"AN EFFICIENT BUDGET ALLOCATION POLICY FOR DECENTRALISATION OF RESPONSIBILITY FOR SITE DECONTAMINATION PROJECTS"

by

C. CORBETT*
F. DEBETS**
and
L. VAN WASSENHOVE ***

94/44/TM
(revised version of 93/66/TM)

* PhD Student at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France.

** Research Associate at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France.

*** Professor of Operations Management and Operations Research at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France.

An Efficient Budget Allocation Policy for Decentralisation of Responsibility for Site Decontamination Projects

Charles Corbett, Frank Debets,

Luk Van Wassenhove

INSEAD, Fontainebleau, France

September 5, 1994

Correspondence address:

Charles Corbett

INSEAD

77305 Fontainebleau Cedex

France

tel: (33) 1 60724000

fax: (33) 1 60724242

Running title: **Budget allocation for site decontamination**

Abstract

Selection and execution of site decontamination projects is often best left to local authorities, in accordance with the subsidiarity principle, even though the budget for such projects is made available through a central authority. In this paper we suggest a practical budget allocation policy which a central authority can employ to allocate budgets to local authorities, while still optimising the central authority's environmental objective function. The procedure is fully consistent with the principle of decentralisation of responsibility for selection and execution of projects, and requires a minimum information exchange between local and central levels. Despite the information asymmetry between local and central levels, incentive compatibility problems can be (partially) prevented by choosing an appropriate evaluation mechanism. At the same time, the procedure is computationally effective and efficient, and can guarantee a fair budget allocation, making it easy to implement and politically acceptable.

Keywords: budget allocation, decentralisation, site decontamination, mathematical programming, Groves mechanism

1 Introduction

Much attention has been paid to economic modelling of environmental problems. Such modelling efforts frequently lead to important insights concerning these problems and the trade-offs they present. One of the primary intentions is often to provide support for environmental policy-making, for example by showing the qualitative effects of different types of policies. Having settled on a particular policy, another important question is how that policy can be made operational. The aim of this paper is to show how the actual implementation of a given policy can also be aided by mathematical modelling.

The problem considered in this paper is the allocation of a central budget to regional authorities for cleaning of contaminated sites, with the aim of achieving maximum overall (national) improvement in the site contamination problem. The problem is modelled after the Dutch situation, but the approach outlined applies to any similar hierarchical budget allocation problems, national or international. In the Dutch case, new polluted sites are continuously detected throughout the country, and coordinated action is urgently required. Indeed, a recent study of the regional approaches to the problem has shown that there is little consistency in decontamination policies. This results in a very small percentage of site decontamination projects actually being completed according to central standards and a poor budget allocation (see Beumer et al., 1991). In a newspaper article[1] the Dutch Ministry of Environment was criticised for using the available money injudiciously, and it was argued that the procedure for controlling the decontamination program is not satisfactory.

The related issue of cleaning up hazardous waste sites has been tackled differently in the US. There, the polluter pays principle has been applied, leading to the well-known "Superfund" legislation. This legislation empowers the EPA to require parties deemed responsible for the pollution to clean up those hazardous waste sites, or to do so itself with funds obtained from lawsuits against those parties. Not only current owners of the site can be held responsible, but also past owners or any firm whose waste is found on the site. Liability is strict; negligence is not necessary for firms to have to pay. It is also joint and several, so any one responsible party

[1] In the Dutch newspaper *Trouw*, 29/1/1993.

can be sued for the total costs, leaving it to them to retrieve as much as possible from the other responsible parties. Even more dramatically, it is retroactive, so firms in full compliance with legislation at the time of dumping can be held responsible now (see, e.g., Bloom & Scott Morton, 1991). The claims have driven firms to bankruptcy, in which case ownership of the site, and thereby liability for the cleanup costs, are transferred to the banks involved.

However, despite the huge amounts of money spent as a result of Superfund, only a very small number of the high-priority sites have actually been decontaminated; RAND has estimated that as much as 88% of insurers' costs related to Superfund have been spent on legal fees, due to the excessive amount of litigation involved.[2] Indeed, new legislation to curb some of the Superfund inefficiencies has been tentatively accepted by environmentalists, chemical firms and insurers.[3] The US situation suggests that the polluter pays principle will not work for the site decontamination problem, rendering a publicly funded cleanup program unavoidable. The Dutch situation, however, clearly shows that such a program is by no means easy to implement; a central authority funding the program will want to be sure that the regional authorities executing the program employ the funds to maximum effect. All these observations lead to the conclusion that there is a need for a mechanism to allocate funds as effectively and efficiently as possible, and to coordinate regional actions, either implicitly or explicitly.

Clearly, if the budget allocation is decided on purely political grounds, it is unlikely to bring maximum environmental benefit. By using an approach as outlined in this paper, decisions can be made in a more objective manner. Moreover, if targets are set for each region, the central authority has to know how much budget each region needs in order to achieve a given target. Coordinating regional targets and budgets is a key benefit of a procedure as suggested here. It should be emphasised, however, that what this paper discusses is not a decision-making mechanism to replace the central authority, but a decision support tool to aid that central authority. The procedure we propose is not restricted to the case of budget allocation for site decontamination. For instance, the Dutch government has recently decided (Vellekoop, 1994) to render the combat of industrial noise more effective by decentralisation of

[2] *The Economist*, August 8, 1992, p. 11.

[3] *The Economist*, May 7th, 1994, p. 87.

responsibility and allocating budgets accordingly to regional authorities; our procedure would be equally applicable in such a context.

The problem of resource allocation in a decentralised setting is of course far from new. The decomposition algorithm for linear programs proposed by Dantzig and Wolfe (1961) was quickly recognised, by Baumol and Fabian (1964) as having the economic interpretation of decentralisation of decision making through transfer prices. Similarly, Kornai and Liptak's (1965) decomposition procedure, later improved by Ten Kate (1972), was interpreted by Burton, Damon and Loughridge (1974) as decentralisation through resource allocation, based on bids for central resources submitted by divisions. Atkins (1974) interprets and compares various decomposition schemes with respect to their organisational implications. Christensen and Obel (1978) compare the computational performance of the Dantzig-Wolfe and Ten Kate procedures in two real organisational settings, and conclude that both can perform well provided good bounds are available and that the procedures are cleverly initiated; generally, however, this requires much more information to be available at the central level than is needed just to perform the iterations.

The mathematical programming literature on resource allocation generally ignores the issue of incentive compatibility, though. Groves and Loeb (1979), for instance, quote Jennergren (1971) who shows that with some of the procedures mentioned above, including the Dantzig-Wolfe decomposition procedure, regional managers may have incentives to report incorrect information. There is, of course, a vast literature on mechanism design to counter this type of problem; see, for example, Green and Laffont (1979) or Laffont and Tirole (1993). In this paper we show how a decomposition approach based on mathematical programming can be supplemented by a Groves mechanism (Groves 1976, Groves and Loeb 1979) to achieve incentive compatibility under appropriate assumptions.

This paper is organised as follows. First, in Section 2, we describe the budget allocation problem in more detail. In Section 3, the problem is formulated as a mixed integer linear programming model. Section 4 outlines the criteria any practically useful procedure should meet, and discusses the solution procedure we propose. A simple numerical example is included.

Section 5 addresses practical issues related to implementation of the procedure, and section 6 discusses the moral hazard problems resulting from the decentralisation of responsibility. Specifically, we discuss the two types of moral hazard that may occur, and how a Groves mechanism can be used to counter one of them and reduce the impact of the other. The final section contains some conclusions and areas for further research.

2 The budget allocation problem

For the purpose of site decontamination, the Netherlands is divided into 16 regions: 12 provinces and the four big cities (Amsterdam, Rotterdam, The Hague and Utrecht). Each region has hundreds of polluted sites, which can be decontaminated in a number of ways, including containment to prevent spreading of the pollution, removal and temporary storage of the polluted soil, and total in situ decontamination of the soil. There is certainly no universally applicable cleanup method. For example, in the case of a former chemical plant site in Hamburg, Germany, authorities wished to preserve the existing buildings, making excavation impossible; the material was extracted using boreholes and high-pressure water jets (Johnson, 1989). The costs and environmental effects obviously vary strongly between these approaches, which require different amounts of personnel and storage capacity for polluted soil. The problem is how to allocate the total available budget to the regions, in order to achieve maximum overall environmental improvement.

It is important to note at this point that the number of polluted sites far exceeds the number of decontamination projects that can be undertaken under the current budgetary constraints. This situation will persist in the foreseeable future in spite of the large efforts of the government. The reason is not only lack of money on the part of the central government but especially the fact that the consequences of careless environmental behaviour in the previous decades are now quickly beginning to show themselves through hundreds of contaminated sites popping up all over the country.

It should be clear that the information requirements to characterise all these potential site decontamination projects are huge. An entirely centralised approach to this problem would require all information about each of several possible variants of thousands of possible projects to be centrally available, which would be highly impractical. Moreover, this option would often be politically undesirable or even infeasible. In particular, a fully centralised approach would be in conflict with the subsidiarity principle, according to which executive responsibility for selection and actual cleaning of sites should be delegated to regional or local authorities. However, the information exchange at all levels is still in a preliminary stage, highly manual, and therefore far from perfect. Two major concerns in developing a hierarchical budget allocation procedure for such a situation are therefore:

- *Decentralisation.* Regional authorities are responsible for selection and execution of decontamination projects, for which a budget is allocated to them. Central government determines these budgets on the basis of limited aggregate information provided by the regions. It also specifies the environmental improvement it expects the regions to achieve with their budgets.
- *Minimum information flow.* The information flow in both directions should be minimised, without jeopardising the efficiency of the final allocation.

In The Netherlands, the bottlenecks in the current situation are as follows.

- The limited budget allows only a very small proportion of projects to be carried out.
- Many projects require a specialist, scarce workforce.
- There is insufficient capacity to store all polluted soil.
- The workforce and storage capacity are strictly regional, they cannot be exchanged between the regions.

The precise nature of the restrictions concerning workforce and soil storage capacity is not critical; any other restrictions on regional resources can be incorporated. Restrictions such as those on the exchangeability of regional storage capacity and workforce will be at least equally relevant in an international context.

We obviously need a quantitative measure for the environmental improvement, or environmental effect, of any site decontamination project. A measure which has been developed by Jacobse & Wolbert (1988) in collaboration with one of the Dutch regions includes the following factors:

- the risk caused by the pollution to public health and the environment, before and after the project is carried out;
- the intended use of the site after decontamination;
- the method used to decontaminate the site (the project variant).

Together, these factors determine the environmental effect of the project. For more detail, see the Appendix. The final aim of any project is to change a site under consideration from polluted to multifunctional, i.e. available for multiple and different uses.

It should be clear that no quantitative measure of the environmental effect to society for each project variant at each contaminated site will be wholly uncontroversial. Note, however, that the issue is not so much that the environmental effect measure be very precise, but rather that all regions agree, for budget allocation purposes, on an environmental effect measurement system which is reasonably accurate and easy to determine and verify. Whatever method is used, the regions have to perceive it as fair, and not favouring any region over the others. For our procedure, the only essential requirement is that it can be treated as an additive utility function.

The problem will be formulated as a mathematical programming problem, and as such is typical of an operational research/management science (OR/MS) approach. Rather than focus on structural properties of solutions this paper is more concerned with how policy-makers can actually find and implement solutions to a given problem. In fact, the OR literature has shown a growing degree of interest in environmental problems, as illustrated recently in this journal by Van Beek, Fortuin & Van Wassenhove (1992). The June 1992 issue of *OR/MS Today* focused on environmental issues, and a special environmental issue of the *European Journal of Operational Research* is in preparation.

3 The mathematical model

The problem described above can be formulated as a mathematical programming problem. In doing so, we should realise that it is intended as a decision support tool for a central government's environmental policy, combining central efficiency with the decentralisation principle. Thus, the model should suggest a budget allocation and the corresponding environmental effect levels which each region can be expected to achieve with that allocation; the actual selection of individual projects is a regional matter. However, a procedure suggesting how to select projects is included to enhance consistency and comparability between regions.

The situation described before includes two different decision variables (the index r denotes regions, p projects, v variants, and t time periods):

b_{rt} : budget allocated to region r in period t

x_{rpvt} = 1 if region r carries out variant v of project p in period t

0 otherwise

Furthermore, the parameters used are:

B_t : total budget available in period t

E_{rpvt} : environmental effect of project variants

C_{rpvt} : costs of project variants

WG_{rpvt} : waste generated by project variants

LR_{rpvt} : labour requirement for project variants

WS_{rt} : regional waste storage capacity

LS_{rt} : regional labour supply

Using this notation the allocation problem can be formulated as follows:

$$(BAP) \quad \max \quad \sum_r \sum_p \sum_v \sum_t E_{rpvt} x_{rpvt} \tag{1}$$

$$\text{subject to} \quad \sum_r b_{rt} \leq B_t \qquad \forall\, t \tag{2}$$

$$\sum_p \sum_v C_{rpvt} x_{rpvt} \leq b_{rt} \qquad \forall\, r, t \tag{3}$$

$$\sum_p \sum_v WG_{rpvt} x_{rpvt} \leq WS_{rt} \qquad \forall\, r,t \qquad (4)$$

$$\sum_p \sum_v LR_{rpvt} x_{rpvt} \leq LS_{rt} \qquad \forall\, r,t \qquad (5)$$

$$\sum_v \sum_t x_{rpvt} \leq 1 \qquad \forall\, r,p \qquad (6)$$

$$x_{rpvt} \in \{0,1\} \qquad \forall\, r,p,v,t \qquad (7)$$

The period index t has been included to emphasise that the problem is a multi-year one; it would take many years indeed to decontaminate all sites. In the rest of the paper, however, we will treat the problem as a sequence of one-year allocation problems, rather than attempt to optimise over the entire horizon simultaneously. This also reflects political reality more closely, as it is often difficult for governments to commit to future budget levels, particularly when elections will take place in the meantime.

In (BAP) the total environmental effect is maximised. Constraints (2) state that the total available central budget cannot be exceeded. Constraints (3) define the budget constraints for the regions, and constraints (4) and (5) are the regional waste storage capacity and labour supply constraints, respectively. Constraints (6) ensure that each site is decontaminated at most once.

We assume that each project can be carried out within one year. Labour supply and storage capacity are given for each region and for each year; they cannot currently be exchanged between regions. This observation that, for any given budget allocation, the regions are independent, allows us to make use of the hierarchical structure of the model. It is particularly likely to be relevant in an international situation.[4]

4 Solving the problem

In the Dutch situation, there are 16 regions, some 100 potential projects per region, on average 3 variants per project, and a 5 year planning horizon. This is a large mixed integer

[4]If a limited amount of exchange between regions would be permitted in practice, the model used here would provide a lower bound on achievable environmental effect.

programming problem, which in general cannot be solved to optimality within a reasonable amount of computer time and memory. In practice, the problem is solved on a year-by-year rolling horizon basis. Before a new planning round for year t, projects completed in year $t-1$ are deleted from the database, and new contaminated sites are added. The problem is then solved for year t. Note that the integer variables x_{rpvt}, which make the model hard to solve, are not of primary importance at the central level. In decentralised decision making, the regions are free to select specific projects, while central government is only concerned with allocating regional budgets b_{rt} so that the total environmental effect over all regions is maximised.

Any procedure to tackle this problem should be evaluated on a number of general and technical criteria:

1. The objective function value must be shown to be good. Although the global optimum will generally be too hard to find, the procedure must be demonstrated to provide solutions which are acceptable in comparison with the global optimum.
2. The target environmental effects corresponding to the resulting allocation must be attainable by the regions. If, for example, the linear programming relaxation of (BAP), where the x_{rpvt} are not longer constrained to be 0 or 1, would be used to determine b_{rt} and the corresponding regional targets, it will generally not be the case that a set of integer x_{rpvt}, allowing the regions to reach their environmental effect targets, exist. The central authority would then be setting unreachable targets, which is clearly undesirable.
3. The information flow required between regional and central governments should be kept to a minimum, given the highly imperfect nature of current information exchange procedures.
4. The solution procedure should be implementable on a small computer or PC, i.e. it should be efficient in terms of computer time and memory. This is to allow scenario analysis, e.g. to study the effects of changing the total budget level B_t.

Some of the more obvious approaches to the problem are ruled out by these requirements. For the linear programming relaxation, it can be shown that an unacceptably large proportion of the x_{rpvt} will be fractional, which conflicts with requirement 2. Additionally, to keep information requirements low, we would like to have a once-only information flow. Well-known iterative price-directed or resource-directed decomposition procedures (see e.g. Kornai & Liptak (1965) or Ten Kate (1972)), are therefore also instantly ruled out.

To solve this problem, we propose the following two-stage procedure. In the first stage, the regions specify the maximum environmental effect they can achieve for different levels of budget, different numbers of 'budget brackets', that might be allocated to them. In the second stage, the central authority determines which budget allocation leads to the highest total environmental effect, given the information specified by the regions in the first stage. More precisely, the procedure is as follows:

Stage 1

The regions make an inventory of their projects and resources. Then, the regions construct their effect functions, indicating what environmental effect they expect to achieve for any given number of budget brackets.

Stage 2

Given the information from stage 1, the central authority divides the total available budget, setting the regional budgets and optimising the total expected environmental effect.

Execution

Given the allocation from stage 2, the regions select and carry out some of their projects.

Feedback

The central authority measures each region's performance according to an evaluation measure as discussed in Section 5, for use in future budget allocation decisions.

The procedure is schematically represented in Figure 1. In Section 5 we discuss how it lives up to the above requirements.

FIGURE 1 ABOUT HERE.

We next describe stages 1 and 2 in some more detail. Assume there are R regions, and that the total budget is divided into K brackets of size L. In choosing the value of K, the increased accuracy provided by smaller brackets must be traded off against increased computational requirements. The following description presupposes some knowledge of mathematical programming; it can be skipped without loss of continuity. To illustrate the procedure, a simple numerical example is included, with $R = 3$ regions and $K = 10$ budget brackets of size $L = 1$. For the data used, see Table 1. In the example, we assume that, for political reasons, each region must receive a budget of at least 2. The procedure now consists of two stages:

TABLE 1 ABOUT HERE.

1. For each possible budget allocation ($b = 0, L, 2L, \ldots, KL$), each region specifies the maximum environmental effect it can achieve, resulting in so-called environmental effect functions. For our example, these are shown in the left-hand side of Figure 2. This involves solving a multi-constraint multiple choice knapsack problem, still a very hard problem to solve optimally. Tests have shown, however, that it is not difficult to find upper and lower bounds for these effect functions. Upper bounds can easily be obtained by solving the linear programming relaxation of the regional problems, lower bounds by using the general integer programming heuristic proposed by Kochenberger et al. (1974). The output of this step is a matrix (A) which contains as elements $A(r, k)$ the maximum effect attainable by region r with a budget of kL.

FIGURE 2 ABOUT HERE.

2. Central government now has a matrix containing the attainable environmental effect for each region and for each budget allocation. This can be transformed into an acyclic graph of the form shown on the right in Figure 2. The vertices are the different cumulative budget levels for the regions, and the edge $(\{r_1, k_1\}, \{r_1 + 1, k_2\})$ has a value that is equal to the maximum attainable environmental effect if region $r_1 + 1$ receives $k_2 - k_1$ brackets of money, i.e. a budget of $(k_2 - k_1)L$. The objective is now to find a path through the graph which maximises total cumulative environmental effect. The corresponding

budget allocation can then be derived by backtracking through the graph via the optimal path. The optimal path can be found by using the following dynamic programming recursion formula:

$$F(r+1,k) = \max_{i=0,..,k} \{F(r,i) + A(r+1,k-i)\} \quad k = 0,..,K \tag{8}$$

where $F(r,k)$ is the optimal environmental effect at vertex (r,k), and $F(0,k) = 0$ for $k = 0,..,K$. We are looking for $F(R,K)$. The interpretation of the recursion is the following: when k budget brackets are to be allocated and a new region is included in the allocation decision, the procedure determines how many of the k brackets should be allocated to the regions already considered previously and how many to allocate to the new region. The complexity of this recursion is $O(RK^2)$. Note, however, that many budget allocations can be ruled out a priori on political grounds, since regional budgets will undoubtedly have lower and upper bounds for a number of reasons. This reduces the number of paths to be evaluated. This explains why some edges have not been included in the graph in Figure 2.

The main advantage of this two-stage approach is its minimum level of information flow. Nearly all the work is done at the regional level; finding the optimal path in step two is not demanding at all, neither practically nor technically. If we extend the procedure to a longer planning horizon, it cannot be used in this form. The effect functions for later years depend on the budget allocation decisions for previous years, which would drastically expand the state space of the graph in Figure 2. What we can do, however, is use the one year procedure on a year-to-year basis. In that way, the 'best' projects will always be chosen first, which is not necessarily optimal, but will certainly correspond with real-life short-term interests. It will also be very easy to update the project information each year, by scrapping the projects that have been carried out, including new projects, and adjusting project characteristics (e.g. urgency). In any year, stages 1 and 2 of the procedure can be repeated for several consecutive periods, to obtain a first estimate of future budget allocations.

5 Practical aspects of the procedure

In this section, we evaluate the procedure in terms of the general and technical criteria mentioned earlier, in Section 4. First of all the procedure has a very straightforward structure and is therefore easy to carry out. The information requirements are minimal: each region must communicate its regional environmental effect function to the central authority, which only communicates the final budget allocation (implying expected environmental effect) to the regions. There is also a high degree of decentralisation, because the central authority does not need to know anything about the actual projects (specification nor selection). Requirement 3 is therefore satisfied.

Extensive computational experiments, modelled after the Dutch situation, have shown that the procedure also performs well in terms of computation speed and closeness to the optimal solution. On randomly generated problems, with up to 15 regions, 100 potential projects per region, and 3 variants per project, the procedure came within 5% of the optimal solution and generally much closer still. Given the level of uncertainty that will be inherent in any environmental effect measure, a deviation of under 5% seems fully acceptable. In all cases, the solution was found in at most some 30 seconds of CPU time on a VAX 8650 mainframe computer. For a long-term decision as the budget allocation one, this is a quick response indeed. (Details on the experiments are given in Debets (1992).) A result of these technical characteristics is that the procedure facilitates evaluating and comparing different scenarios. The experiments clearly demonstrate that the procedure easily meets requirements 1 and 4.

Moreover, the central authority sets targets and allocates budgets in a mutually consistent fashion, ensuring that each region can meet the target with the budget allocated to them. This is because budgets and targets are both based on the environmental effect functions reported by the regions themselves. This shows that requirement 2 is met.

6 Incentive compatibility

An extremely important issue for successful implementation of this budget allocation procedure, which has not yet been addressed, is that of incentive compatibility. In this section we provide an informal discussion of the extent to which incentive compatibility problems can arise as a result of the decentralisation of responsability and information in our procedure. The procedure relies heavily on the assumption that the regions construct their environmental effect functions based on optimisation of the environmental effect measure specified by the central authority, and that they then report these environmental effect functions truthfully. It should be clear that it is not necessarily in the regions' best interests to do so, unless the central authority appends an appropriate incentive mechanism to the budget allocation procedure. Let us assume that the central authority is responsible for evaluating the performance of the managers of each of the regional authorities. Promotion, salary, bonus, job security etc. of the manager of region r are decided by the central authority, on the basis of a quantitative performance evaluation measure P_r.[5] There are two ways in which a moral hazards problem can arise:

1. A regional manager can compute his environmental effect function A_r correctly, but deliberately report a different function $\hat{A}_r$, e.g. with a view to obtaining a larger budget allocation.
2. When computing his environmental effect function, a regional manager can accord particular projects a higher priority than they deserve under the central authority's environmental effect measure, e.g. if some project has a high local political visibility but in fact yields little environmental improvement at a high cost.

We now address each of these cases in turn.

In the first case, each regional manager acts so as to maximise his performance evaluation measure P_r. To induce truthful reporting of effect functions by regional managers, we can use the scheme proposed in Groves (1976) or Groves and Loeb (1979). Let A_r be the true

[5] For notational convenience, we omit the period index t.

environmental effect function of region r, and $\hat{A}_r$ the reported effect function, with A and $\hat{A}$ denoting the vectors of all A_r and $\hat{A}_r$ respectively. So, $A_r(b_r)$ solves region r's subproblem if they receive a budget level b_r; let b denote the vector of allocations b_r. The central authority now needs to find an evaluation measure $P_r(A_r, \hat{A})$ such that when region r solves

$$\max_{\hat{A}_r} \qquad P_r(A_r, \hat{A}) \tag{9}$$

and the central authority computes $b^*(\hat{A})$ by solving

$$\max_{b=(b_1,\ldots,b_R)} \qquad \sum_{r=1}^{R} \hat{A}_r(b_r) \tag{10}$$

$$\text{subject to} \qquad \sum_{r=1}^{R} b_r \leq B \tag{11}$$

the central authority's original objective function (1) is indeed maximised. The Groves mechanism, which has the following form, can easily be seen to produce the desired effect:[6]

$$P_r(A_r, \hat{A}) = A_r + \sum_{s \neq r} \hat{A}_s(b_s^*(\hat{A})) \tag{12}$$

Region r's manager's performance measure is not directly influenced by the $\hat{A}_r$ he reports, only indirectly through its effect on the budget allocation b_s^*. If region r reports truthfully, the centre computes b^* so as to maximise $A_r(b^*) + \sum_{s \neq r} \hat{A}_s(b_s^*)$, so region r's performance measure P_r is maximised. A more complete treatment of this mechanism for the resource allocation problem is given in Groves (1976).[7] An important and highly attractive feature of the Groves mechanism is that it yields truth-telling as an equilibrium in dominant strategies, so that one regional manager's beliefs about the other regional manager's true and reported effect functions are irrelevant for his own behaviour.

The second type of cheating is impossible to prevent completely. If a regional manager acts not to maximise his performance evaluation measure P_r but some other utility function unknown

[6] In fact, a class of mechanisms, transformations of the one given here, will yield the same result.

[7] The mechanism as presented here supposes that regions communicate with the centre through complete functions. If, in some context, information exchange between regions and centre were sufficiently well-developed to enable iterative procedures, such as the Dantzig-Wolfe procedure for price-directed decomposition, the mechanism could be modified to work in that context too: see Groves and Loeb (1979), who refer to Cohen (1977).

to the central authority, there is nothing the central authority can do to enforce full compliance from the regional manager. If a regional manager accords a much higher value to performing a particular project than dictated by the centre's environmental effect measure, he can act as follows in computing the environmental effect function. For each hypothetical budget level, he selects projects so as to maximise *his own* environmental objective, but then computes the corresponding environmental effect according to the *central* measure. This way, he will be able to achieve exactly the environmental effect promised by his environmental effect function, but will not be maximising the central objective. However, this course of action will be rendered less attractive by our budget allocation procedure than it would be if there were no formal procedure at all. First, the 'pseudo' environmental effect function constructed by a regional manager on the basis of his own objective need not be an increasing function of budget level, and any region promising less effect with more budget is bound to raise suspicions. Second, the effect promised by a region, and the budget allocated to it as a result, are determined by the relative efficiency, according to the central measure, with which central budget can be deployed in that region. By forcing politically expedient but 'inefficient' projects into the regional selection, that region's 'efficiency' is decreased, which will generally lead to a lower resulting allocation of budget. We see that, although a Groves mechanism cannot prevent this second type of cheating, implementing a formal procedure will reduce the gravity of the problem.

The above mechanism is based on the assumptions that project information is known to the regions and the environmental effect measurable. Several problems arise if these assumptions do not hold. Most important is the environmental effect: there has to be a generally accepted system of specifying the environmental impact of the projects. If that is not agreed upon, there is no basis to compare performances at all, in whatever control mechanism. Given an effect measure, the regions have to be able to estimate the outcome of any particular project variant. Although this is not an easy task, steps have already been made in this direction. For example, a computer model has been developed to help users estimate contaminant concentrations, based on information provided by the user about the site and its intended future use (Daniel, 1989). Another problem arises if it is not possible to measure the actually achieved effect for all projects carried out, in the Dutch case some 150 each year. There has to be a way to compute the evaluation measure, as otherwise the control mechanism breaks down. Failing that, though, one can hardly expect any logical budget allocation and evaluation mechanism to exist. Fortunately, technology does already exist for site monitoring: fiber optic chemical sensors allow in situ real-time readings of various pollutants.[8]

However, for all these measures it is not so much important that they be absolutely accurate, but more that they be consistent. In that case, the evaluation mechanism is fair to the regions: if a region does not act according to the overall objective, it will be punished and the others will profit. In real life, fairness is often one of the best incentives one can think of, and that can be a first step towards agreement on the issues just mentioned. The main difficulty seems to lie in the performance evaluation. If we do not want to give up decentralization, this depends crucially on cooperation from the regions: there is no possibility of checking the environmental situation before and after each project, although there is a possibility of imposing good documentation requirements, so that specifications can be checked on a sampling basis if necessary. The report by Beumer et al. (1991) shows that most of the specification is done to central authority's standards, although it is not always very well documented. It is recommended to improve record keeping systems so as to decrease sensitivity to fraud. It should be noted that the whole planning scheme also depends quite heavily on the environmental effect measure, a measure which is not at all uniquely defined. However, it is not essential that the measure is absolutely accurate, but more that it is consistent between regions and the central authority, so that it is perceived to be fair as a basis for comparison. We can conclude that the main strength as well as the main weakness of the procedure is the fact that nearly all responsibilities lie with the regional, not the central authorities.

7 Conclusions

We have argued that there is a need for an objective, formal procedure to support and implement budget allocation policies for site decontamination projects. In this paper we have

[8] *Environment Today*, Vol. 4, No. 2, February 1993, p. 8.

discussed the requirements such a procedure should meet, and outlined a procedure based on an operational research approach that meets these requirements. The problem of allocating budgets to regions to optimise the total national environmental effect of soil decontamination projects is a difficult one, and an OR approach cannot, by itself, give the optimal solution to a politically sensitive problem. It can, however, provide an objective basis for decision making and for negotiation. The two-stage procedure proposed in this paper, where central authorities divide the budget on the basis of summary regional information concerning achievable environmental effects given different budget levels, satisfies political as well as technical requirements that might be imposed: it facilitates decentralisation, minimises information requirements, it is fast and efficient, produces good and robust results, and it is transparent. It is also equitable, as it facilitates control over the spending of the budgets: a feedback procedure, incorporating penalties or bonuses for (not) managing to reach target environmental effects, can be used. If nothing else, it can certainly result in a more efficient budget allocation than any manual method, without unrealistic information or computer requirements. Besides, using the procedure proposed in this paper will result in a better insight in and understanding of the problem, especially when it is embedded in a decision support system that allows for easy scenario analysis.

Extensions to the problem considered here are to allow pooling of resources between the regions, cooperation between regions in 'similar' or geographically overlapping projects, rolling horizon planning, allowing projects to take fractions of time periods, and further incorporation of dynamic allocation policies with information feedback and incentive schemes for meeting specifications. More theoretical work can be done on the mechanism design issues related to the procedure. The impact of issues such as reputation effects and the ratchet effect (see e.g. Laffont and Tirole 1993) can then be examined in this context.

The value of a mathematical approach to environmental problems has already been shown elsewhere, and is also demonstrated by the problem discussed in this paper. We believe that there are many more areas where OR can be of similar aid, making extensive further research along these lines a promising enterprise as well as a social necessity.

Acknowledgements

We are grateful to Paul de Bijl for his helpful comments on an earlier draft of this paper.

References

Beumer, P.A.M., L.K. Slager, A.A.A. van der Schraaf, J.C.M. van Eindhoven (1991), "Landelijk onderzoek bodemsanering budgetgevallen", Report no. W90039, Vakgroep Natuurwetenschap en Samenleving, University of Utrecht (NL) (in Dutch).

Cohen, S. (1977), "Incentive compatible control of the multidivisional firm with iterative communication", unpublished Ph.D. thesis, Northwestern University.

Daniel, D. (1989), "AI Decides Safety of Contaminated Sites", *Computing Canada,* **15**(July 20), 1, 4.

Debets, F.J.C. (1992), "Budget allocation for decontamination projects: a mathematical programming approach", unpublished Master's Thesis, Econometric Institute, Erasmus University Rotterdam (NL).

Green, J.R. and J.-J. Laffont (1979), *Incentives in public decision-making,* North-Holland, Amsterdam.

Groves, T. (1976), "Incentive compatible control of decentralized organisations", in Y. Ho & S. Mitters, eds., *Directions in large-scale systems: many-person optimization and decentralized control,* New York, Plenum.

Groves, T. and M. Loeb (1979), "Incentives in a Divisionalized Firm", *Management Science* **25**(3), 221-230.

Jacobse, A.J. & P.P.G. Wolbert, (1988), "Saneringsproject en saneringstechniek: een keuzeprobleem" (in Dutch), Unpublished Master's Thesis, Agricultural University, Wageningen (NL).

Johnson, E. (1989), "Answers Surface for Europe's Underground Woes", Chemical Engineering, **96**(6), 47-58.

Kochenberger, G.A., B.A. McCarl, and F.P. Wyman (1974), "A heuristic for general integer programming", *Decision Sciences*, 5(1), 36-44.

Kornai, J. and T. Liptak (1965), "Two-level planning", *Econometrica*, **33**(1), 141-169.

Laffont, J.-J. and J. Tirole (1993), *A Theory of Incentives in Procurement and Regulation*, The MIT Press, Cambridge MA.

Ten Kate, A. (1972), "Decomposition of linear programs by direct distribution", *Econometrica*, **40**(5), 883-898.

Van Beek, P., L. Fortuin and L.N. Van Wassenhove (1992), "Operational Research and the Environment", *Environmental and Resource Economics*, **2**, 635-639.

Vellekoop, M. (1994), "Via decentralisatie naar krachtiger bestrijding van industrielawaai", *Ruimtelijke Ordening en Milieu*, May, 9-11 (in Dutch).

Appendix: an environmental effect measure

A measurement system for environmental effect has been developed at the Agricultural University of Wageningen (NL), by Jacobse and Wolbert (1988), in collaboration with one of the Dutch regions. As mentioned before, three factors determine the environmental effect of decontaminating a site using a specified method:

- Risk scores are determined for each combination of site and decontamination method. Both prior and estimated posterior scores, before and after project execution, are determined, R_{before} and R_{after} respectively. The scores are based on degree of contamination of soil and ground water, surface and volume of contamination, the rate at which the contamination spreads horizontally and vertically, the frequency with which people are present at the site, the extent to which exposition to the contamination is possible (through air, water, or direct), and the presence of farm animals and plants and their exposition to the contamination. The resulting total scores lie within the interval $[0, 100]$.
- A score L in $[1, 10]$ for the site's location is determined, where e.g. 10 indicates that the site will be used for housing, 6 indicates offices etc., and 2 indicates a transportation function.

- The "applicability factor" T, in $[0, 10]$, indicates how suitable that particular decontamination method is for the project concerned. Criteria included here are among others reliability of the method, technical feasibility, and the extent to which the method has a permanent effect.

The relative importance of the subcriteria can be included by using various weighting schemes. The total environmental effect, E, of a proposed project, is now given by

$$E = L(R_{\text{before}} - R_{\text{after}})\frac{20 + T}{20}$$

so that $E \in [0, 1500]$.

Table 1: Numerical example

project	region 1 costs	effect	LR_{1p}	WG_{1p}	region 2 costs	effect	LR_{2p}	WG_{2p}	region 3 costs	effect	LR_{3p}	WG_{3p}
1	1	20	2	3	1	25	5	5	1	15	5	2
2	1	15	3	2	2	20	3	2	1	15	3	5
3	2	20	2	2	2	10	2	2	1	10	2	2
4	2	15	1	2	3	40	6	6	2	25	5	5
5	3	30	3	4	3	35	4	3	2	20	4	4
6	4	30	2	3	3	30	3	3	2	15	3	3
	LS_1:	10			LS_2:	10			LS_3:	10		
	WS_1:	10			WS_2:	10			WS_3:	10		

Figure 1: Schematic representation of the two-stage procedure

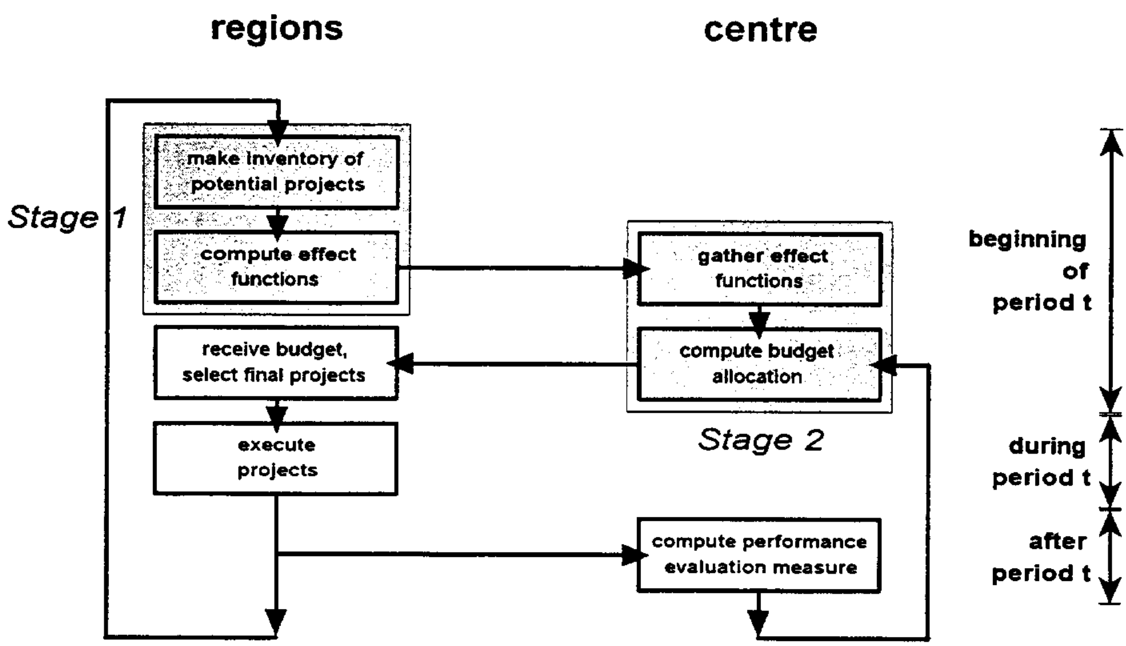

Figure 2: Regional effect functions and central budget allocation procedure

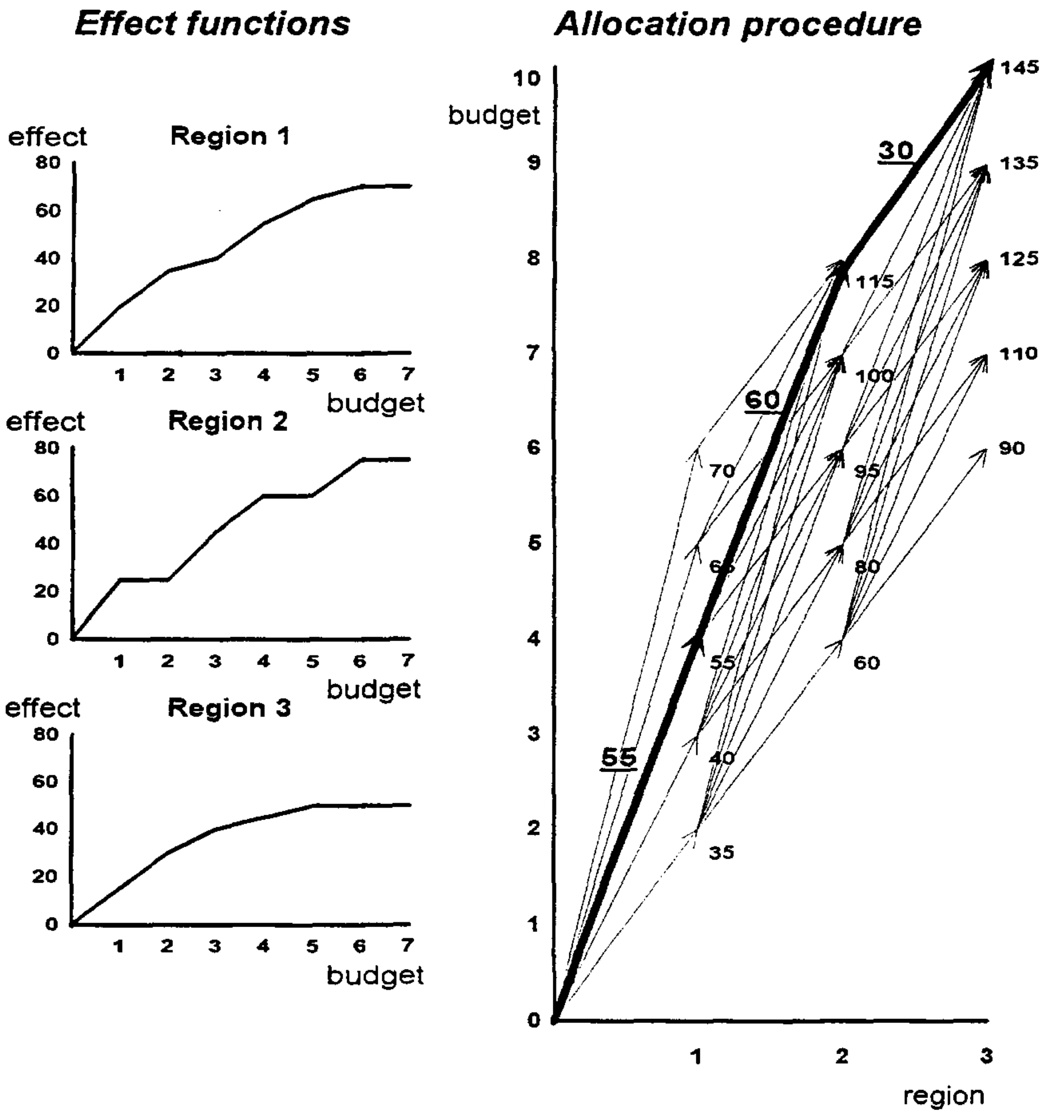

Note: The effect functions and the allocation procedure are explained in more detail in the text. The figures at the vertices (r, k) of the graph to the right are the maximum obtainable effects when allocating a budget of k brackets to the first r regions; the underlined figures beside the bold arrows correspond to the environmental effect achieved by each region in the optimal solution.