Project selection using payback ratios

Ken Buckingham

HESG January 2005, Oxford.

Summary

At the core of the paper lies the premise that opportunity costs should be used in preference to accounting costs. Accountants may regard savings as negative costs. To an economist a saving is a gain to the organisation, and the opportunity cost of that saving is the appropriate cost concept. Thus when we talk of cost-effectiveness ratios, both cost savings (which can be used to purchase health gains) and health gains themselves must be put on an equal footing. This is done using ‘payback ratios’, the ratio of the gain to the opportunity cost of achieving the gain.

The legitimacy of payback ratios is tested by considering whether they can be used to solve project selection problems. Since payback ratios were designed because conventional cost-effectiveness ratios are incommensurable across quadrants, any test of a payback ratio must ensure that it is able to solve problems where comparisons need to be made across quadrants. The paper argues that an appropriate treatment of opportunity costs should enable such problems to be solved and illustrates how this might be achieved.

Introduction

The use of ratios of resource use to health effects remains popular as a guide to decision makers. For example the National Institute for Clinical Excellence[1] in its guidance notes comments that,

> ‘Incremental cost-effectiveness ratios should be calculated as appropriate.’ ………

and rather enigmatically

> ‘Standard decision rules should be followed in combining costs and QALYs.’

Perhaps the most commonly used decision rule employs the cost-effectiveness ratio. This term is somewhat misleading, since the ‘costs’ are not necessarily true economic costs (ie they are not necessarily opportunity costs). In the South-West quadrant of the conventionally portrayed cost-effectiveness plane, where resources are saved at the expense of reductions in health (see Figure 1), the opportunity cost is that reduction in health and not the resources saved.

There are four principal difficulties with cost-effectiveness ratios.

Firstly, in relation to the representation of uncertainty, where measured health effects are small, the ratio of ‘costs’ to health effects is large, and at the extreme, where effects are estimated to be zero, the cost-effectiveness ratio is infinite (there are an infinite number of nothings in any finite something). This is particularly problematic when we are examining estimates stochastically. Some estimates will take extremely high values, rendering the statistics unstable. Briggs and Fenn[2] make the following observation,

> 'In recognition of the problem of applying standard methods of confidence interval estimation to ratio statistics, a number of analysts have proposed alternative methods for estimating confidence limits for the ICER (Incremental Cost-effectiveness Ratio). These include methods based on confidence boxes, confidence ellipses, the Taylor series expansion, Fieller's theorem and non-parametric bootstrapping.'

Cook and Heyse ([3]) provide a solution based on angular transformations in which the full 360^o of the cost-effectiveness plane is uniquely identified as angular deviations.

Secondly, and in relation to project selection, ratios do not provide information on the extent of the welfare gain. In contrast, a measure of net benefit (or net cost) would provide such information. However, if net benefits are to be determinate we require information about the societal value to be assigned to the measure of health (for example, the price that society is prepared to pay for a QALY). We also require that funding be available for all projects in which the cost per health gain is less than this societal value. If funds are constrained below this amount, a higher shadow price will arise reflecting the scarcity of resources. This paper assumes that funds are constrained, and that a shadow price for resources exists in the health care sector. I refer to the shadow price as lambda.

Thirdly, again in relation to project selection, cost-effectiveness ratios can be misleading if there are indivisibilities in projects. With indivisibilities, it may be the case that the implementation of a project with a low ratio of costs to effects would preclude the implementation of other smaller projects that might otherwise be used to exhaust the budget more efficiently. The issue has been raised, among others by Birch and Donaldson[4] who offer an integer programming solution which,

> '… acknowledges that budget constraints be recognised when evaluating alternative projects and that one should determine whether the implementation of each project would result in different amounts of resources being left over. If different amounts of resources do remain then their alternative uses should be identified and evaluated.'

and contested by others who question whether decision makers operate with sufficient certainty to know whether an individual project will indeed result in breaking a budget constraint. In practice it may be that decision makers have little control over much of their expenditure (the 'must-do' or 'too hard to challenge basket'). They will probably be making decisions over a relatively small 'basket' of projects which are under review (the 'might-do basket' of 'marginal' projects as advocated by Mooney[5]). Given uncertainties about the level of expenditure required for the 'must-dos', the residual budget for the 'might-dos' will be, to say the least, 'fuzzy'. The baseline case described in this paper assumes that indivisibilities are unlikely to be of practical significance to decision makers.

Finally, the interpretation of numerically identical ratios differs across quadrants. In the North-East quadrant, policy makers prefer low ratios of costs to effects than high ratios. In the South-West quadrant, high ratios of 'costs' to effects are preferred, since in this case, these negative costs represent savings. This issue has been noted by other authors, for example. Heitjan et al[6] ,

> ‘The ICER is an ill-defined parameter because it identifies pairs of cost and effectiveness differences that can have altogether different interpretations.’

Laska et al[7]

> ‘Thus, knowledge of the value of the magnitude of the cost-effectiveness ratio alone is insufficient to ascertain which treatment is preferred. To make that determination, the sign of ε (the incremental measure of effect) or of γ (the incremental measure of cost), must also be known …….’

Stinnett et al[8],

> ‘Thus, R hat (the estimator for the cost-effectiveness ratio) has no meaningful interpretation unless it is presented in the context of the quadrant of the incremental effect – incremental effect plane to which it corresponds.’

This issue is the concern of this paper and in the opinion of the author, relates back to the inappropriate definition of ‘cost’.

Payback ratios defined

This paper advocates the use of payback ratios in preference to cost-effectiveness ratios as a guide to decision making. Payback ratios have been discussed previously by the Health Economics Study group[9], but their use in project selection under resource constraints has not been considered. The main aim of this paper is to examine whether ‘payback ratios’ can assist in project selection, particularly where we need to compare projects that arise both in the North-East and the South-West quadrants.

Payback ratios are defined as the ‘incremental gain’ achieved by a course of action divided by the incremental opportunity cost of achieving that gain. In the North-East quadrant the incremental gain is the value of the incremental health improvement and the opportunity cost is the incremental resource required to achieve the gain. In the South-West quadrant, the gain is the incremental saving in resources, and the opportunity cost is the incremental reduction in the value of health.

To enable comparisons to be made between payback ratios from the North-East quadrant, with those that arise in the South-West quadrant, both costs and effects must be expressed in the same units. To do so, we need to apply a value to the measure of health gain. Depending upon the context, that value might be either a shadow price determined endogenously within a decision problem, or an exogenously determined social value. Payback ratios are illustrated in the Figure 1 and Table 1 which show two projects, one in the North-East, the other in the South-West quadrant.

Algebraically we have:

In the North-East quadrant the payback ratio is

(lambda x Q) / R

In the South-West quadrant the payback ratio is

R / (lambda x Q)

We can think of Q as the health effect measured in QALYs, lambda as the shadow price of a QALY, and R as representing resources.

Table 1 Cost-effectiveness ratios compared with payback ratios

Project number	Resources	Health effects	c/e ratio	Payback ratio at a shadow price of £40,000 per QALY
1	£100,000	2 QALYs	£50,000 per QALY	0.8 (£80,000 /£100,000)
2	-£100,000	-2 QALYs	£50,000 per QALY	1.2 (£100,000 /£80,000)

Figure 1 Two projects on cost-effectiveness plane

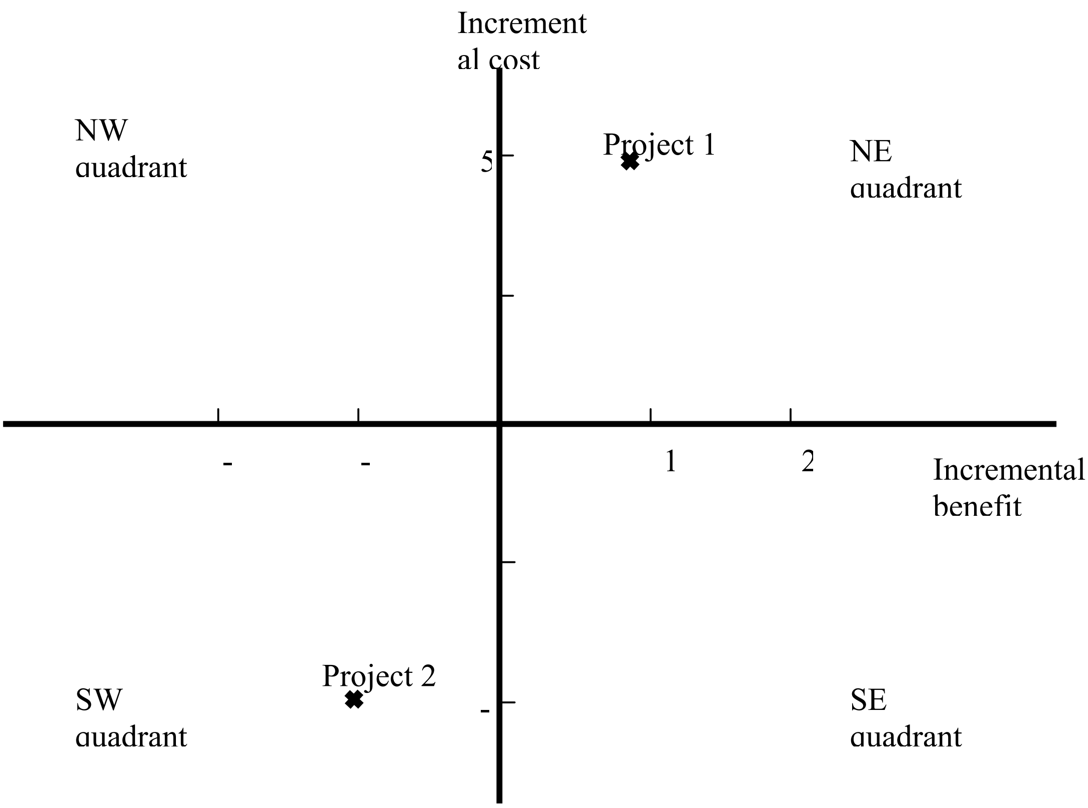

If, in Table 1, we decide to implement projects in which the payback ratio is greater than or equal to 1 (ie we get back at least as much as we put in), then we would implement project 2, but not project 1.

Project selection – the basic case

We begin by considering a simple case, and then consider how the procedure must be adapted to allow for the existence of some mutually exclusive sets of projects within the projects under consideration. The basic case is described as follows:

- Resources are approximately constrained by the current level of spending on a set of projects under review.
- A set of health care projects not presently implemented are also under consideration, some of these can save resources at the expense of a loss of health, others can increase health but require additional resources to do so.
- Projects may be indivisible, but, in the face of 'fuzzy' budget constraints, decision makers fund as closely as possible to their understanding of what the budget might be, without exceeding it.
- No projects are mutually exclusive (mutual exclusivity can arise, for example, if we are considering alternative ways of treating the same group of people for the same condition or using the same piece of land or equipment).
- All resource and health changes are incremental in comparison with the alternative care that would be provided in the absence of the project.
- Projects have already been screened, such that all existing projects in the North West quadrant (ie projects which result in both health losses and resource costs) have been abandoned, while all projects in the South East quadrant (ie projects which result in both health gains and resource savings) have been implemented and are not under consideration.
- Projects currently implemented and projects currently not implemented are treated equally (the role of the currently implemented projects being simply to define the available budget).

Project selection – the algorithm

Our approach is to search iteratively for the value of lambda that is consistent with the objective of maximising health within a specified budget (in this case defined by the value of existing projects). For any given value of lambda I consider the effect of implementing all projects in which the payback ratio is greater than or equal to unity (ie all projects in which gains exceed losses). We need not distinguish whether the project is in the North-East or the South-West quadrant of the cost-effectiveness plane.

Note that a low shadow price of the health gain favours projects in which savings are achieved at the expense of health gains and mitigates against projects in which health gains are achieved at the expense of resources. The algorithm searches progressively from low values of lambda to higher values. As we move from low to high values the balance shifts; such that we would implement more projects in which health is purchased and fewer projects in which savings are achieved. As a result, the net use of resources across all projects tends to increase as the shadow price increases. The algorithm maximises health gain by iterating with increasing the value of lambda, until no more health gain can be achieved within the available budget.

Table 2 Some specimen projects

Project no.	Health effects	resource effects	Currently funded?
1	10	2	No
2	-9	-2	No
3	8	3	Yes
4	7	4	No
5	6	5	No
6	5	6	Yes
7	4	7	No
8	-4	-7	No

In practical applications it is straightforward to program this procedure and to select an appropriate level of resolution in the search for the health maximising value of lambda.

The process is illustrated in simplified form in Tables 2 and 3. Table 2 shows eight projects. Projects 3 and 6 are currently funded, providing 13 units of health gain at the expense of 9 units of resources.

The uppermost section of Table 3 shows what happens to payback ratios as we vary lambda.

The middle section shows which projects would be implemented at the various levels of lambda, assuming that we implement projects in which the payback ratio is greater than or equal to 1 (ie the gain is at least as much as the opportunity cost).

The lower section of Table 3 shows how the health and resource implications that result from implementing projects in which the payback ratio is greater than 1, as lambda is varied. We see that as we increase lambda, more projects that improve health are implemented, so consuming resources, while fewer projects that release resources are implemented. As lambda is increased, we reach a point at which any further increase would result in implementing projects that would exceed the available budget. The maximum health gain achievable arises where 27 units of health gain are obtainable (in comparison with the 13 units before any changes are made). At this point, lambda equals 1, projects 1,3,4,5,6 and 8 are implemented, and 7 units of resources are used. (With hard budgets and divisibility, policy makers might choose to use funds on the project in the North-East quadrant having the highest payback ratio.)

Table 3 Project selection by search over values of the shadow price (lambda).

		payback ratios									
Project no.	At start	Lambda = 0.2	lambda = 0.4	lambda = 0.6	lambda = 0.8	lambda = 1	lambda = 1.2	lambda = 1.4	lambda = 1.6	lambda = 1.8	lambda = 2
1		1.000	2.000	3.000	4.000	5.000	6.000	7.000	8.000	9.000	10.000
2		1.111	0.556	0.370	0.278	0.222	0.185	0.159	0.139	0.123	0.111
3		0.533	1.067	1.600	2.133	2.667	3.200	3.733	4.267	4.800	5.333
4		0.350	0.700	1.050	1.400	1.750	2.100	2.450	2.800	3.150	3.500
5		0.240	0.480	0.720	0.960	1.200	1.440	1.680	1.920	2.160	2.400
6		0.167	0.333	0.500	0.667	0.833	1.000	1.167	1.333	1.500	1.667
7		0.114	0.229	0.343	0.457	0.571	0.686	0.800	0.914	1.029	1.143
8		8.750	4.375	2.917	2.188	1.750	1.458	1.250	1.094	0.972	0.875
	Implementation (0 = not implemented, 1 = implemented)										
Project no.											
1	0	1	1	1	1	1	1	1	1	1	1
2	0	1	0	0	0	0	0	0	0	0	0
3	1	0	1	1	1	1	1	1	1	1	1
4	0	0	0	1	1	1	1	1	1	1	1
5	0	0	0	0	0	1	1	1	1	1	1
6	1	0	0	0	0	0	1	1	1	1	1
7	0	0	0	0	0	0	0	0	0	1	1
8	0	1	1	1	1	1	1	1	1	0	0
	Net consequences										
Total health gain	13	-3	14	21	21	27	32	32	32	40	40
Total resource use	9	-7	-2	2	2	7	13	13	13	27	27

Project selection with mutually exclusive sets

There is nothing particularly surprising about the adaptation to a situation in which some projects can be freely selected, while others are part of mutually exclusive sets. When choosing projects we must choose between all that can be freely selected and the optimum project within each mutually exclusive set. As with constrained optimisation using cost-effectiveness ratios, when we are choosing from within a mutually exclusive set, we examine the highest ratio then consider whether the incremental effects from the next best ratio are nevertheless worthwhile. If the next best project increases welfare (has a payback ratio greater than unity in comparison with the project having the highest ratio) it should be implemented. This process may need to be repeated if a third alternative has worthwhile incremental effects in relation to the second, and so on). With payback ratios we simply expand the analysis into the South-West quadrant.

The multi-stage process required to solve for mutually exclusive sets may be better illustrated graphically than in tables. We see the first stage, the selection within each mutually exclusive set, in Figure 2. This shows project selection within a set of 4 mutually exclusive projects. The fact that preferred ratios in either the North-East or the South-West quadrant always like closer to the dominant South East quadrant, implies that in the North-East quadrant, the more clockwise the project, the more it is preferred, while in the South-West quadrant, the more anticlockwise the project, the more it is preferred.

Also note that the payback ratio of any project which lies clockwise of the shadow price line in the North-East quadrant is greater than unity and in the absence of mutual exclusivity would be chosen (subject to the budget constraint). (It is easy to see that for any given level of health gain a project on a radial lying clockwise to the shadow price line would cost less to implement than a project above the shadow price line and hence would be preferred.)

Also note that the payback ratio of any project which lies anti-clockwise of the shadow price line in the South-West quadrant is greater than unity and in the absence of mutual exclusivity would be implemented. (Again, it is easy to see that for any given level of saving in resources, a project on a radial anti-clockwise to the shadow price line would result in a smaller reduction in health than a project above the shadow price line and hence would be preferred.)

Figure 2 shows the effect of varying lambda on the choice of the optimum project from a set of four mutually exclusive projects. In successive panels, the value of lambda increases. In each panel, the hatched line allows us to examine the incremental effect in relation to the project with the highest payback ratio.

As the search for the optimising lambda begins, we have low values indicated by the relatively shallow gradient of the shadow price line in Panel 1. In Panel 1, both A and B have payback ratios greater than unity (they are anti-clockwise to the shadow price line through the origin. Although project A has a higher payback ratio than B, it is nevertheless worthwhile implementing B, since B has an incremental payback ratio greater than unity in relation to A.

In Panel 2, at the higher shadow price, project A has the highest payback ratio and both projects A and B have payback ratios greater than unity, however the incremental payback ratio of project B in relation to A is less than one. Project A should be selected at this shadow price.

In Panel 3, project C has a higher payback ratio than project A, and would be selected from the set. Neither project C to the South-West of A, nor project D, to the North-East of A, have incremental payback ratios in relation to C that are greater than one.

In Panel 4, projectC has a higher payback ratio than project D, however with the high shadow price for the value of health, the small incremental health gain from project D in relation to project C is worthwhile achieving. At this shadow price project D would be preferred.

If we do identify a project that might be selected because its incremental effect in relation to a project with a higher payback ratio would justify its conclusion, then we should also check to ensure that no other project has a payback ratio greater than unity in comparison with it.

At each iteration of the search for the optimising shadow price, the selected project (provided of course that there is one with a payback ratio greater than unity that would be selected) within each mutually exclusive set is considered together with all other projects having a payback ratio greater than one. If at that shadow price we can achieve the maximum health gain, while staying within the budget, then that group of projects is the optimal group.

Figure 2 Selecting from a mutually exclusive set at different shadow prices.

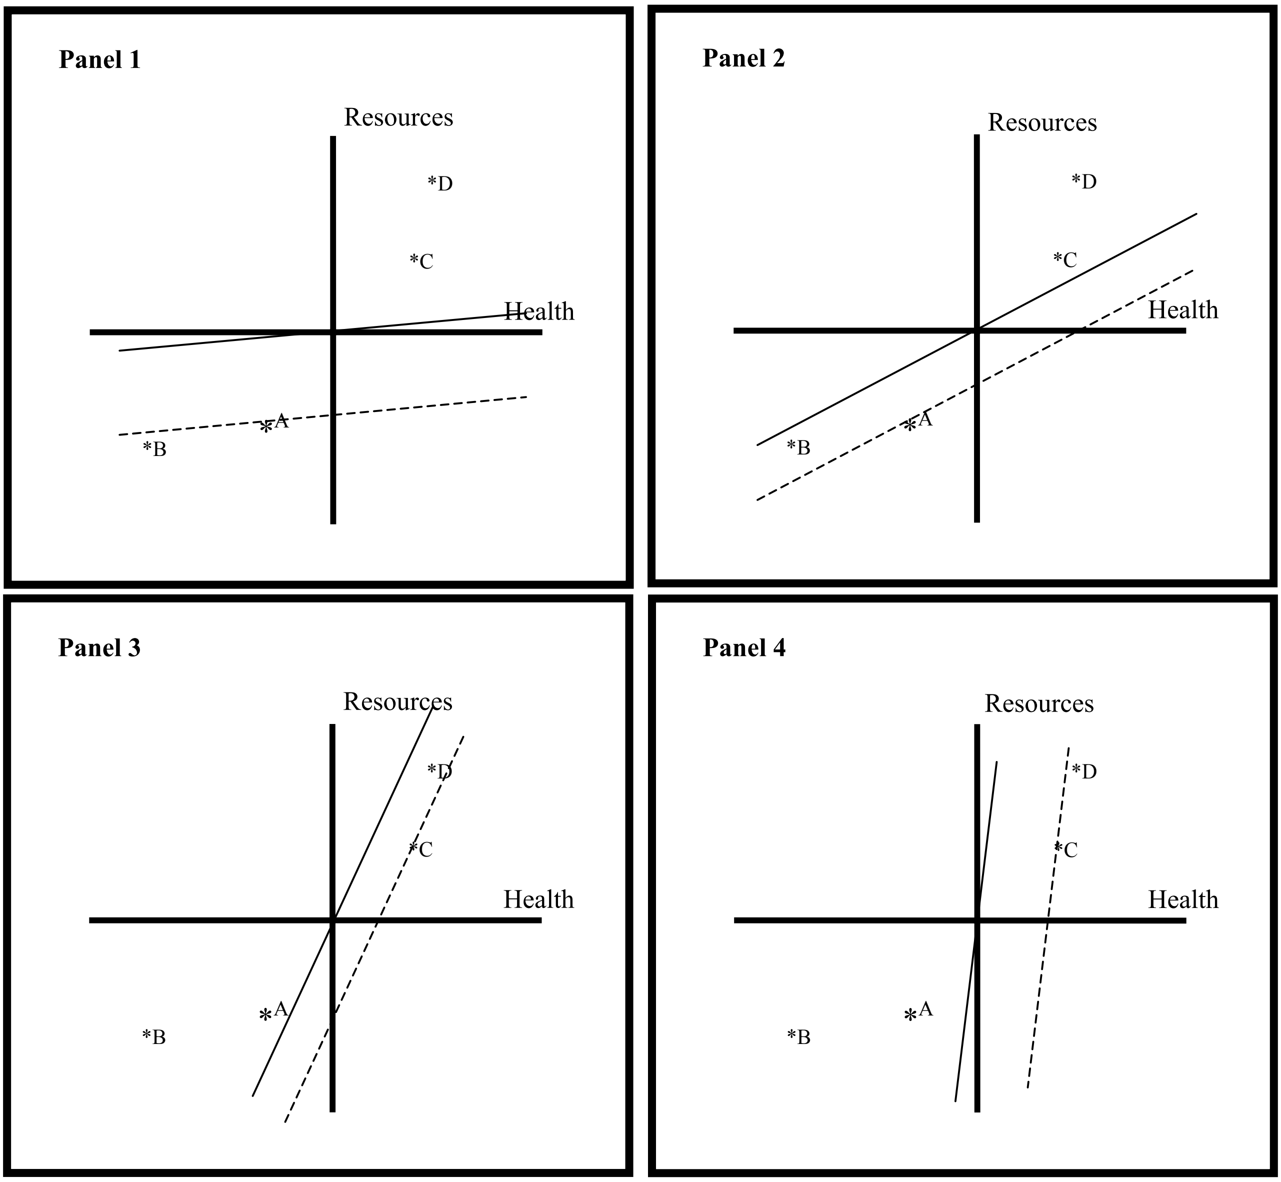

Discussion

The method described in this paper illustrates an attempted solution to a selection problem, in which the final budget constraint is not defined in advance, but rather is determined within the problem itself, and achieved by the release of funds from projects in the South-West quadrant. The underlying aim of the paper is not to present an alternative to other programming approaches to optimisation. I would be surprised if programmes do not exist that simultaneously consider resource releasing and resource consuming projects. It may indeed be the case that the method I describe, if correct, is in fact already used, and that the payback ratio is implicit within their solutions. I would be grateful for comments on whether this is the case.

However, my main intention is to reiterate the point that others have made, ie that cost-effectiveness ratios are misleading in the South-West quadrant, and to suggest that payback ratios should be considered as an alternative, an alternative that is more in line with economic concepts of cost.

Some further comments may be in order.

The size of the budget constraint and the consequent value of lambda

Since the only purpose of identifying projects that are currently funded is to define the initial budget constraint, it would be relatively straightforward to amend the budget constraint, either upwards or downwards. If it were judged to be politically difficult or socially disadvantageous to discontinue existing projects, this might be incorporated into the model, either by increasing the cost of implementation, if for example the withdrawal of services needs to be accompanied by public education; or by increasing the opportunity cost of health lost, if, for example, the discontinuation of a service results in adverse psychological effects. Obviously a combination of these two options is also possible.

If the initial budget constraint is changed, then the resulting optimising value of lambda will be affected. It is thus possible to examine the effect of different budgets on the shadow price of health.

The use of net benefits

Because the reason behind this study was to examine the legitimacy of payback ratios as an alternative to cost-effectiveness ratios, I have used payback ratios, in preference to net benefits in the algorithm described. It should be said, however, that either measure could equally have been used in the algorithm. At each iteration over values of lambda, we simply need to identify those projects with a payback ratio greater than unity. We could instead have used a net benefit greater than zero as the relevant criterion, since a net benefit greater than zero, always implies a payback ratio greater than unity (and vice-versa) as shown below.

Net benefits can defined as:

$$\textit{Net benefit} = \textit{lambda} \times Q - R$$

To examine the inequality condition that Net Benefit is greater than 0. In the North-East quadrant we have:

$$\textit{lambda} \times Q - R > 0$$

hence $\textit{lambda} \times Q > R$

or $(\textit{lambda} \times Q) / R > 1$ *(ie the payback ratio is greater than 1.*

In the South-West quadrant we have:

$$\textit{lambda} \times (-Q) - (-R) > 0$$

hence $R > \textit{lambda} \times Q$

or $R / (\textit{lambda} \times Q) > 1$ *(ie again, the payback ratio is greater than 1.*

(Naturally, the fact that the inequalities around net benefit = zero and payback ratio = unity are simultaneously true, does not imply that net benefits are the same as payback ratios.)

Integer programming

It might be possible to adapt the methods described here to an integer format. I have not seriously examined that possibility, however, it seems to me that the indivisibilities that give rise to need for integer programming are possibly something of a ‘red herring’. Perhaps my experience in health authorities is not typical, but there the situation appears to be more one of the ‘fuzzy’ (and usually greatly overspent) budgets that I described above. Health Authorities must be some of the most complex organisations. They provide a vast range of ‘product lines’, many ‘products’ can be provided using alternative technologies, in a variety of settings and with very uncertain outcomes. Any individual projects will only comprise a small part of total expenditure, and even if all projects were indivisible, the idea that ‘lumpy’ indivisibilities are problematic in the light of all the other complexities that health authorities must face seems unlikely. Moreover, it seems that many projects are, in fact, not ‘indivisible’, for example, we can do more or fewer hip operations if we choose to do so.

The practical significance of indivisibilities

I would welcome the thuoghts of others about whether the situation I have described is very different in PCTs, or even in the Department of Health. At operational levels, I suspect that the concerns are more with intermediate outputs such as the number of patients seen and are not amenable to the sort of analysis described here;

Non-linearity.

Of far greater practical importance might be the question of non-linearity. If costs and benefits alter with the scale at which a project is implemented, the simple solutions provided here would give misleading results. A simple, if imperfect, adaptation might be to include the different scales at which projects are implemented, as a discrete set of projects within mutually exclusive sets.

Multiple constraints

Another, possibly important, deficiency in the method described is that it fails to account for resource constraints other than as financial aggregates. In practice there may be issues such as skill shortages, that need to be considered. The true opportunity cost of using such staff is not fully considered in the method described. The method would need to be adapted for multiple constraints.

But, as I have already said, my main concern has not been to solve all the problems of resource allocation. Less ambitiously, it is to demonstrate the use of a statistic that has more intuitive appeal to managers. If they are getting a large health gain for little cost, they are getting a good ‘return’ on their investment, and the same is true if they are saving a lot of money for a relatively little loss in health. That too will represent a good return in the health gains that it would allow them to purchase elsewhere.

Endnote

Programming techniques offers a computationally efficient way of selecting the optimal mix of projects with indivisibilities. A cruder way of identifying the optimal mix is simply to generate results for all possible combinations of projects, exclude those combinations that exceed the budget and other requirements such as mutual exclusions, and then look for the combination with the greatest health gain. This rather naïve approach does have the advantage, if indeed it be an advantage, that it incorporates indivisibilities. It just tries every bit of the 'jig-saw puzzle' in every position to arrive at the best fit. Actually this is exceedingly simple to programme and quite quick to run if the number of projects under consideration is not too large. With the eight projects from which we select in the example above, there are 256 possible combinations. The set producing the highest health gain within the existing budget, is the same set as selected using the payback ratio algorithm that I have described. The same is true when we re-examine the full set of possible combinations, but exclude cases in which there are more than a single project from a mutually exclusive set. This does not prove that these algorithms will work correctly on all occasions. On some occasions, where the budget constraint is defined and rigid, it might be the case that the payback ratio algorithm identifies a sub-optimal case. It is the author's opinion that this is unlikely to be a serious practical drawback.

References

[1] National Institute for Clinical Excellence, Guide to the Methods of Technology Appraisal, April 2004.

[2] Briggs A, Fenn P. Confidence intervals or surfaces? Uncertainty on the cost-effectiveness plane; Health Economics, 7: 723-740, 1998.

[3] Cook JR, Heyse JF. Use of an angular transformation for ratio estimation in cost-effectiveness analysis; Statistics in Medicine, 19:2989-3003, 2000.

[4] Birch S, Donaldson, C. Applications of cost-benefit analysis to health care; Journal of Helath Economics, 6: 211-225, 1987.

[5] Mooney GH, Russell EM, Weir RD, Choices for health care Edition 2, Macmillan, 1986.

[6] Heitjan DF, Moskowitz AJ, Whang W. Bayesian estimation of cost-effectiveness ratios from clinical trials. Health Economics. 8(3): 191-201, 1999.

[7] Laska EM, Meisner M, Siegel C. Statistical inference for cost-effectiveness ratios.[comment]. Health Economics;6(3): 229-42 1997.

[8] Stinnett AA, Mullahy J. Net health benefits: a new framework for the analysis of uncertainty in cost-effectiveness analysis.[comment]. Medical Decision Making; 18(2 Suppl):S68-80 1998.

[9] Buckingham K, Musings on the meaning of incremental opportunity cost, HESG, July 2003.