"A PROCEDURE FOR EFFICIENT BUDGET ALLOCATION TO SITE DECONTAMINATION PROJECTS"

by

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

93/66/TM

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

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

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

Printed at INSEAD, Fontainebleau, France

A Procedure for Efficient Budget Allocation to Site Decontamination Projects

Charles Corbett, Frank Debets,
Luk Van Wassenhove

INSEAD, Fontainebleau, France

September 15, 1993

Abstract

Part of the environmental policy of a central government concerns the decontamination of polluted soil. Since the number of polluted sites is usually quite substantial, these decontamination projects have to be monitored regionally. The problem considered in this paper is that of allocating a central budget to the regions so as to obtain the "best" overall environmental effect. One of the most important side constraints is that the project information flow from regions to centre should be as small as possible. A model is formulated from which an upper bound to the actual achievable environmental effect can easily be obtained. This serves as a benchmark for the effect actually attainable with any budget allocation proposed. To find such an allocation, a procedure is proposed where the regions specify their achievable effect for every possible level of budget allocated. Using this as a basis, the central authority determines the most preferable budget allocation. Both steps in this two-step procedure are based on mathematical decision processes. Some computational results are presented for the situation in the Netherlands. Finally some attention is paid to an incentive scheme that urges regions to specify true information and carry out overall desirable operational programs, consistent with central and other regions' standards.

1 Introduction

Economics is often defined as the science of scarce resources; operational research (OR) can be thought of as a technology enabling optimal allocation of these scarce resources. OR originated during World War II, when it was concerned with the urgent problems of deploying scarce military resources efficiently and effectively. Now that the threats confronting our environment are generally considered to be among the most serious problems faced today, it is only natural to employ the techniques of OR to tackle these problems. For an overview of environmental applications of OR, see e.g. Van Beek, Fortuin & Van Wassenhove (1991) and Debets & Van Wassenhove (1992).[1]

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) environmental benefit. The problem is similar in spirit to that studied in the RAINS project (see, e.g., Hordijk (1988)), where simulation was used to evaluate various budget allocation policies to combat acidification in Europe. An important factor in these allocation decisions is the so-called subsidiarity principle, delegating executive responsibility for selection and actual cleaning of sites to the regional authorities.

Clearly, if the budget allocation is decided on purely political grounds, it is unlikely to bring maximum environmental benefit. By using a mathematical approach, decisions can be made in a more rational manner. It should be emphasized, 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 problem is modelled after the Dutch situation, but the approach outlined applies to any similar hierarchical budget allocation problem, national or international. New cases of polluted sites are continuously detected throughout the country, and coordinated action is urgently required. For this purpose, the Netherlands is divided into 16 regions.[2] 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. The costs and environmental effects obviously vary strongly between

[1]INSEAD has recently designated environmental issues as a key area for research and teaching, by setting up the Management of Environmental Resources Programme. A range of projects concerning various forms of interaction between managerial and environmental issues has been and is being undertaken within this framework.

[2]The Netherlands is divided into 12 provinces; for this particular problem the four big cities (Amsterdam, Rotterdam, The Hague and Utrecht) have also been given the status of "region".

these approaches, which also have different requirements for 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 effect.

It should be clear that the information requirements to characterise all potential site decontamination projects are huge. An entirely centralized approach to this problem would require all this information to be centrally available, which would be highly impractical. Moreover, this option would often be politically undesirable or even infeasible. Two major concerns in developing a hierarchical budget allocation procedure for such a situation are therefore:

- *Decentralization.* 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 effect it expects the regions to achieve with their budgets.
- *Minimum information flow.* The information flow in both directions should be minimized, without jeopardising the efficiency of the final allocation.

The bottlenecks in the current situation are as follows:

- The limited budget allows only a small proportion of projects to be carried out.
- Many projects require a specialist but scarce workforce.
- There is insufficient capacity to store polluted soil.
- The information exchange at all levels is still in a preliminary stage, and therefore far from perfect.

We obviously need a quantitative measure for the environmental effect of any site decontamination project. One measure which has been proposed includes the following factors:

- The risk caused by the pollution to the public health and the environment, before and after the project is carried out.
- The intended use of the site (equivalent to the fragility of the area).
- The method used to decontaminate the site (henceforth called the project variant).
- The urgency of the project.

These factors are combined and result in an environmental score for each variant of the project.

A more elaborate discussion on how the environmental effect is "determined" can be found in Deijmann (1991) or in Jacobse & Wolbert (1988). It should be clear that quantifying the environmental effect to society for each project variant at each contaminated site is a subject for debate. Note however that the issue is not so much that the environmental effect measure be precise (which is impossible for many reasons), but rather that all regions agree on a reasonably accurate and easy to determine environmental effect measurement system for budget allocation purposes.

The eventual aim is to change a site under consideration from polluted to multifunctional, i.e. available for multiple and different uses. A recent study of regional approaches to the problem has shown little consistency in decontamination policies. The result is that only a very small percentage of projects is completed according to central standards and that the current budget allocation is far from optimal. In a recent newspaper article[3] the Dutch Ministry of Environment was criticized for using the available money injudiciously, and it was argued that the procedure for controlling the decontamination program is not satisfactory. A new allocation system is said to be put into operation soon.

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 subordinate actions, either implicitly or explicitly. In such a case, an OR approach can be a natural support.

The remainder of this paper is organised as follows. In Section 2, the problem is formulated as a mixed integer linear programming model. The resulting model is far too large to be solved to optimality; in Section 3 we propose a two-step heuristic procedure to obtain good allocations. Section 4 presents some computational results, and in Section 5 we evaluate the strength of the procedure and address some issues related to implementatiom. The final Section contains some conclusions and areas for further research.

[3] Appeared in the Dutch newspaper *Trouw*, 29/1/1993

2 The mathematical model

The problem described above can be formulated as a mathematical model. In doing so, we should realise that it is 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 a corresponding achievable environmental effect; the 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 (henceforth, the index r denotes regions, p projects, v variants, and t time periods):

b_{rt}	:	budget for region r in period t
x_{rpvt}	=	1 if region r carries out variant v of project p in period t
		0 otherwise

Furthermore, all the parameters to be considered are:

B_t	:	total budget available in period t
E_{rpvt}	:	environmental effect of projects
C_{rpvt}	:	costs of projects
WG_{rpvt}	:	waste generated by projects
LR_{rpvt}	:	labour requirement for projects
WS_{rt}	:	waste storage capacity
LS_{rt}	:	labour supply

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

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

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

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

$$\sum_p \sum_v WG_{rpvt} x_{rpvt} \le WS_{rt} \qquad \forall\, r,t \tag{4}$$

$$\sum_p \sum_v LR_{rpvt} x_{rpvt} \le LS_{rt} \qquad \forall\, r,t \tag{5}$$

$$\sum_v \sum_t x_{rpvt} \le 1 \qquad \forall\, r,p \tag{6}$$

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

In (BAP) the total environmental effect is maximized. 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 waste storage capacity and labour supply constraints, respectively. Constraints (6) ensure that each project is carried out 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 assumption allows us to make use of the hierarchical structure of the model. It is most likely to hold in an international situation, and means that the performance measures derived from the model will be conservative and can be improved upon if the condition is relaxed.

In the Dutch situation, there are 16 regions, some 100 projects per region, on average 3 variants per project, and a 5 year planning horizon. This is a large mixed integer programming problem, which in general cannot be solved to optimality within a reasonable amount of computer time and memory. Note that the 0-1 variables, which make the model hard to solve, are not of primary importance at the central level. In decentralized decision making, the regions are free to select specific projects, while central government is only concerned with allocating the budget so that the global environmental effect over all regions is maximized. Since the global problem cannot be solved to optimality, we now turn to developing procedures that yield lower and upper bounds to the optimal solution value.

3 Solving the model

Recall that the solution we are looking for is not a full, optimal solution to (BAP), specifying in detail all projects to be executed by the regions, but rather a budget allocation with maximum practically attainable environmental effect. All procedures to arrive at allocations and lower and upper bounds for the optimal attainable effect should be evaluated on a number of general and technical criteria:

1. The objective function value must be shown to be good.
2. The environmental effect corresponding to the resulting allocation must be attainable by the regions.

3. The information flow required between regional and central governments should be kept to a minimum.
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.

Additionally, a procedure that is easy to understand, e.g. by having a straightforward intuitive interpretation, will be more easily accepted. This section considers two ways of obtaining allocation and project proposals. First we tackle the continuous (linear programming) relaxation of (BAP), $(LPBAP)$, to find an upper bound on actually achievable environmental effect. This bound can then be used to evaluate any feasible solution found by a heuristic approach. Later some computational results will be given for tests performed on the Dutch case.

Even $(LPBAP)$, the linear programming relaxation, will be difficult to solve on a small computer or a PC. What is worse is that it will also allow a solution to include partially completed projects. As this is not allowed in practice, the environmental effect specified in the solution will, in general, not be actually attainable. It is known that for the linear relaxation of a knapsack-type problem such as (BAP), every resource constraint may allow one variable to be fractional. In the Dutch case this means that up to 15% of the x_{rpvt} may be fractional. If we try to obtain a feasible solution by simply rounding the fractional variables, the upper bound on the environmental effect achievable cannot be expected to be very tight, nor can the implied operational program be expected to be very efficient or effective. One may try to improve upon the initial solution obtained by rounding the fractional LP output using optimization techniques based on interchanging project choices either within a region or between regions. Different versions of this strategy have been tested in Deijmann (1991). They all leave a gap of about 5% on average between their best environmental value and that of the optimal solution to the LP relaxation. For these heuristics, all project data need to be available at the central level, which is not very useful in a decentralization context.

To improve on this information requirement, we would like to have a once-only information flow. We therefore propose the following two-step procedure. In the first step, the regions specify the maximum environmental effect they can achieve for different levels of budget that might be allocated to them. In the second step, the central authority determines which budget allocation leads to the highest total environmental effect, given the information specified by the regions in the first step. (This can be done using a simple dynamic programming algorithm.)

We next describe this two-step procedure in some more detail. The total budget is divided into K brackets of size L. In choosing the value of K, increased accuracy must be traded off against increased computational requirements. To illustrate the idea we look at only one time period. The procedure now consists of two steps:

1. For each possible budget allocation ($b = 0, L, 2L, \dots, KL$), each region specifies the maximum environmental effect it can achieve, resulting in so-called effect functions (see figure 1). This means solving a multi-constraint multiple choice knapsack problem, a very hard problem to solve. Tests have shown, however, that it is not very difficult to bound these effect functions with upper and lower limits. Upper limits can be obtained by solving the linear relaxation of the regional problems, lower limits by using the general integer programming heuristic proposed by Kochenberger et al. (1974). The output of this step is a matrix $(A)_{R\times(K+1)}$, where R is the number of regions, which contains these maximum effects.

FIGURE 1 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 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. The objective is now to find a path through the graph which maximizes 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 $F(1, k) = A(1, k)$ for $k = 0, .., K$. We are looking for $F(R, K)$. 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. The interpretation of this step is the following: for each new region included in the allocation decision it is determined how much of the total budget should be withdrawn from the regions already considered and be allocated to the new region.

FIGURE 2 ABOUT HERE

The main advantage of this two-step 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 very demanding at all, neither practically nor technically. If we extend the procedure to a longer planning horizon, we cannot use it as before. The effect functions for later years depend on the budget allocation decisions for previous years, which would severely blow up 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 appeal to real-life short-term interests. It will also be very easy to update the project information each year, by scrapping projects carried out, including new projects and adjusting project characteristics (e.g. urgency).

4 Some computational results

(BAP) has been "solved" for different problem sizes using various forms of procedures mentioned above, each for multiple sets of artificial parameters. A summary of the relevant test results is given in this section, accompanied by comments on the performance of the different procedures. A more comprehensive report can be found in Debets (1992). For a quick reference we shall first list the procedures used:

LP : A direct simplex approach to solve ($LPBAP$); yields an upper bound for the maximum achievable environmental effect.

LP/DP : The two step procedure where the linear relaxations of the regional problems are solved (hence providing an upper limit to the effect functions), and the dynamic programming method is used for determining the best budget allocation. This procedure will be the same as LP if the bracket size L equals one monetary unit, otherwise it will yield a lower bound to the LP bound.

KOCH/DP : The same type of procedure as LP/DP, but now using the Kochenberger algorithm to solve the regional problems, hence providing lower bounds to the effect functions. This will yield a lower bound to the achievable effect.

The results are now presented for two different sets of problems. All entries in the tables are averages of a number of differently seeded randomized problem instances. This means that the data used are not from real life situations, but designed to give an equal spread from small to large projects, while all regions are considered to be equally large (which does not structurally affect the performance of the procedures). The computing times (CPU) shown are those in seconds on a VAX 8650 mainframe computer. We will come back to the PC requirement later on.

Table 1 shows the results for different one-year problems. OFV stands for the objective function values obtained by the different procedures. For ease of comparison these values are scaled with respect to the LP solution which is given a value of 100. The regional problems have also been solved to optimality using a branch and bound enumeration technique for the budget suggested by the allocation procedure, in order to evaluate how good this allocation is in reality (in terms of actually achievable effect, denoted AAE). One should keep in mind that the procedures LP and LP/DP provide only budget allocations, i.e. their OFV's in Table 1 do not correspond to actual project selections by the regions as they do for the KOCH/DP procedure.

In general LP/DP overbids by 2.0%, i.e. the actually achievable environmental effect (AAE) corresponding to the budget allocation is about 98% of the upper bound (OFV) resulting from the procedure. KOCH/DP underbids by 0.94%, i.e. the heuristic project selection underlying the allocation procedure (OFV) could be improved upon by almost one percent on average if the regional subproblems could be routinely solved to optimality using a branch and bound procedure (as in AAE). However, computational requirements of doing so are prohibitive. Note that although the LP and LP/DP methods promise a lot more, their actually achievable effect is not significantly higher than that for KOCH/DP. This means that the budget allocation achieved by the (quick)

Table 1: Table of results for 1 year test problems

SIZE	LP	LP/DP		KOCH/DP	
r_p_v	OFV	OFV	AAE	OFV	AAE
5 _30_2	100	100.0	97.2	95.4	96.8
_3	100	100.0	97.7	96.1	97.4
_50_2	100	100.0	97.9	95.4	96.8
_3	100	100.0	97.9	96.7	97.7
_100_3	100	100.0	98.0	97.2	98.2
7 _30_3	100	100.0	98.1	96.8	97.7
_50_2	100	100.0	97.8	96.6	97.7
_70_3	100	100.0	98.2	97.5	98.1
_100_3	100	100.0	98.0	96.8	98.1
10 _30_3	100	100.0	97.7	96.5	97.2
_50_2	100	100.0	97.8	96.3	97.3
_70_3	100	100.0	98.2	97.7	98.1
_100_3	100	100.0	97.9	97.2	98.3
15 _30_3	100	100.0	97.8	96.9	97.8
_50_2	100	100.0	98.0	97.0	97.6
_70_3	100	100.0	98.1	96.9	97.8
_100_3	100	100.0	98.1	97.4	98.2
average	100	100.0	97.9	96.9	97.8

r : # regions, p : # projects, v : # variants
OFV : objective function value (scaled, LP=100)
AAE : actually achievable effect (regional optimum) for given allocation

KOCH/DP procedure is not worse than those from the linear programming procedures. The difference is in the bidding, not in the achievability.

The linear optimality lost by setting bracket sizes larger than one is very small. The LP and LP/DP columns are very similar, and the actually achievable effect with the budget allocation proposed by the two methods is the same for all practical purposes. The bracket sizes have been chosen to achieve an efficient pay-off between objective function value and computing time. Decreasing the bracket size has been shown to result in a smaller under- or overbid, rather than a better budget allocation.

Extending the planning horizon to multiple years, and solving the problem year by year, yields results as shown in Table 2. The big problems for three years as well as five year problems were too large to solve directly via LP relaxation. The LP/DP procedure has not been used anymore as its results do not substantially differ from the LP procedure. The gap between the LP and KOCH/DP procedures increases as we look at more years. However, this effect is smaller if there is more freedom of choice. This means that if new projects are added as the time horizon is rolled forward, the greediness of the KOCH/DP procedure will not much worsen its performance. The actually achievable effect with the resulting budget allocation is no longer computable because of excessive computational requirements in solving all the regional allocation problems to optimality using a branch and bound procedure. However, given the other similarities to the one year case, there is no reason to assume that the budget allocation resulting from the LP procedure will be significantly better than the one resulting from the Kochenberger procedure.

Table 2: Table of results for multiyear test problems

SIZE	1 YEAR		2 YEARS		3 YEARS		5 YEARS
r_p_v	LP	KOCH/DP	LP	KOCH/DP	LP	KOCH/DP	KOCH/DP
5 _30_3	100	96.1	168.2	159.0	213.3	199.5	-
_50_2	100	95.6	171.3	161.6	234.4	221.2	279.3
_70_3	100	96.6	177.7	168.7	236.6	222.1	299.1
_100_3	100	96.9	171.3	163.8	230.5	216.5	300.2
10 _30_3	100	96.5	162.6	152.9	214.1	200.7	-
_50_2	100	96.3	166.2	157.2	220.0	205.9	278.2
_70_3	100	97.7	166.9	158.9	224.0	211.3	290.4
_100_3	100	97.2	172.6	165.8	231.2	217.9	302.5
15 _30_3	100	96.9	166.6	157.0	224.0	209.1	-
_50_2	100	97.0	167.3	158.9	225.4	211.6	281.4
_70_3	100	96.9	169.0	161.9	-	212.3	294.1
_100_3	100	97.4	175.0	167.8	-	224.8	311.5
average	100	96.8	169.6	161.1	225.4	212.7	293.0

r : # regions, p : # projects, v : # variants
Shown are objective function values (scaled, LP (1 year) =100)

Table 3: CPU times for test problems (in seconds)

SIZE	1 YEAR			2 YEARS		3 YEARS	
r_p_v	LP	LP/DP	KOCH/DP	LP	KOCH/DP	LP	KOCH/DP
5 _30_3	5.9	6.3	2.9	12.9	4.1	21.7	4.8
_50_2	7.0	6.9	3.8	15.4	5.4	39.2	7.6
_70_3	14.6	11.6	7.6	31.4	11.8	61.7	15.3
_100_3	22.1	15.5	12.4	50.2	18.2	101.4	24.9
10 _30_3	13.5	11.7	5.5	30.0	8.0	58.3	9.4
_50_2	16.2	13.2	7.0	38.0	10.8	78.5	13.7
_70_3	28.5	16.2	14.7	90.7	24.1	207.3	31.0
_100_3	47.0	28.5	22.2	154.9	38.1	376.9	51.6
15 _30_3	23.3	17.1	8.0	87.2	12.0	466.9	15.0
_50_2	29.3	19.0	10.5	83.9	17.4	307.1	22.4
_70_3	55.5	33.3	21.0	215.7	37.4	-	48.2
_100_3	82.4	44.6	32.0	311.3	56.7	-	77.6

r : # regions, p : # projects, v : # variants

Table 3 shows some computational requirements for the suggested procedures. Computing times seem to be roughly linear in number of regions and number of projects, which is promising for tackling bigger problem instances. Especially the KOCH/DP procedure appears to provide excellent results at a moderate computational price.

5 Implementation

To summarize, we have proposed a procedure consisting of the following steps:

1. The regions make an inventory of their projects and resources.
2. The regions construct their effect function, indicating what environmental effect they expect to achieve for any given number of budget brackets.
3. Given the information from step 2, the central authority divides the total available budget, setting the regional and optimizing the total expected environmental effect.
4. Given the allocation from step 3, the regions select and carry out some of their projects.
5. Optionally steps 1-4 are repeated for several consecutive time periods, so as to give the regions some lookahead facility.

If we evaluate this procedure following the general and technical criteria specified earlier its attractiveness will be clear. First of all the procedure has a very straightforward structure and is therefore easy to carry out. The information requirements are absolutely minimal: strictly needed are one set of regional effect functions and one budget allocation (implying expected environmental effect). There is also a very strong level of decentralization, because the central authority does not need to know anything about the actual projects (specification nor selection). The computational results have shown that a good environmental effect can and in many cases will be attained, although this cannot be absolutely guaranteed since the procedure is heuristic in nature.

Although earlier we have argued that a procedure for allocating the budget should be implementable on a PC, all the results so far have been presented for tests performed on a VAX mainframe. However, one should realize that we were only comparing different approaches, and not really evaluating absolute performance. A PC version of the KOCH/DP procedure is available. Computing times for such a procedure are small, certainly given the fact that the procedure only needs to be used once every time period. The regional problems of constructing the effect functions can be solved within minutes, the exact time very much depending on the number of projects and variants and on the size of the monetary brackets. The central (dynamic programming) problem takes an even smaller amount of time to be solved. Memory requirements are very low: only the effect functions and the set of projects that can be carried out given a certain level of budget allocated need to be stored. Another result of these technical characteristics is that the procedure facilitates running through different scenario's.

An extremely important issue for successful implementation of the two-step procedure which has not been addressed yet is the possibility of regional non-cooperation. There are basically two forms one can think of. First of all regions can misspecify their effect functions on purpose, in order to get more money than they would normally be entitled to. Second, their project selection might be dominated by other interests rather than efficiency and effectiveness (e.g. local popularity). What we need to ensure desirable behaviour is an incentive system that implicitly forces the regions to take overall (near-)optimal decisions. Groves (1976) presents a mechanism for performance evaluation that is particularly useful for the problem discussed in this paper. We need the following definitions:

$\overline{A}$	:	joint reported effect functions
$b(\overline{A})$	:	budget allocation given $\overline{A}$
$A_r(b(\overline{A}))$	:	actually achieved effect for region r given $b(\overline{A})$
$\overline{A}_r(b(\overline{A}))$	:	reported achievable effect for region r given $b(\overline{A})$
E_r	:	evaluation measure for region r

For any region $\hat{r}$, the evaluation measure to use is the following:

$$E_{\hat{r}} = A_{\hat{r}}(b(\overline{A})) + \sum_{r \neq \hat{r}} \overline{A}_r(b(\overline{A})) \tag{9}$$

This measure does not depend on other regions' actions, but only on their reported information. It is used on a relative basis: regions are judged on their performance in relation to that of the other regions. Assuming the regions know this evaluation mechanism, it can be seen that they are best off reporting truthfully and selecting projects accordingly. If only true information is sent, the value of the performance measure will be the same for all regions: $E_r = \sum_r \overline{A}_r(b(\overline{A})) = \sum_r A_r(b(\overline{A}))$, for all r. In case of under- or overestimation there will be a deviation from the average, the size and direction of which depends on the correctness of other regions' information. This deviation can be taken into account in the budget allocation for the following year, assuming the actual allocation takes place before the projects are carried out. This can be incorporated in step 3 of the above procedure. Of course it is still possible to cheat, but this would only lead to a decrease in the region's own (relative) performance. A more elaborate and technical discussion of the mechanism can be found in Groves (1976), although it is not discussed exactly how to adjust budgets given the evaluation measure. This depends on specific circumstances and reasons for deviations encountered, as well as on general policy objectives. The aim is to encourage regions to achieve good decontamination results, not to just evaluate them and take punishing or rewarding actions.

The above mechanism is based on the assumptions that project information is known 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. 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. The same holds for the project information: if one does not know beforehand what the outcomes of selecting different project alternatives are, one can hardly expect a logical budget allocation and evaluation mechanism. 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 cooperate, it will be punished and the others will profit from that. 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.

6 Conclusions

The problem of allocating budgets to regions to optimize 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 a rational basis for decision making and for negotiation. The two-step procedure proposed in this paper (as summarized at the start of the previous section), 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 decentralization, minimizes information requirements, works 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.

Before implementing a procedure within a decision support system, the governmental bodies concerned have to decide upon the eventual choice of solution method and the actual structuring of the information flow. One also has to choose the length of the planning horizon. 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.

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.

References

Atkins, D., (1974), "Managerial decentralization and decomposition in mathematical programming"', Operational Research Quarterly, Vol. 25, No. 4, pp. 615-624

Van Beek, P., L. Fortuin & L.N. Van Wassenhove, (1991), "Operational Research and Environment", INSEAD Working Paper No. 91/02/TM/SM

Debets, F.J.C. & L.N. Van Wassenhove, (1992), "A Guided Tour Through Applications of OR-Techniques to Environmental Problems", INSEAD Working Paper No. 92/24/TM

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

Deijmann, G.J., (1991), "Een budgetallocatiemodel voor bodemsanering" (in Dutch), Unpublished Master's Thesis, Econometric Institute, Erasmus University Rotterdam (NL)

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

Hordijk, L., (1988), "A model approach to acid rain", Environment, Vol. 30, No. 2, pp. 17-42

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

Kornai, J. & T. Liptak, (1965), "Two-level planning", Econometrica, Vol. 33, No. 1, pp. 141-169

Ten Kate, A., (1972), "Decomposition of linear programs by direct distribution", Econometrica, Vol. 40, No. 5, pp. 883-898

Figure 1

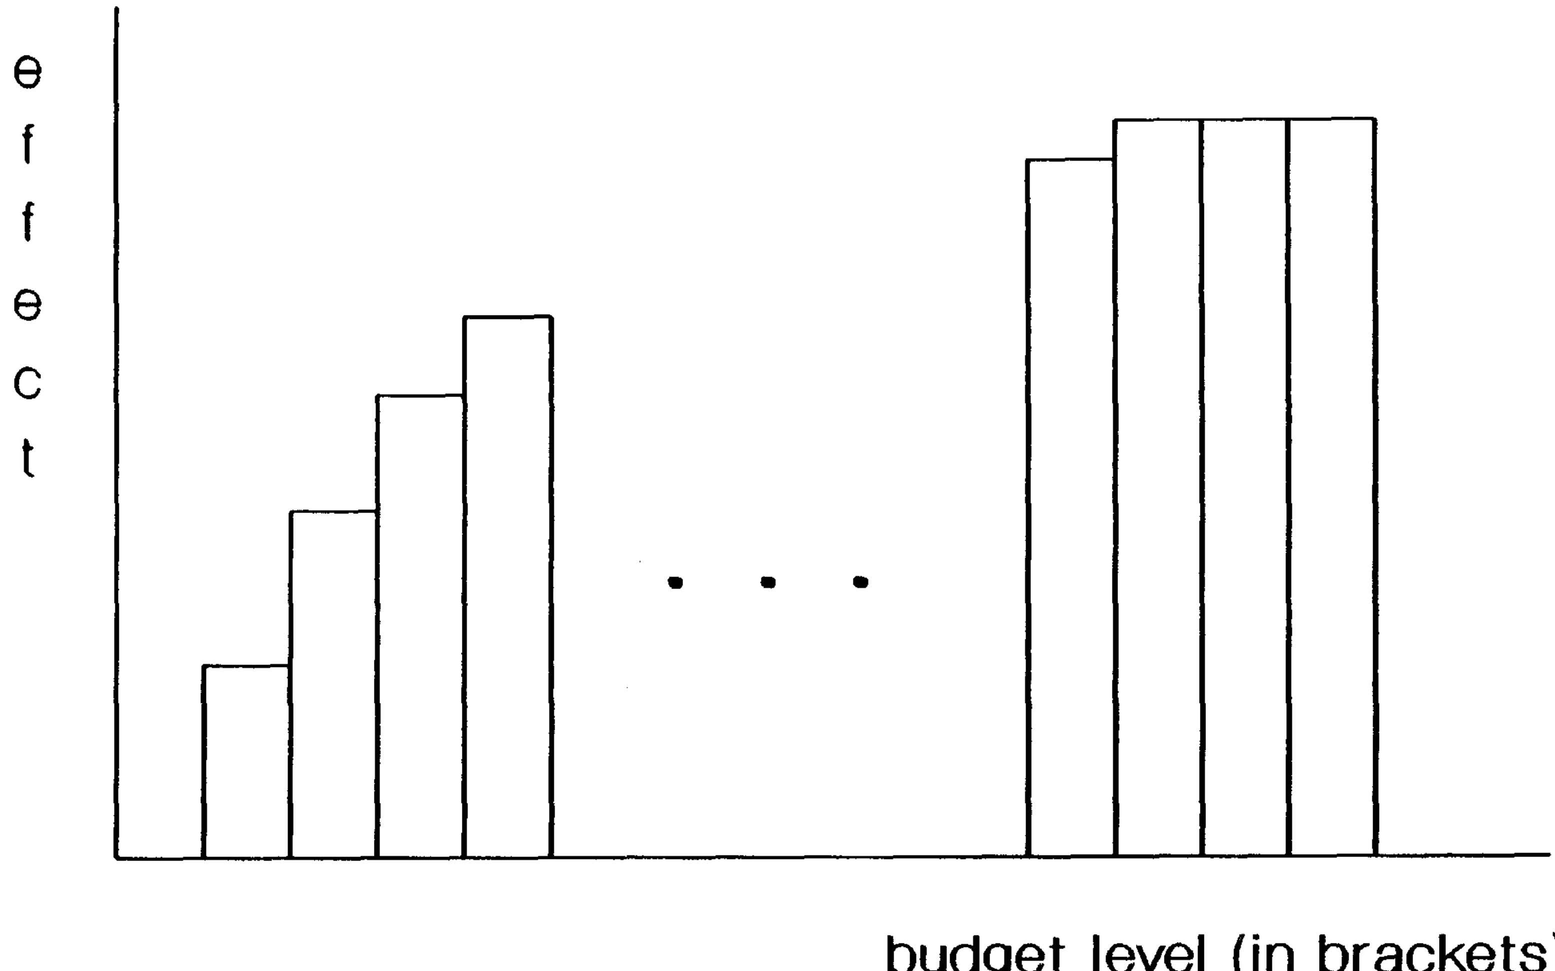

Figure 2

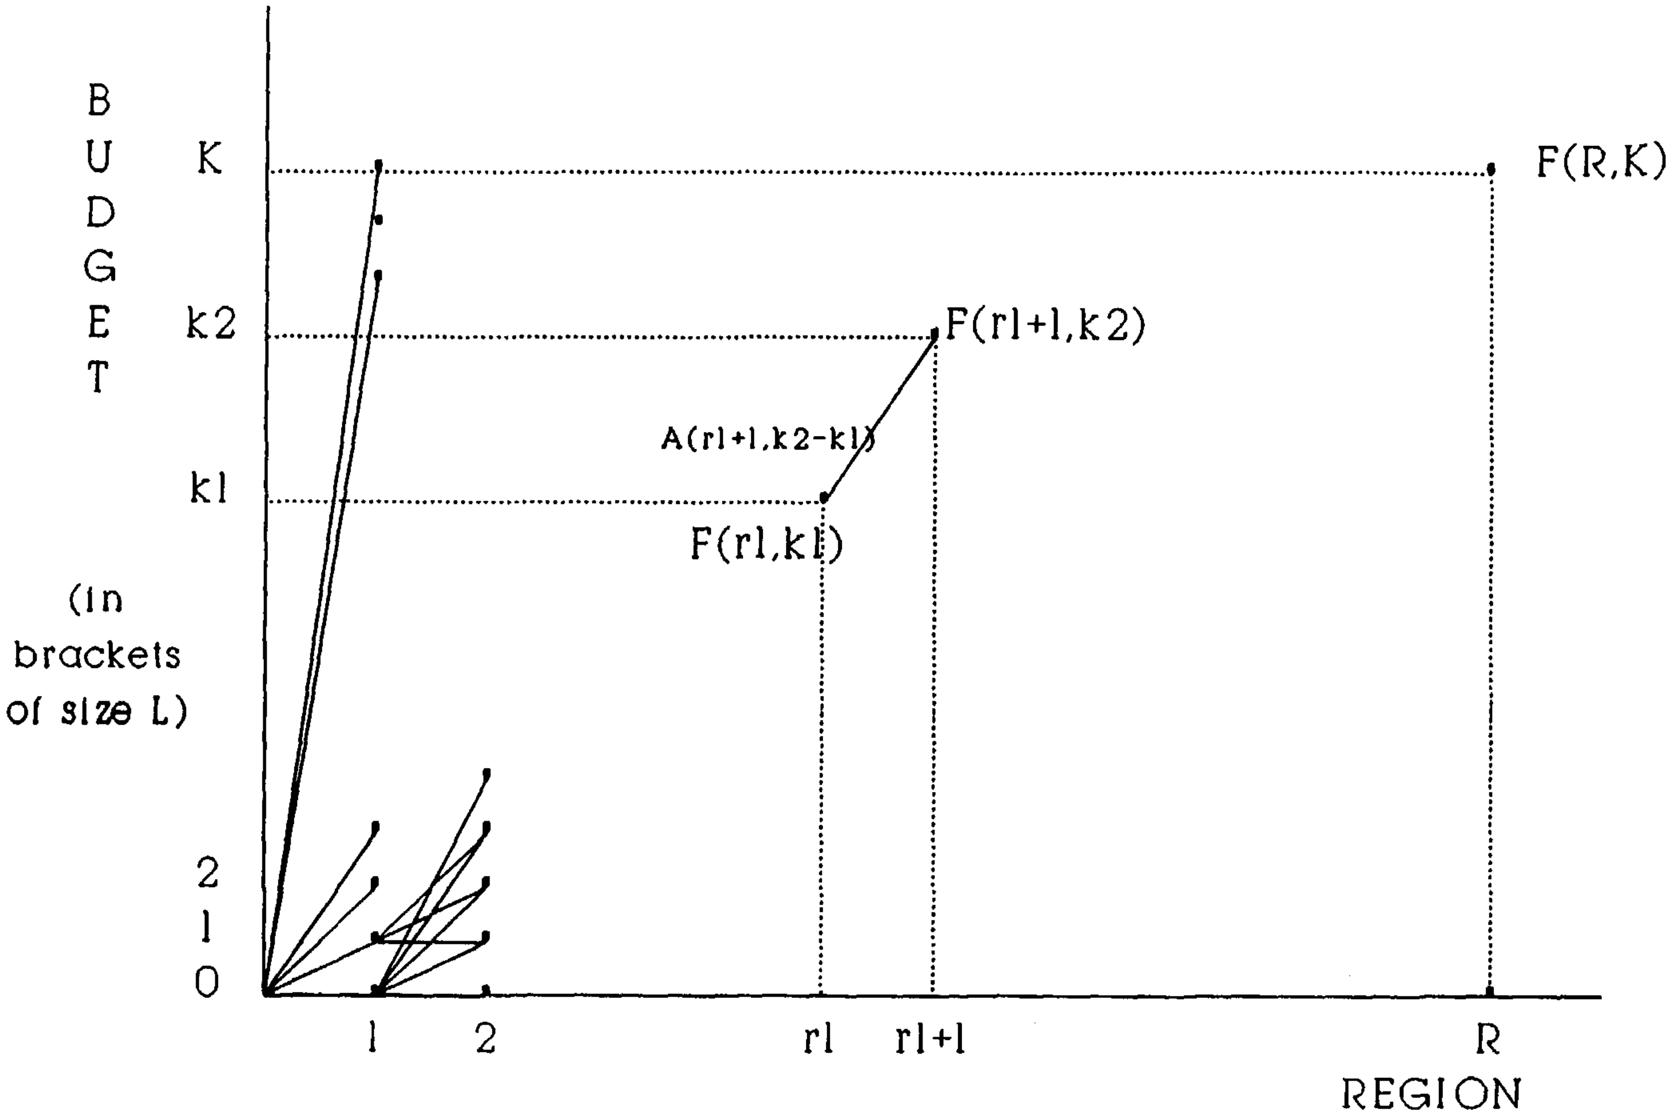
B U D G E T
(in brackets of size L)
K
k2
k1
2
1
0
F(R,K)
F(r1+1,k2)
A(r1+1,k2-k1)
F(r1,k1)
1
2
r1
r1+1
R
REGION