referee, newer forms of Research Assessment whereby departments are ranked not by some weighted average of 'productivity' but by the whole distribution may modify the details of the analysis performed here, but not its main conclusions, as long as the overall effect is to favour the concentration (as opposed to the dispersion) of talent.

[9] The term 'researcher' ought not to be interpreted literally, in so far as there may be economies of scale in tackling any given task and thus there may be a minimum efficient scale for each research team assigned to a task.

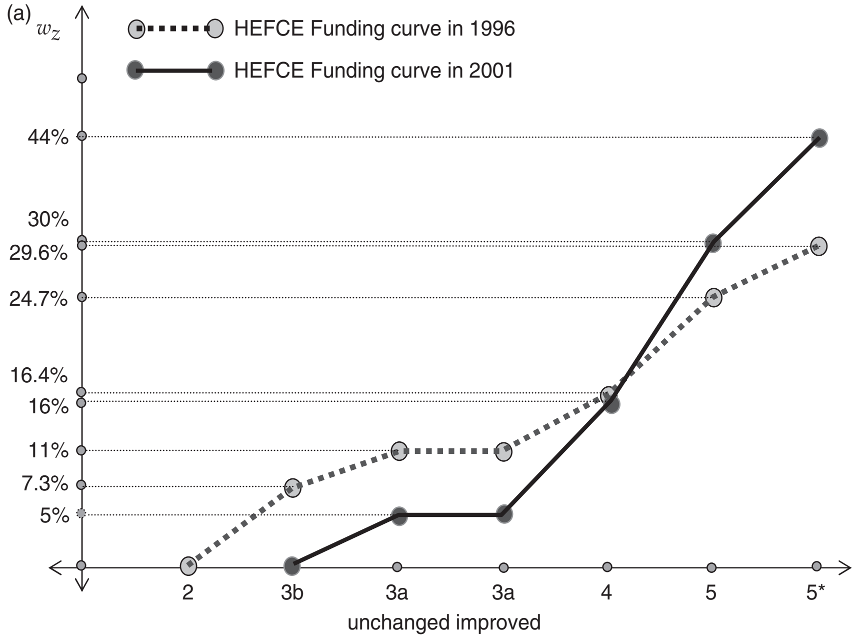

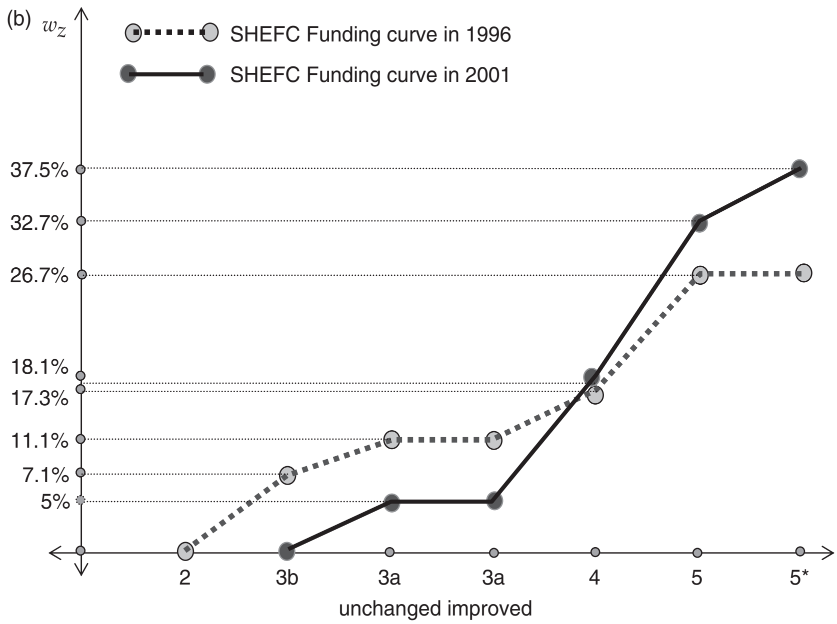

Figure 2. Funding gradients.

number of departments. For the case of generalist departments we can establish the following:

Proposition 1: if a particular discipline is organized by means of D generalist departments, each of which tackles all T tasks that have to be overcome to achieve overall success, then the allocation of $N = D \times T$ researchers of productivity p_i, (as measured by the probability of success in any individual research task) with $p_1 \geq p_2 \geq \ldots \geq p_N$ that *maximizes* the probability of overall success is to rank researchers according to their individual productivity p_i and then assign them to departments in *decreasing* order.

In other words, the first T 'best' researchers would go to the top-ranked department, the next T 'best' researchers to the second best-ranked department, and so on. The same logic applies even if the various routes to success pursued by each department involve a different number of obstacles. The proof of this remarkable result is surprisingly simple if use is made of the powerful technique of majorization and of the notion of Schur convexity.[10] (see La Manna, 2002 for details.)

Proposition 1 suggests that, *for the case of generalist departments*, the RAE provides the best incentives for the allocation of research talent across departments. Of course, the likelihood of *all* 69 disciplines to follow the 'generalist' route is quite small and thus it is worth exploring how well the RAE incentives work for those disciplines in which the organization of research is more akin to the 'specialist department' blueprint, whereby departments concentrate all their resources in one specific research area (obstacle) and pursue all the possible alternatives that can yield a successful solution. Remarkably, we can show (as a kind of dual to Proposition 1) that:

Proposition 2: the allocation of resources that maximizes efficiency when research is carried out by generalist departments *minimizes* the overall probability of success when departments specialize in individual research tasks.

The intuition behind Proposition 2 is quite straightforward: the other side of the 'centres of excellence' policy is, of course, the creation of centres of mediocrity, which, in the extreme case when all departments specialize in different obstacles on the path from ignorance to knowledge, inevitably leads to the creation of bottlenecks where the least productive researchers effectively block (in stochastic terms) the flow of knowledge.

Propositions 1 and 2 show that, if the organization of research within each discipline is taken as exogenous (i.e., whether it is carried out through generalist or specialist departments), the incentive mechanism implicit in the RAE has sharply contrasting effects: in so far as it is optimal for generalist

[10] See La Manna (2002) for details, where theorems originally devised for systems reliability (in particular by El-Neweihi *et al*., 1986) and the theory of majorization (see Marshall and Olkin, 1979) are applied to research design.

Table 3
Overall probability of success and relative performance under Optimal and RAE-induced research assignments[a]

	Probability of success under Optimal assignment (π^O)	Probability of success under RAE-induced assignment (π^{RAE})	$\frac{\pi^O - \pi^{RAE}}{\pi^{RAE}}$ (%)
$P_H = 0.9$, $P_M = 0.3$, $P_L = 0.1$	0.834	0.198	320
$P_H = 0.9$, $P_M = 0.4$, $P_L = 0.1$	0.876	0.26	236
$P_H = 0.9$, $P_M = 0.5$, $P_L = 0.1$	0.912	0.302	202
$P_H = 0.8$, $P_M = 0.3$, $P_L = 0.2$	0.721	0.34	111
$P_H = 0.8$, $P_M = 0.4$, $P_L = 0.2$	0.788	0.446	76
$P_H = 0.8$, $P_M = 0.5$, $P_L = 0.2$	0.849	0.518	63

Note: [a]The probabilities of overall success under the two regimes can be easily computed as $\pi^O = [(1-(1-p_H)(1-p_M)^2(1-p_L)]^4$ and $\pi^{RAE} = [1-(1-p_H)^4][1-(1-p_M)^4]^2[1-(1-p_L)^4]$.

departments but is the worst imaginable for specialist ones. The obvious question is: what is the efficiency loss of the RAE-induced incentives on disciplines that organize their basic research programme by means of specialist departments?

The analytical apparatus sketched above can be used to obtain estimates for any given discipline of the efficiency gains (or losses) of any allocation of talent across departments compared with the optimum. To obtain a feel for the magnitudes involved, consider the following highly simplified example.

Suppose that 16 research units have to be allocated among four departments, each specializing in one of four research tasks. Suppose that one quarter of research units are highly productive, i.e., their individual probability of success is p_H, one quarter have low productivity p_L and the rest are of medium research performance p_M, i.e., $p_H > p_M > p_L$. In this case the optimal assignment of research resources requires talent to be spread uniformly across departments, so that each of the four departments would include one high-productivity unit, two medium-productivity ones, and one 'deadwood'. On the other hand, the RAE-induced allocation of talent ('centres of excellence') would produce one 'star' department wholly staffed by p_Hs, two departments staffed by p_Ms, with the remaining one being staffed by p_Ls. For any distribution of talent $\{p_H, p_M, p_L\}$, we can now compute the overall probability of success, $\pi(p_H, p_M, p_L)$, under the optimal research assignment and compare it with the probability of success induced by the RAE under alternative scenarios. Table 3 summarizes the relative efficiency of the two regimes for some combinations $\{p_H, p_M, p_L\}$ and calculates the percentage gain of switching from an RAE-induced allocation to the optimal one.

Thus, the answer to Problem 1 is that, unless all the 69 broad disciplines assessed in the RAE turn out to be made up *entirely* of *generalist* departments (so that no two or more researchers from different departments ever work on the same type of 'research obstacle'), then the rigid application of the same (implicit) reward scheme across all disciplines is definitely sub-optimal. It is hardly surprising that if research arrangements vary across disciplines, then for some disciplines the diffusion of talent is more efficient than its concentration. Moreover, the calculations reported in Table 3 suggest that, especially in disciplines where certain research obstacles may be less 'fashionable' or less likely to yield measurable research outputs, the misallocation of research resources due to the implementation of a 'centres of excellence' strategy to specialist research organizations may involve substantial costs.

Extension to correlated productivities

At this stage of the argument, we can relax the assumption that the productivity of an individual researcher is uncorrelated with the productivity of his/her colleagues and thus consider the effects of externalities. A full analysis of the whole range of possible externalities is beyond the scope of this paper, but some significant examples may suffice to convince the reader that the superiority of dispersing talent across specialist departments is a fairly robust result.

The strategy followed in this section is to introduce productivity externalities so that the analysis is strongly biased *against* talent dispersed across specialist departments.

As a first indirect test of the robustness of our results, we can perform the following thought experiment. Consider the same allocation problem summarized in Table 3, where 16 research units ('researchers') have to be allocated among four departments. Again one quarter of research units have an individual probability of success p_H, one quarter have low productivity p_L, and the rest are of medium research performance p_M, i.e., $1 > p_H > p_M > p_L > 0$. Propositions 1 and 2 guarantee that the system formed by four specialist departments is always superior to the system made up of four generalist departments, *if there are no* externalities. Two types of externality can be introduced now: a *positive* externality, whereby lower-productivity researchers gain from working with higher-productivity colleagues (the 'infectious enthusiasm' effect) and a *negative externality* whereby the productivity of the ablest researchers is reduced by the interaction with lower-ability colleagues (the 'sinking in the back waters' effect). It ought to be obvious that the superiority of dispersion of talent across specialist departments is *strengthened* by introducing the former externality,[11] in so far as the RAE-induced scheme *minimizes* heterogeneity within each department and hence nullifies the positive externality. A more telling test for the superiority of specialist departments is to handicap them by assuming

[11]Bhattacharya and Newhouse (2005) construct a simple model where researchers of lower ability are assumed to benefit from joining a department with higher average ability, whereas the productivity of high-ability researchers is not affected by the research environment in which they operate (they mention the example of Einstein working out the general theory of relativity in the uninspiring confines of the Swiss Patent Office). Not altogether unsurprisingly, they find that a scheme such as the RAE, which favours the creation of 'centres of excellence', may be inefficient in so far as it discourages heterogeneous (and hence more productive) departments. (I owe this reference to a referee.)

Table 4
Overall probability of success under Arrangement by Specialist Department with No externality and by Generalist Department with Maximal Externality[a]

	Probability of success under Generalist assignment with no externality	Probability of success under Specialist Department with Maximal externality	Maximal externality coefficient α
$P_H = 0.9,\ P_M = 0.3,\ P_L = 0.1$	0.834	0.627	3
$P_H = 0.9,\ P_M = 0.4,\ P_L = 0.1$	0.876	0.511	2.25
$P_H = 0.9,\ P_M = 0.5,\ P_L = 0.1$	0.912	0.43	1.8
$P_H = 0.8,\ P_M = 0.3,\ P_L = 0.2$	0.721	0.18	1.5
$P_H = 0.8,\ P_M = 0.4,\ P_L = 0.2$	0.788	0.6	2
$P_H = 0.8,\ P_M = 0.5,\ P_L = 0.2$	0.849	0.511	1.6

Note: [a]The probabilities of overall success under the two regimes can be easily computed as $\pi^O = [(1-(1-p_H)(1-p_M)^2(1-p_L)]^4$ and $\pi^{EX} = 1-(1-p_H(\alpha p_M)^2\alpha p_L)^4$.

(somewhat counterfactually) that the positive externality mentioned above does not apply to them but *only* to generalist departments, while at the same time allowing for dispersion of talent.[12] Specifically, the individual probability of success of less productive researchers in generalist departments is multiplied by an externality coefficient $\alpha \gg 1$. If it is assumed (realistically) that the externality does not result in lower-productivity units 'leapfrogging' their abler colleagues, the upper bound of the multiplier coefficient is given by $\alpha = Min\left[\frac{p_H}{p_M}, \frac{p_M}{p_L}\right]$.

Table 4 shows that even with positive externality multipliers ranging from 3 to 1.5, the heavily handicapped arrangement by specialist departments is still superior.

A second indirect test of the superiority of dispersion of talent across specialist departments is provided by introducing a *negative* externality applied *only* to specialist departments, in the sense that the productivity of higher-ability researchers is reduced if they work with lower-ability colleagues. To get a feel of how robust our model is to the introduction of this negative externality, we can compute the *coefficient of minimum degradation* δ by which the productivity of the ablest researchers has to be reduced in order to make the centre-of-excellence strategy superior to the dispersion of talent. Using the same range of parameters as in Tables 3 and 4, it can be seen that in order to reverse the superiority of specialist departments with dispersed talent, the extent of the negative externality must reach unrealistically high levels. For example, for all but one of the parameter sets in Table 5, the productivity of the ablest researchers must

[12] Notice that if no handicap is introduced, the arrangement by specialist department is unambiguously superior. The proof is trivial: suppose the effect of the positive externality is to raise the probability of success of each researcher to a common value, $\bar{p} < 1$, then if in the simplest 2 × 2 scenario described in Figure 1 we compare the overall probability of success for the best arrangement by specialist departments $\pi_4(\bar{p}) = (2\bar{p} - \bar{p}^2)^2$ and by generalist departments $\pi_1(\bar{p}) = 2\bar{p}^2 - \bar{p}^4$, we find that the former always outperforms the latter (as $\pi_4(\bar{p}) = \bar{p}^2(2-\bar{p})^2 > \bar{p}^2(2-\bar{p}^2) = \pi_1(\bar{p})$).

Table 5
Minimum degradation coefficient necessary to reverse the superior efficiency of specialist departments with dispersed talent

	Coefficient of minimum degradation δ	$\tilde{p}_H = \delta p_H$
$P_H = 0.9$, $P_M = 0.3$, $P_L = 0.1$	0.274	0.246
$P_H = 0.9$, $P_M = 0.4$, $P_L = 0.1$	0.132	0.119
$P_H = 0.9$, $P_M = 0.5$, $P_L = 0.1$	(0[a])	(0[a])
$P_H = 0.8$, $P_M = 0.3$, $P_L = 0.2$	0.497	0.398
$P_H = 0.8$, $P_M = 0.4$, $P_L = 0.2$	0.458	0.366
$P_H = 0.8$, $P_M = 0.5$, $P_L = 0.2$	0.302	0.242

Note: [a]Even if the productivity of ablest researchers is reduced to zero by the negative externality, the arrangement by specialist departments is still superior.

fall *below* the productivity of their next-best colleagues and in all cases it must be more than halved.

The examples provided above suggest that the assumption of uncorrelated productivity among researchers is not crucial to establishing the inherent superiority of (a) organizing research activities by specialist departments and then (b) ensuring that talent is dispersed across them.[13]

V The RAE as an Instrument to Improve the Organization of Research

In this section, I shall examine the question of whether the RAE can be deployed as an instrument *to alter* the organization of research in some disciplines. In fact, whereas in some fields of research the type of research structure can be considered as exogenously fixed, it seems likely that, at least for some areas of scientific endeavour, the way in which research is organized may be amenable to being influenced by careful design of an appropriate reward mechanism. In line with the rest of the paper, which aims at highlighting the potential benefits of examining the issue of research design, I shall contrast the two extremes of a pure ‘specialist departments’ set-up and of a pure ‘generalist departments’ scheme.

Unsurprisingly, the same techniques used above to establish the optimal assignment of research inputs of varying efficiency for a *given* research

[13] In their analysis of long-term partnership formation with two-sided matching with heterogeneous players, Burdett and Coles (1999) report the possibility of both positive or negative assortative matching in equilibrium. More specifically, high-productivity researchers would prefer forming partnership with similar individuals if low-productivity players are uniformly ‘bad’ (i.e., have low probabilities of success in whichever task they perform), whereas if low-productivity researchers are good at one task, then high-productivity individuals may want to be matched with them (because they can extract most of the surplus). On the contrary, in our model, negative matching (resulting in heterogeneous departments) is always preferable on efficiency grounds. One suspects that the difference is due mainly to the assumption of two-sided matching in Burdett and Coles and the emphasis on system-wide efficiency in our model. (I owe this reference to a referee.)

arrangement can be used to determine which type of research design is more efficient. It is reasonably straightforward to show that, *ceteris paribus*,

(a) research by specialist department is *always* more efficient than research by generalist departments.[14]
(b) If research is carried out by specialist departments, then talent should be spread across *tasks* as evenly as possible.

Two observations are in order here: point (a) applies quite generally, e.g., when different disciplines have different numbers of critical tasks, different numbers of alternatives paths, etc. Point (b) has an interesting practical application in so far as it draws attention to the importance of assigning resources to individual (basic) tasks, as opposed to departments. This could be done by providing a finer partition of individual disciplines, allocating resources evenly to each sub-discipline. Although at one level this is a more interventist approach than the one currently adopted by the RAE, at another level it is less so, because the allocation of resources to individual departments within the same sub-discipline becomes immaterial.[15]

What magnitude of efficiency gain would be achieved if the way a given discipline carries out its programme is changed from (optimally assigned) *generalist* departments to (optimally assigned) *specialist* departments?

To answer this question, again it is revealing to consider a numerical example. Suppose that for a given discipline the basic research programme requires the overcoming of four research obstacles (or tasks) and that research inputs are available in four quality levels, as measured by their probability of success in any one task: $p_E \geq p_G \geq p_S \geq p_U$, where the subscripts stand for Excellent, Good, Satisfactory, and Unsatisfactory.[16] Suppose that there are eight departments of four units each and that research ability is distributed as follows: one-eighth of the units are excellent, one quarter are good, another quarter satisfactory, and three-eighths unsatisfactory. If all departments pursue a 'generalist' research strategy (as is implied by the RAE-induced reward scheme), then the most efficient allocation of talent is by centres of excellence, as proved in Proposition 1, whereas if research is organized by means of specialist departments then talent should be spread across *tasks*. The improvements in the overall probability of success due to the switch from generalist to specialist departments are very significant for any reasonable distribution of talent. Indeed, the superiority of specialist departments is independent of the distribution of talent, as can be seen comparing the probabilities of overall success when the quality of research units is uniform.

The calculations reported in Table 6 show that the switch from a generalist to a specialist type of organization of research can increase the overall probability of success by margins ranging from the respectable to the spectacular. Table 6

[14] See La Manna (2002) for details.

[15] Provided that, as it is likely to be the case, there are more departments than sub-disciplines.

[16] I am simplifying here by ignoring the self-selection aspect of the RAE whereby departments are allowed to choose the researchers included in the assessment. On this, see Talib (1999).

Table 6
Overall probability of success and relative performance of Specialist vs. Generalist Research Arrangements

	Probability of success under 'generalist' strategy $= \pi^G$	Probability of success under 'specialist' strategy $= \pi^S$	Efficiency gain: $\frac{\pi^S - \pi^G}{\pi^G}$ (%)
$P_E = 2/3$, $P_G = 1/3$, $P_S = 1/6$,	0.218	0.718	225
$P_U = 1/12$ ($\bar{p} = 0.315$)	(0.073)	(0.814)	(1004)
$P_E = 0.7$, $P_G = 0.5$, $P_S = 0.3$,	0.343	0.897	161
$P_U = 0.1$ ($\bar{p} = 0.4$)	(0.187)	(0.934)	(398)
$P_E = 0.75$, $P_G = 0.5$, $P_S = 0.25$,	0.404	0.909	124
$P_U = 0.125$ ($\bar{p} = 0.40625$)	(0.198)	(0.939)	(374)
$P_E = 0.8$, $P_G = 0.6$, $P_S = 0.4$,	0.577	0.976	69
$P_U = 0.2$ ($\bar{p} = 0.5$)	(0.403)	(0.984)	(144)
$P_E = 0.9$, $P_G = 0.7$, $P_S = 0.2$,	0.802	0.983	22
$P_U = 0.1$ ($\bar{p} = 0.475$)	(0.34)	(0.977)	(186)
$P_E = 0.9$, $P_G = 0.7$, $P_S = 0.5$,	0.829	0.996	20
$P_U = 0.3$ ($\bar{p} = 0.6$)	(0.67)	(0.997)	(48)

Note: $\bar{p} \equiv (p_E + p_G + p_S + p_U)/4$.

also confirms that the greater efficiency of specialist departments does not depend on the uneven distribution of talent, but is in fact reinforced for the case of uniform distributions.

VI Conclusions

The model used in this paper to provide a new perspective on the theoretical foundations of the RAE is not meant to capture all the aspects of the multi-faceted process of research. It merely highlights the potential benefits of analysing the issue of research design by drawing on the concepts and techniques of systems reliability theory. In order to focus on the essentials, I have opted throughout for simplicity over completeness. To mention but one simplifying assumption used throughout, a more complete model would relax the assumption that research inputs are perfectly adaptable and can be applied to tackling any research task. However, accounting for this, and other factors, is unlikely to alter the main implications of the simple model:

(i) unless *all* disciplines carry out their basic research programmes through the very same strategy, namely by *generalist* departments, then the reward structure implicit in the RAE is sub-optimal;
(ii) for any discipline that carries out research by means of *specialist* departments, RAE incentives point in the completely wrong direction, by promoting the concentration of talent when the optimal allocation calls for its even distribution; and
(iii) if research design is regarded as amenable to change, then the RAE ought to promote research by specialist departments, by applying a

finer definition of 'discipline' broken down by essential tasks and by allocating resources evenly across sub-disciplines.

If one takes the perspective of considering the University research sector as a substantial part of the process for the creation and testing of new ideas, thereby taking into account the inherent multi-path and multi-task nature of research, then the RAE becomes an (indirect) incentive mechanism for allocating research talent across research tasks and path. The notion of 'department' does not fit perfectly into this classical resource assignment problem, but it is close enough as a first approximation. Thus, the ultimate message of this model is that University research funders cannot ignore the large indirect effect their incentive mechanisms have on the organization of research and indeed should devise schemes that encourage efficient research design. Finally, mention should be made of one factor that has been omitted from our model, following its neglect in the RAE incentive scheme, namely an analysis of the *relative* importance of different disciplines. As convincingly argued by Chatterji and Seaman (2006), it is unlikely that an optimal economy-wide incentive scheme for the promotion of research should assign equal weights to all disciplines. The exciting task of modelling the interactions among different disciplines (and sub-disciplines) and their response to alternative incentive mechanisms is the subject of future research.

References

Ball, D. F. (1997). Quality measurement as a basis for resource allocation: research assessment exercises in UK universities. *R&D Management*, **27**, 3, pp. 281–9.

Ball, D. F. and Butler, J. (2004). The implicit use of business concepts in the UK RAE. *R&D Management*, **34**, 1, pp. 87–97.

Bessant, J., Birley, S., Cooper, C., Dawson, S., Gennard, J., Gardiner, M., Gray, A., Jones, P., Mayer, C., Mc Gee, J., Pidd, M., Rowley, G., Saunders, J. and Stark, A. (2003). The state of the field in UK management research: reflections of the RAE panel. *British Journal of Management*, **14**, 1, pp. 51–68.

Bhattacharya, A. and Newhouse, H. (2005). Allocative efficiency and incentive schemes for research. Mimeo, Department of Economics, University of York, UK.

Burdett, K. and Coles, M. (1999). Long-term partnership formation: marriage and employment. *Economic Journal*, **109**, pp. F307–34.

Chatterji, M. and Seaman, P. (2006). Research assessment exercise results and research funding in the United Kingdom: a comparative analysis. *Education Economics*, **14**, 3, pp. 259–79.

Chatterji, M. and Seaman, P. (2007). Research assessment exercise results and research funding in the United Kingdom: a regional-territorial analysis. *Education Economics*, **15**, 1, pp. 15–30.

De Fraja, G. and Iossa, E. (2002). Competition among universities and the emergence of elite institution. *Bulletin of Economic Research*, **54**, 3, pp. 275–93.

Doyle, J. R. et al. (1996). The judge, the model of the judge, and the model of the judged as judge - analyses of the UK 1992 RAE for business and management studies. *Omega*, **24**, 1, pp. 13–28.

El-Neweihi, E., Proschan, F. and Sethuraman, J. (1986). Optimal allocation of components in parallel-series and series parallel systems. *Journal of Applied Probability*, **23**, pp. 770–7.

Gautier, A. and Wauthy, X. (2007). Teaching versus research: a multi-tasking approach to multi-department universities. *European Economic Review*, **51**, pp. 273–95.

Hare, P. (2002). The UK's research assessment exercise: impact on institutions, departments, individuals. Meriot-Watt University, Mimeo.

HAYRI, A. (1997). The research assessment exercise and transfer of academics among departments. *Risk Decision and Policy*, **2**, 1, pp. 71–86.

HEFCE (2004). RAE 2008: Initial decisions by the UK funding bodies (available at http://www.rae.ac.uk/pubs/2004/01/).

JOHNES, G. (1990). Measures of research output: university departments of economics in the UK, 1984–1988. *Economic Journal*, **100**, pp. 556–60.

JOHNES, G. and JOHNES, J. (1993). Measuring the research performance of UK economics departments: an application of data envelopment analysis. *Oxford Economic Papers*, **45**, pp. 332–47.

LA MANNA, M. (2002). R&D as a multi-component system: theory and applications to research joint ventures. University of St. Andrews, Mimeo.

MARSHALL, A. W. and OLKIN, I. (1979). *Inequalities: Theory of Majorization and Its Applications*. New York: Academic Press.

TALIB, A. A. (1999). Simulation of the submission decision in the research assessment exercise: the 'Who' and 'Where' decision. *Education Economics*, **7**, 1, pp. 39–51.

TAYLOR, J. and IZADI, H. (1996). The 1992 RAE: outcome, outputs and inputs in economics and econometrics. *Bulletin of Economic Research*, **48**, 1, pp. 1–26.

Date of receipt of final manuscript: 14 April 2008