

OPTIMAL RESPONSE TO OIL SPILLS: THE STRATEGIC DECISION CASE

HARILAOS N. PSARAFTIS

Massachusetts Institute of Technology, Cambridge, Massachusetts

GEVERGHESE G. THARAKAN

The World Bank, Washington, D.C.

AVISHAI CEDER

Technion, Israel Institute of Technology, Haifa, Israel

(Received March 1984; revision received December 1984; accepted May 1985)

In this paper we develop a model for the problem of (a) locating appropriate levels and types of cleanup capability to respond to oil spills and (b) allocating such capability among points of high oil spill potential. The model takes into account frequency of spill occurrence, variability of spill volumes, different cleanup technologies, equipment efficiency and operability, fixed costs to open a facility, equipment acquisition, transportation and operating costs, and costs of damage as functions of spill volume and level of response. The model can also accept policy stipulations on response times. We present an illustrative application of the model in the New England region and discuss its possible uses within existing and alternative policy environments.

Oil spill response concerns the emergency action that must be taken to mitigate damages caused by an oil spill. Governments, industry and society in general are concerned about massive spills such as the Torrey Canyon spill off the coast of Britain, in 1967; the Amoco Cadiz spill off the French coast, 1978; the Ixtoc-1 spill in the Gulf of Mexico, 1979; and about smaller spills that occur on a day-to-day basis. Part of the emergency action to mitigate damages concerns the dispatching of specialized cleanup equipment to the spill site in order to contain, recover or disperse the spilled oil.

This paper deals with the *strategic* aspect of the oil spill response problem, that is, with the problem of deciding where to locate adequate capability to respond to potential oil spills. In addition to locational considerations, the strategic oil spill response problem generally calls for decisions concerning the proper levels and types of equipment to be stockpiled, as well as for policies regarding the allocation of such capability among points or zones of high oil spill potential. The purpose of the paper is to formulate a model for the strategic oil spill response problem, present some

illustrative applications, and discuss uses of the model within the existing or alternative policy and regulatory environments.

Strategic oil spill response decisions typically involve planning horizons of considerable duration (e.g., 5–15 years). Since decisions must be made *before* actual spill incidents, the strategic planners must base their decisions, among other things, on *probabilistic* information about the number and volume of spills, as well as on assessments of the potential *consequences* of any particular spill event under a prescribed response. Of course, not all oil spill response decisions are strategic in nature. *Tactical* (or *operational*) decisions deal with aggregate (or detailed) actions that are taken *after* the occurrence of a particular spill, such as what equipment should be sent to the scene, what facility should dispatch that equipment, how long it should stay on the scene and how it should be operated. As Anthony (1965) points out, such a multilevel hierarchical decomposition is desirable (in fact, often necessary) whenever a decision problem possesses a structure that lends itself to natural separation into several decision levels. Decomposition is equally

Subject classification: 165 optimal response to oil spills, 185 siting of oil spill response facilities, 448 optimal response to oil spills.

imperative whenever the problem cannot be formulated or solved by simultaneously capturing all of its facets, from the most aggregate to the most detailed, into a single, monolithic model. The oil spill response decision problem possesses such a structure.

Detailed analytical modeling of tactical and operational decisions is beyond the scope of this paper and can be found elsewhere (see, for instance, Psaraftis and Ziogas 1985, and Ziogas 1982). Still, since such decisions *are* connected with strategic decisions, and for the sake of making this paper self-contained, we shall also include some simplifying assumptions regarding spill response *tactics*, which will be represented in a fairly aggregate (but realistic) fashion. It is also important to stress that the emphasis of this paper is on model formulation, assumption justification and discussion of illustrative examples, rather than on a detailed description of the solution algorithm. All the algorithmic details can be found elsewhere (see, for instance, Tharakan 1982).

The rest of the paper is organized as follows: Section 1 provides some background on this problem to motivate our modeling philosophy. Section 2 models the strategic oil spill response problem and formulates it as a Mixed Integer Programming (MIP) problem. The model, despite its simplicity relative to the real world, accounts for the following problem features: spill frequency of occurrence, variability of spill volumes, different equipment types, fixed costs to open a facility, equipment acquisition, transportation and operating costs, equipment efficiency and operability as a function of weather conditions, and damage costs as functions of spill volume and level of response; the model can also accept policy stipulations on response times. To our knowledge, no other model in the OR/MS literature to date has attempted to integrate all of these features of this problem and explicitly evaluate trade-offs between system and damage costs, as they affect locational decisions. Section 2 also outlines briefly the solution approach used to solve the problem and discusses computational experience with the procedure. Section 3 discusses an illustrative application of the model in the New England region. This analysis seems to indicate, among other things, that (a) it is generally less cost-effective to invest in expensive, large-scale cleanup technology that would be used only rarely for large (>50,000 gallons) spills than to acquire capability to cope with smaller spills that occur on a daily basis, and (b) equipment to combat small spills should be geographically dispersed across many small local response sites, whereas capability to respond to large spills should be consolidated at one or a few large-scale strike centers. Finally,

Section 4 comments on the results of this work and suggests directions for further research.

1. Background, Structure of the Strategic Problem, and Modeling Philosophy

1.1. Prior Research

Although the *general* oil spill literature is very rich, it is fair to say that this topic has received to date far less attention on the OR/MS literature than it really deserves. However, several related studies are worthy of note. These include Charnes et al. (1976, 1979), TSC (1979), Bellardo, Karwan and Wallace (1984a), and Belardo et al. (1984b). Of these studies, only TSC, Belardo et al. (1984b), and, to a lesser extent, Charnes et al. (1979) addressed the locational aspect of the strategic problem. Charnes et al. (1979) developed a "chance-constrained, goal programming" model to aid U.S. Coast Guard (USCG) managers in formulating policies with respect to planning for various types of equipment required to control oil spills. However, they simultaneously considered strategic and low-level operational decisions and applied their model only to a small illustrative example (3 equipment sites, 3 equipment types, 3 spill sites). TSC investigated how much equipment should be stockpiled by the USCG in order to satisfy the Presidential directive set forth by Jimmy Carter (1977) in his message to Congress, March 1977: "the goal is to respond adequately to a 100,000-ton spill within 6 hours." TSC used a rather "ad hoc" procedure, which began by determining "optimal" equipment *locations* solely from information on spill *frequencies* across the United States, *and then* estimated "optimal" equipment *levels* at those locations by taking spill *volume* information into account. As a result, this study recommended that the overwhelming bulk of cleanup equipment in the United States should be stockpiled in the port of Philadelphia. Belardo et al. (1984b) developed a "partial covering" approach to siting response resources for major maritime oil spills, and applied their model to a case study of a spill in Long Island Sound. Their model considered a multiple objective function consisting of the probabilities of covering spills of various "groups" (in terms of environmental/economic harm), subject to some relatively simple cleanup resource availability constraints. The model's basic limitation was that equipment needs were determined on the basis of a *single* (standard) spill volume (300,000 gallons for the Long Island case). They thus neglected the large variability of the volume of an individual spill, which is perhaps one of the most important

probabilistic features of this problem (see the following subsection). Finally, none of these studies explicitly considered spill response and damage costs, or analyzed the trade-offs between these two on a cost/benefit basis.

1.2. Large Variability of Spill Volumes

The probabilistic characteristics of oil spills exhibit certain peculiarities that ultimately play a major role in the structure of the strategic problem. Such peculiarities have been reported by various researchers who analyzed oil spill statistics in the past—Devanney and Stewart (1974), among others. For instance, it is generally accepted that the distribution of the volume of an individual spill is extremely skewed. In that respect, it becomes meaningless to consider statistics such as average spill size, let alone base strategic decisions on such statistics. To state one example, the average spill size of all spills in and around U.S. waters in 1982 was reported by the USCG to be 2,290 gallons (USCG 1982). However, 55% of these spills were at least 20 times smaller than this average, while 61% of all oil spilled was spilled in just two spills, each of which was over 4,000 times the average (USCG 1982). It therefore becomes obvious that in strategic planning for future spills, one should devise a method that explicitly accounts for the large variability of spill volumes, rather than their average values.

1.3. Nonsimultaneity of Spills

It is reasonable to assume that a simultaneous occurrence of two (or more) spills that would cause queuing or congestion in the use of the response resources of a particular geographical area is *very unlikely* (although not impossible). The general validity of the nonsimultaneity assumption was tested in TSC, which estimated that the probability of *one* overlap of two or more spills of 50,000 gallons or more within a year and *within the entire United States* was on the order of 0.01–0.001 for a typical spill duration of 5 days. The same study indicated that there were only about 20 spills in this size range each year during the 1974–1977 period. Similar observations can be made about smaller spills (they are of much shorter duration): the total number of spills (all sizes) reported by the USCG for the five maritime New England states in 1981 and 1982 was 566 and 455, respectively, that is, in the neighborhood of 1–2 spills per day *for the entire region* (USCG 1982). Similar orders of magnitude can be observed in other U.S. areas. Under such circumstances, queuing or congestion considerations concerning equipment utilization become minor.

1.4. Complementarity of Response

History has shown that it is not at all unlikely that a particular spill be responded to from more than one facility. In particular, for very large, catastrophic spills, the response typically originates from more than one mutually supportive (or “complementary”) location because each facility is rarely able to handle such spills by itself. Notice that this situation contrasts with the usual assumption in more “classical” facility location problems (e.g., *K*-median, *K*-center), for which only one facility serves each demand point (typically its “closest” median or center). In strategic planning for oil spills, one should therefore pay attention to the complementary nature of the response.

1.5. Objective(s) of the Strategic Problem

In this paper we shall assume that the objective of the strategic problem is to minimize the expected sum of response system costs and the costs due to damages from spills that may occur in the area, the latter costs multiplied by a user-specified “weight.”

The consideration of both system costs and damage costs as part of the problem objective makes sense intuitively, because any response system requires funds and one would like to know not only how much a system would cost, but also how much damage that system would avert. However, it is clear that such an objective function implies not only risk neutrality on the part of the decision maker, but equally important, that oil spill damage costs can be evaluated with some confidence. We have assumed risk neutrality both for analytical convenience, and because very little or nothing has been reported to date regarding the risk preference structure of “society” regarding oil spills (the possible exception is the recent work of a Norwegian research group, who seem to have obtained some insights on this issue—see Fredrikson 1983). At the same time, and for reasons similar to those discussed earlier regarding expected volumes, we need to introduce constraints that somehow capture the decision maker’s risk aversion, particularly when it comes to very large, very rare, catastrophic spills. This requirement will be discussed in Section 1.6.

Regarding the evaluation of damages, despite the general consensus that “oil spill cleanup and damage assessment are not, and are never likely to be precise arts” (White and Nichols 1983), researchers *have* made some progress in this area in recent years. In this paper we shall take advantage of recent related work at MIT, the purpose of which has been to quantitatively evaluate the damage costs of an oil spill under a variety of scenarios. Such an approach takes

as input spill-specific information (location, size, sea state, wind, oil type), area-specific information (inventory of environmental and economic resources), and information about the response, and produces estimates of damages, broken down into several categories (value of lost oil, organisms, beaches, marshes, recreation, and so forth). The description of the damage assessment algorithms is beyond the scope of this paper and can be found elsewhere (see, for instance, Baird et al. 1982 for the original version, and Demis 1984 for a more recent version). This paper assumes that these damages can be predicted as a function of several spill parameters and of the response to the spill and describes how such information can be used in the strategic decision making process. We shall see that this approach exhibits considerable flexibility with regard to possible future improvements in damage assessment.

Finally, the role of the damage weight is twofold: first, it can be used to represent how much the decision maker is willing to pay in system costs in order to reduce damage costs by \$1 (and, in that respect, a high value of that weight increases the relative importance of damage costs vis-à-vis system costs). Second, the weight can be used to perform sensitivity analysis on the value of damages, which, as mentioned before, is never likely to be precisely known. We shall see the effect of weight variations in Section 3.

1.6. The "Benign Neglect" Response and Policy Stipulations on Response Times

The following sections will show that our model assumes that a "benign neglect" response to a spill is a *permissible* response. Of course, this assumption is patently false in the real world, where "benign neglect," or, essentially "doing nothing," is likely to be not only politically unacceptable, but also illegal under existing regulations. Nevertheless, and since oil pollution combat funds are by no means unlimited, we feel that one cannot reject such an assumption a priori, particularly if a "benign neglect" response can be justified on a cost/benefit basis (if the latter is the case, one can essentially rename that response as "surveillance and monitoring" so as to make it politically more palatable). Our model will give us the flexibility to examine the merits of such an option.

At the same time, there are constraints that the decision maker may wish to impose so as to guarantee some minimum level of response. An example of such constraints is the 1977 Presidential directive to "respond adequately to a 100,000-ton spill within 6 hours," provided of course that the phrases "respond adequately" and "within 6 hours" are unambiguously

defined. In this paper we shall have the option to consider similar "policy stipulations," which will be seen to translate into probabilistic constraints on response times. Such policy stipulations reflect the decision maker's risk aversion in an *indirect* way, not only in terms of insuring against the adverse consequences of a catastrophic spill in a particularly vulnerable area, but also in cases for which the risk is a "political" one—for example, when the system might be unable to handle a very large spill.

2. The Model

In generic terms the strategic oil spill response model can be described as follows:

1. Let J be a known and finite set of oil spill "risk points," that is, points in the area of interest where oil spills are likely to occur. At each $j \in J$, oil spills occur according to a Poisson process of known parameter f_j . The volume of any spill at $j \in J$ has a known probability mass function defined over a finite set K : the volume is equal to v_{jk} with probability p_{jk} with $k \in K$, and $\sum_{k \in K} p_{jk} = 1$ for all $j \in J$. Each spill of volume v_{jk} at $j \in J$ has a known duration T_{jk} , measured from the beginning of the *discharge* to the end of the *event*, the latter being the time beyond which no events related to that spill are significant enough to be accounted for (typically, T_{jk} is longer than the duration of the discharge). Assume that spill occurrences and volumes at $j \in J$ and $j' \in J$ are mutually independent for $j \neq j'$, and that the probability of simultaneous occurrence or overlap is negligibly small.

2. Let I be a known and finite set of candidate stockpiling facilities, none of which has any initial response capability. Opening $i \in I$ entails a fixed cost of FC_i (equivalent annualized value). There is a known and finite set E of equipment types that may be used to comprise the response capability at each $i \in I$. Assume that the acquisition cost of $e \in E$ stockpiled at $i \in I$ is a_{ie} per annum and per unit of capability stockpiled. FC_i and a_{ie} include all committed fixed and variable annual costs once i is opened or e is acquired (such as manning or maintenance), but do not include spill-specific costs (see item 5 to follow).

3. The capability u_{ie} of an equipment package of type $e \in E$ located at $i \in I$ is defined as the maximum volume of spill that the package is capable of fully recovering in the hypothetical situation in which (a) the package can operate on scene for the entire duration of the spill, and (b) its recovery efficiency (ratio of oil volume recovered versus encountered) is equal

to 1.0. We assume that capability is infinitely divisible (u_{ie} is continuous).

We note that this definition of capability involves a fair degree of aggregation and simplification regarding the cleanup process, for it condenses all four principal “capability” attributes of a real-world cleanup package (skimming rate, pumping rate, storage capacity and containment boom length) into *one* number (an equivalent spill volume). It also implies that increased on-scene time cannot make up for insufficient on-scene capability. In that respect, we note that aggregating four attributes into one makes sense *from a strategic perspective* and was also proposed elsewhere: in TSC, each of the four attributes was considered as a linear function of capability. Furthermore, compensating insufficient on-scene capability by more on-scene time is considered inappropriate (again from a strategic viewpoint) because oil that cannot be readily recovered by an insufficient capability is likely to dissipate by winds and currents, making any prolonged application of such a capability ineffective.

4. Conditions (a) or (b) of item 3 will not apply in general. We assume the following *simplified* characterization of the performance of types $e \in E$ on a spill at $j \in J$ (a more detailed characterization can be found in Ziogas: first, via a coefficient of “operability” OP_{ej} ($0 \leq OP_{ej} \leq 1$), defined as the *a priori* probability that weather conditions at j will allow any operation of equipment type e (typically, OP_{ej} is the probability that the sea state at j is below a certain threshold that depends on e). Second, via a coefficient of a “recovery efficiency” RE_e ($0 \leq RE_e \leq 1$), which is the fraction of oil volume encountered by the device that is actually recovered, *given* the equipment is in fact, operable. At first glance, the description of capability and performance by the triplet (u_{ie}, RE_e, OP_{ej}) may seem redundant, since an equipment package of such characteristics would seem exactly equivalent to another equipment package described by the triplet $(RE_e u_{ie}, 1.0, OP_{ej})$. We remove this redundancy by assuming that if (u_{ie}, RE_e, OP_{ej}) is applied to a spill of size u_{ie} , and if condition (a) of the previous paragraph is satisfied, then the equipment would recover an oil volume of $RE_e u_{ie}$, while the amount $(1 - RE_e)u_{ie}$ would escape and *become completely nonrecoverable by the entire response*. In general, if u_{ie} , deployed alone for the entire duration of a spill of volume v_{jk} , is operable, it would recover a volume of oil equal to $RE_e \min(v_{jk}, u_{ie})$ (and *not* $\min(v_{jk}, RE_e u_{ie})$). The remainder of the oil, $v_{jk} - RE_e \min(v_{jk}, u_{ie})$, would *escape*. Of the escaped oil, a volume equal to $(1 - RE_e) \min(v_{jk}, u_{ie})$ would in fact be nonrecoverable, while the rest, $\max(0, v_{jk} - u_{ie})$, could potentially be

recoverable by other equipment packages. Of course, if such capability is not operable, then the entire volume of v_{jk} would escape. Finally, if such capability arrives on scene late, it would be effective only for a fraction of the duration of the spill (see also item 8 to follow).

5. The dispatching time of equipment $e \in E$ from $i \in I$ to $j \in J$ is equal to d_{iej} , and the equivalent transportation cost is TC_{iej} per unit of capability transported. d_{iej} (and TC_{iej}) include duration (and costs) of mobilization and set-up delays at i as well as of deployment delays at j . TC_{iej} does not include on-scene operational costs, which are equal to b_{iek} per unit of capability if spill volume is v_{jk} .

6. The model assumes that escaping oil from a spill of volume v_{jk} at $j \in J$ that cannot be recovered creates *damage* of cost equal to DP_{jk} per unit volume of nonrecoverable oil, with the “damage potential” DP_{jk} being a *known* function of j and k (and, in general, a nonlinear function of v_{jk}). The ratio $w_{jk} = DP_{jk}/T_{jk}$ is defined as the “damage rate” for (j, k) and reflects the rate at which damage costs accumulate through time per unit volume of nonrecoverable oil and throughout the duration of the event. Notice that this formulation assumes (a) that damage costs vary linearly with the volume of nonrecoverable oil (although the factor of proportionality is a general function of v_{jk}), and (b) that w_{jk} is constant through time. Both (a) and (b) may be false in the real world, for obvious reasons. It turns out that assumption (b) is nonbinding in the formulation of the model and can easily be relaxed. Assumption (a) is more drastic (despite the fact that DP_{jk} and w_{jk} are general functions of v_{jk}), and can potentially cause significant approximation errors if damages are *highly* nonlinear. We have tested this assumption with realistic damage cost data and have concluded that although nonlinearities do exist, the error introduced by this approximation is not significant. DP_{jk} s are computed by the MIT Damage Assessment subroutine (Demis) after averaging over all possible oil types, wind directions and seasons corresponding to v_{jk} .

7. The response to a spill of volume v_{jk} and duration T_{jk} at $j \in J$ consists of capabilities r_{iej} (some or all of them zero) of various equipment types $e \in E$ being dispatched from various locations $i \in I$, arriving on scene at various times, and recovering oil until the end of the event (if, in fact, such capabilities *can* operate on scene). Note that since T_{jk} and d_{iej} are typically longer than the actual duration of the *discharge*, this response scheme implies that equipment on scene would essentially recover oil *after* the latter

is released, in a cumulative fashion. We assume that r_{iej} is continuous.

8. Focusing for the moment on the contribution of an individual equipment $e \in E$ (dispatched *alone* from $i \in I$ to v_{jk} at $j \in J$), we can note the following: first, it would not make sense to have $r_{iej} > v_{jk}$, hence $r_{iej} \leq v_{jk}$ (no excess capability on scene). Second, irrespective of whether e is operable at j or not, a spill volume of $v_{jk} - \text{RE}_e r_{iej}$ would certainly escape recovery and create damage of cost equal to $\text{DP}_{jk}(v_{jk} - \text{RE}_e r_{iej})$. The remainder of oil volume $\text{RE}_e r_{iej}$ would either fully escape recovery (with probability $1 - \text{OP}_{ej}$, if e cannot operate) and create damage of cost equal to $\text{DP}_{jk} \text{RE}_e r_{iej}$, or would be *partially* recovered (with probability OP_{ej}). In the latter case, the recovery would be partial because e cannot arrive on scene earlier than d_{iej} . The additional damage cost in the latter case would be equal to $w_{jk} \text{RE}_e d'_{iej} r_{iej}$, where $d'_{iej} = \min(d_{iej}, T_{jk})$. Hence, the total expected damage costs for this scenario are equal to

$$\begin{aligned} & \text{DP}_{jk}(v_{jk} - \text{RE}_e r_{iej}) + (1 - \text{OP}_{ej})\text{DP}_{jk} \text{RE}_e r_{iej} \\ & \quad + \text{OP}_{ej} w_{jk} \text{RE}_e d'_{iej} r_{iej} \\ & = \text{DP}_{jk} v_{jk} - \text{OP}_{ej} \text{RE}_e (\text{DP}_{jk} - w_{jk} d'_{iej}) r_{iej}. \end{aligned} \quad (1)$$

Since the first term in (1) is the damage cost under no response, the second term is the expected damage cost *averted* due to r_{iej} . It is possible to verify that, if $d_{iej} \geq T_{jk}$ (equipment cannot arrive on scene before the end of the spill), the second term is equal to zero (as expected).

9. Now generalizing expression (1) to more than one equipment package dispatched to v_{jk} at $j \in J$ (where, again, we have no excess capability on scene, that is, $\sum_{i \in I} \sum_{e \in E} r_{iej} \leq v_{jk}$), and defining $x_{jk} = v_{jk} - \sum_{i \in I} \sum_{e \in E} r_{iej}$ (≥ 0 , “unsatisfied demand” from v_{jk} when the response is defined by *array* r_{iej}), we find that the total expected damage costs are equal to

$$\begin{aligned} & \text{DP}_{jk} x_{jk} \\ & + \sum_{i \in I} \sum_{e \in E} [\text{OP}_{ej} \text{RE}_e w_{jk} d'_{iej} + (1 - \text{OP}_{ej} \text{RE}_e) \text{DP}_{jk}] r_{iej}. \end{aligned} \quad (2)$$

Expression (2) has the following intuitive interpretation: its first term, $\text{DP}_{jk} x_{jk}$, can be considered as the damage cost due to any unsatisfied demand, that is, the damage cost that would certainly occur whenever the total on-scene capability is less than the volume of the spill. The second term, represented by the double summation, is the expected damage cost due to (a) delays and (b) inefficiencies in response, both of which would allow some additional oil to escape. Without loss of generality, we shall call the first and

second terms of (2) the “cost of damage due to unsatisfied demand” and “cost of damage due to delay and equipment inefficiency,” respectively.

Based on the previous discussion, we are in a position to formulate the strategic oil spill response problem as an optimization problem. The *decision variables* are (for $i \in I, e \in E, j \in J, k \in K$):

$$\begin{aligned} y_i &= \begin{cases} 1 & \text{if candidate location } i \text{ is opened} \\ 0 & \text{otherwise} \end{cases} \\ u_{ie} &= \text{response capability of equipment of type } e \\ & \quad \text{stockpiled at } i \\ r_{iej} &= \text{response capability of equipment of type } e, \\ & \quad \text{stockpiled at } i, \text{ that is designated to respond to} \\ & \quad \text{a spill of volume } v_{jk} \text{ at } j \\ x_{jk} &= \text{unsatisfied demand at } j \text{ when spill volume is } v_{jk}. \end{aligned}$$

The *objective function* Z of the problem is equal to $(Z_1 + Z_2 + Z_3 + Z_4) + W(Z_5 + Z_6)$, whose terms are defined as follows:

$$Z_1 = \text{cost to open new facilities} = \sum_{i \in I} \text{FC}_i y_i$$

$$Z_2 = \text{cost to acquire capability} = \sum_{i \in I} \sum_{e \in E} a_{ie} u_{ie}$$

$Z_3 =$ expected cost to mobilize and transport equipment to the spill site

$$= \sum_{i \in I} \sum_{e \in E} \sum_{j \in J} \sum_{k \in K} f_j p_{jk} \text{OP}_{ej} \text{TC}_{iej} r_{iej}$$

$Z_4 =$ expected cost of cleanup

$$= \sum_{i \in I} \sum_{e \in E} \sum_{j \in J} \sum_{k \in K} f_j p_{jk} \text{OP}_{ej} b_{iek} r_{iej}$$

$Z_5 =$ expected cost of damage due to unsatisfied demand

$$= \sum_{j \in J} \sum_{k \in K} f_j p_{jk} \text{DP}_{jk} x_{jk}$$

$Z_6 =$ expected cost of damage due to delay and equipment inefficiency

$$= \sum_{i \in I} \sum_{e \in E} \sum_{j \in J} \sum_{k \in K} f_j p_{jk} [\text{OP}_{ej} \text{RE}_e w_{jk} d'_{iej} + (1 - \text{OP}_{ej} \text{RE}_e) \text{DP}_{jk}] r_{iej}$$

and

$W =$ user-specified damage weight.

The *constraints* of the problem are of four types:

1) No capability can be acquired unless the corresponding facility is opened:

$$y_i M - u_{ie} \geq 0 \quad i \in I, e \in E. \quad (3)$$

where $M = \max_{j,k} v_{jk}$

2) No facility can dispatch more than its capability:

$$u_{ie} - r_{iej} \geq 0 \quad i \in I, e \in E, j \in J, k \in K. \quad (4)$$

(Notice that this constraint differs from its equivalent in “warehouse location” formulations that typically require that $u_{ie} - \sum_{j \in J} \sum_{k \in K} r_{iej} = 0$. The nonsimultaneity assumption is the main reason for this difference.)

3) Definition of unsatisfied demand:

$$\sum_{i \in I} \sum_{e \in E} r_{iej} + x_{jk} = v_{jk} \quad j \in J, k \in K. \quad (5)$$

4) Nonnegativity and integrality:

$$u_{ie} \geq 0, \quad r_{iej} \geq 0, \quad x_{jk} \geq 0, \quad y_i \in \{0, 1\} \quad (6)$$

$$i \in I, e \in E, j \in J, k \in K.$$

Collecting terms in the objective function, we arrive at the following Mixed Integer Programming (MIP) formulation:

$$\begin{aligned} \text{minimize } Z = & \sum_{i \in I} FC_i y_i + \sum_{i \in I} \sum_{e \in E} a_{ie} u_{ie} \\ & + \sum_{j \in J} \sum_{k \in K} G_{jk} x_{jk} \\ & + \sum_{i \in I} \sum_{e \in E} \sum_{j \in J} \sum_{k \in K} H_{iej} r_{iej}, \end{aligned} \quad (P)$$

subject to expressions 3, 4, 5 and 6. In this formulation, we have made the following transformations for notational convenience:

$$G_{jk} \equiv f_j p_{jk} DP_{jk} \cdot W \quad (7)$$

$$\begin{aligned} H_{iej} \equiv & f_j p_{jk} [\text{OP}_{ej}(\text{TC}_{iej} + b_{iek} + \text{RE}_e w_{jk} d'_{iej}) \cdot W] \\ & + (1 - \text{OP}_{ej} \text{RE}_e) DP_{jk} \cdot W]. \end{aligned} \quad (8)$$

It is clear that this formulation explicitly allows a “benign neglect” response to some (or even all) spills, a policy that may be difficult to justify politically. One way of guaranteeing some minimum acceptable level of response is via what we call “policy stipulations on response times,” which are defined as follows: for each $j \in J$, let α_j and R_j be user inputs satisfying $0 < \alpha_j < 1$ and $R_j > 0$. Define $Q_j(\alpha_j)$ as the “ α_j -percentile” of the spill volume distribution at j (that is, $Q_j(\alpha_j) = \text{CDF}_j^{-1}(\alpha_j)$, where CDF_j^{-1} is the inverse of the cumulative distribution function of the discretized spill volume at j). A strategic response plan is said to satisfy an (α, R) policy stipulation if it guarantees a response capability of at least $Q_j(\alpha_j)$ to every $j \in J$, and if all equipment designated to comprise such capability can be dispatched to arrive on scene in no more than R_j time.

Such stipulations can be handled as follows: for a given (α, R) stipulation, define the following two

classes of sets of all $j \in J$:

$$\begin{aligned} \text{(a) The “covering set” } \text{COVER}_j(R_j) & \\ = \{i, e) : i \in I, e \in E, d_{iej} \leq R_j\}; & \quad (9) \end{aligned}$$

$$\begin{aligned} \text{(b) The “volume index set” } \text{VOL}_j(\alpha_j) & \\ = \{k : k \in K, v_{jk} \leq Q_j(\alpha_j)\}. & \quad (10) \end{aligned}$$

$\text{COVER}_j(R_j)$ consists of those equipment location-type pairs that lie within a time radius of R_j from j . $\text{VOL}_j(\alpha_j)$ includes all indices $k \in K$ of the spill volume distribution at j that correspond to volumes within the α_j -percentile of that distribution.

Notice that an (α, R) stipulation need not be feasible. A *necessary* condition for feasibility is that $\text{COVER}_j(R_j) \neq \emptyset$ for all $j \in J$. If this is *not* the case, there will be at least one j that cannot be responded to within a time of R_j , and the overall problem becomes infeasible. Of course, that condition alone is not sufficient for feasibility, for there may not be enough capability within $\text{COVER}_j(R_j)$ to satisfy a demand of v_{jk} for those volumes k belonging to $\text{VOL}_j(\alpha_j)$.

Within (P), such stipulations can be handled easily *without* the imposition of further constraints. Indeed, if we *reset* $G_{jk} = +\infty$ for all $j \in J$ and $k \in \text{VOL}_j(\alpha_j)$, and $H_{iej} = +\infty$ for all $j \in J, k \in \text{VOL}_j(\alpha_j)$ and $(i, e) \notin \text{COVER}_j(R_j)$, we effectively prohibit any unsatisfied demand at those (j, k) s for which $v_{jk} \leq Q_j(\alpha_j)$, and, in addition, we consider any response originating outside the appropriate covering set as infeasible. In that respect, (P) is infeasible whenever $\min Z = +\infty$.

The fact that parameters α_j and R_j can be chosen to depend on the risk point j provides considerable flexibility to the decision maker, who can thus design a response system that best reflects his own aversion to catastrophic spills in certain areas by a judicious choice of these parameters. For instance, certain risk points may be particularly “sensitive” from an environmental viewpoint (high damage potential), and for those, one would typically wish to set α_j close to 1.0 and R_j relatively low. Other risk points may be less sensitive, and hence receive more relaxed stipulations, if any. This flexibility permits the model to capture (albeit indirectly) the decision maker’s potential risk aversion, which cannot be reflected in the assumed objective function of (P).

Our *solution approach*, so far, has consisted of a fast, “network synthesis” heuristic that solves for u_{ie} , x_{jk} and r_{iej} , given a combination of y_i s, and of an implicit enumeration procedure that embeds this heuristic and solves with respect to the y_i s. Thus far, the worst-case performance of the algorithm, as well as its

average-case behavior on random instances, are unknown. However, we do have some experience on how the solutions produced by the network synthesis heuristic compare with those produced by solving the associated linear program *exactly*. Such a comparison, made on a set of realistic *oil spill* instances, revealed an average deviation from optimality of about 1–2%. Furthermore, the network synthesis heuristic proved about 4–6 times faster than the exact, simplex-based procedure. Further details about the algorithm can be found in Tharakan.

To date, we have been able to handle problems involving up to 8 facilities, up to 5 equipment types at each facility, up to 19 risk points and up to 3 spill volumes at each risk point. Computationally, the most sensitive parameter seems to be the number of equipment locations i for which $FC_i \neq 0$, that is, the number of *new* facilities to be established. In realistic problems, one can generally expect that number to be small (say, no more than 10). The running time of a typical application involving 8 facilities, 4 equipment types, 12 risk points, 3 volumes and no response time stipulations is a little over 1 CPU minute on a VAX 11/782 if *all* 8 fixed costs are *zero* (Run 4 of Section 3.1). However, the same run takes about 3 CPU minutes if *all* 8 fixed costs are *nonzero* (Run 5 of Section 3.1). Response stipulations generally tend to reduce the running time, because many arcs of the network are then assigned infinite costs and essentially eliminated. For instance, the previous run takes only about 45 CPU seconds if a stipulation is used (and if fixed costs remain zero—Run 8 of Section 3.1). The running times suggest that even larger problems can be tackled with no considerable difficulty, particularly if the codes implementing the model are further refined.

3. An Illustrative Example

This section presents an application of our model in the New England area. The main goal of this investigation is to demonstrate the versatility of the model, perform sensitivity analysis on various parameters, and, generally, provide insights into various issues in strategic planning for oil spill response for this geographical area. In that respect, it is important to emphasize that, although a significant attempt was made to use input data that were as realistic as possible, this exercise should be considered only as an *illustration* of the potential of the model and not a description of its actual *implementation* in the New England area. The scenario that we examined had the

following main features:

1. There are a total of 19 risk points. Their locations and frequencies are displayed in Table I. Frequencies have been derived using information on oil throughputs in the area and the methodology outlined in Devanney and Stewart. The risk points can be divided into two major categories, each of which is further divided into two subcategories. The two major categories are “small” spills (up to about 50,000 gallons) and “large” spills (50,000 gallons and above). “Small” spills are denoted by asterisks. “Small” spills are further divided into harbor spills (risk points 1, 3, 5, 7, 10 and 13) and platform operational spills (risk point 18). “Large” spills are further divided into tankship spills (risk points 2, 4, 6, 8, 9, 11, 12, 14, 15, 16, and 17) and platform blowout spills (risk point 19). The locations of tankship spill risk points have been selected along tanker traffic lanes feeding into major oil terminals in the area, the risk points usually located about 50–100 miles off the corresponding ports. Spill volume distributions have been discretized to 10^2 , 10^3 and 10^4 gallons for small spills and to 10^5 , 10^6 and 10^7 gallons for large spills. The corresponding probabilities (p_{jks}) have been set to 0.895, 0.0855 and 0.0195 for harbor spills, 0.989, 0.01 and 0.001 for platform operational spills, 0.25, 0.60 and 0.15 for tankship spills and 0.61, 0.32 and 0.07 for platform blowout spills. These figures have been calculated so that the corresponding probability mass functions match the parent

Table I
Risk Points and Frequencies^a

j	Location	f_j (Spills/Year)
1*	Penobscot Bay, ME (harbor)	7.00
2	Penobscot Bay, ME (offshore)	0.04
3*	Portland, ME (harbor)	55.00
4	Portland, ME (offshore)	0.35
5*	Portsmouth, NH (harbor)	6.75
6	Portsmouth, NH (offshore)	0.04
7*	Boston, MA (harbor)	60.00
8	Boston, MA (offshore 1)	0.19
9	Boston, MA (offshore 2)	0.19
10*	Providence, RI (harbor)	66.75
11	Providence, RI (offshore 1)	0.21
12	Providence, RI (offshore 2)	0.21
13*	New Haven, CT (harbor)	44.25
14	New Haven, CT (offshore 1)	0.14
15	New Haven, CT (offshore 2)	0.14
16	Cape Cod Canal, MA (offshore 1)	0.06
17	Cape Cod Canal, MA (offshore 2)	0.06
18*	Georges Bank, MA (operational)	56.00
19	Georges Bank, MA (blowouts)	0.02

^a Asterisks denote small spills.

Table II
 “Basic” Dispatching Time Matrix (in Hours)^a

<i>i</i>	Location	<i>j</i>																		
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
1	Portland, ME	3	7	1	2	2	4	3	5	13	6	8	10	7	10	9	7	6	15	15
2	Portsmouth, NH	4	7	2	4	1	2	1	3	11	5	6	8	6	9	8	6	5	13	13
3	Boston, MA	5	7	3	5	1	4	1	2	10	4	6	8	5	8	7	5	4	12	12
4	Woods Hole, MA	6	7	5	7	3	6	2	4	8	2	3	7	4	7	7	2	4	10	10
5	Providence, RI	8	10	7	9	5	7	4	6	10	1	2	7	6	7	7	2	3	11	11
6	New Haven, CT	9	11	8	10	6	8	5	7	11	2	2	7	1	5	5	1	3	13	13
7	Provincetown, MA	5	7	3	7	4	8	3	7	7	3	3	5	6	7	7	3	2	9	9
8	Penobscot, ME	1	7	2	3	3	5	4	6	14	7	9	11	7	11	10	9	8	16	16

^a Does not include mobilization and on-scene delays, which vary with scenario (see text). All times are independent of equipment type deployed.

Table III
 Equipment Cost and Performance Data^a

Equipment Category	Acquisition Cost (\$/gal)	Transport Cost (\$/gal/mile)	Cleanup Cost (\$/gal) at Volume:							Recovery Efficiency	Operability	
			10 ²	10 ³	10 ⁴	10 ⁵	10 ⁶	10 ⁷	Harbor		Non-Harbor	
1	5.0	0.0015	30	15	8	8	8	8	0.75	1.0	0.3	
2	7.0	0.0015	30	15	8	8	8	8	0.80	1.0	0.4	
3	1.0	0.00015	40	20	10	5	2	0.6	0.80	1.0	0.6	
4	1.2	0.00015	40	20	10	5	2	0.6	0.85	1.0	0.7	

^a Notice that transportation costs are given in \$/gal/mile. These figures should be multiplied by the appropriate dispatching distances to obtain TC_{iej}.

continuous distributions in terms of mean and variance, as per Devanney and Stewart’s statistical analysis. Notice the dramatic frequency difference between small and large spills.

2. There are a total of 8 candidate equipment stockpiling facilities. Table II displays the locations of those facilities as well as the dispatching times (in hours) from those points to each of the 19 risk points. This table does not include mobilization and on-scene deployment delays, which vary with each case examined. In this example, we assume that all times are independent of equipment type deployed.

3. We assume a total of 4 candidate equipment types, all in the mechanical removal category. Table III shows information on acquisition, transportation and on-scene operational costs for those types, as well as efficiency and operability, the latter for harbor and non-harbor spills. For illustration purposes, Table III reflects the basic breakdown among oil spill cleanup hardware, with categories 1 and 2 essentially associated with small spills in protected areas, while categories 3 and 4 consist of heavy-duty equipment, typically used in large high-seas spills. Category 2 (4) differs from category 1 (3) only in acquisition costs, efficiency and operability. Notice the economies of scale in on-scene cleanup costs realized by all cate-

gories (up to 10⁴ gallon spills for categories 1 and 2 and throughout the range for categories 3 and 4).

4. Finally, Table IV displays the values of the “damage potentials,” DP_{jk}, obtained by running the MIT

Table IV
 Damage Potential (DP_{jk})^a

<i>j</i>	DP _{j1}	DP _{j2}	DP _{j3}
1*	10.88	7.86	16.30
2	1.39	1.08	1.06
3*	59.20	68.39	51.60
4	1.34	1.18	0.98
5*	10.04	3.67	1.39
6	2.22	1.28	0.99
7*	89.34	91.28	75.90
8	1.32	1.19	1.03
9	1.15	1.12	1.04
10*	66.74	79.87	60.19
11	2.33	1.40	1.08
12	0.96	0.96	1.02
13*	6.87	7.06	41.75
14	1.59	1.03	0.98
15	0.92	0.92	0.94
16	1.14	1.04	0.97
17	1.11	1.45	11.47
18*	0.85	0.85	0.85
19	0.85	0.90	0.95

^a Damage per unit volume of nonrecoverable oil from a spill of volume index *k* at risk point *j* (\$/gallon). Asterisks denote “small” spills.

Damage Assessment subroutine. It is interesting to observe that small spills generally exhibit significantly higher damage potentials than large spills, quite probably because these spills occur much closer to shore and are likely to impact the environmental and economic resources along the coast more severely.

In an early application of the strategic model in New England, we examined both categories of spills (small and large) *simultaneously* (see for instance Tharakan, and Psaraftis, Nyhart and Betts (1983)). In this paper we examine these classes of spills *separately*. Indeed, Tables I and IV already provide sufficient evidence that these two categories of spills are dramatically different. Given this result, and *in addition* the classical set of questions to be addressed by the strategic model, a whole host of new issues become important: Which portion of the overall system and/or damage cost is attributable to each category? How do system/damage cost trade-offs differ between the two? Which of the two categories of spills is most cost-effective to clean up? How do strategic response policies between the two categories compare? How valid or invalid are currently stipulated policies for these two classes? Insights regarding these and other questions were difficult (or impossible) to obtain from our earlier analysis, which lumped together both spill

categories. As the rest of this section shows, a separate investigation sheds more light on these issues.

3.1. "Large" Spill Analysis

Table V summarizes the results of 12 variants of the "large" spill problem. The table contains information on the damage weight chosen, the six objective function components Z_1 to Z_6 (as per Section 2), which of the eight facilities are opened, and which risk points are served by each facility, as well as the level and type of response capability stockpiled at each of the facilities that are opened. Notice that the table *does not* include disaggregate information on how to allocate the capability associated with each facility *and equipment type* among risk points *and spill volumes*. Since this output—which essentially defines the optimal response policy for the problem at hand—is fairly voluminous (array r_{iej} has a dimension of 1,152 here), we shall refer to it only on a selective basis, so as to obtain insights into the structure of the problem and the behavior of its solutions. Table V contains the following results:

1. Run 1 is the "basic" scenario, with the damage weight being equal to 1.0, no fixed costs to open a facility, no response time stipulations, and a "stand-

Table V
"Large" Spill Analysis for New England^a

Run No.	Damage Weight	Cost Components (as per Section 2)						Facilities Opened	Risk Points Served (as per Table I)	Response Capability	Equipment/Technology Type
		Z_1	Z_2	Z_3	Z_4	Z_5	Z_6				
1	1.0	—	—	—	—	4.651	—	—	—	—	—
2	5.0	—	1.0	0.015	1.413	3.109	0.826	Woods Hole	All	1,000,000	3
3	10.0	—	10.0	0.040	2.219	—	2.463	Woods Hole	All	10,000,000	3
4	15.0	—	10.2	0.046	2.455	—	2.292	Woods Hole Provincetown Penobscot	All All All	9,000,000 900,000 100,000	3 4 4
5	15.0	0.5	10.2	0.043	2.455	—	2.292	Woods Hole	All	9,000,000 1,000,000	3 4
6	15.0	0.5	10.2	0.045	2.455	—	2.292	Provincetown	All	9,000,000 1,000,000	3 4
7	15.0	0.5	10.2	0.068	2.455	—	2.293	New Haven	All	9,000,000 1,000,000	3 4
8	1.0	—	10.0	0.040	2.219	—	2.463	Woods Hole	All	10,000,000	3
9	1.0	—	10.0	0.042	2.219	—	2.463	Provincetown	All	10,000,000	3
10	1.0	—	10.0	0.040	2.219	—	2.442	Woods Hole	All	10,000,000	3
11	1.0	—	—	—	2.117	—	2.461	Portsmouth Provincetown	2, 4, 6, 8, 9 9, 11, 12, 14, 15, 16, 17, 19	10,000,000 10,000,000	3* 3*
12	1.0	—	1.0	0.015	8.133	—	2.463	Portsmouth Providence Provincetown	2, 4, 6, 8, 9 11, 12, 14, 15, 16, 17, 19 All	10,000,000 10,000,000 1,000,000	3* 3* 3

^a All costs are per annum and in millions of dollars. All capabilities are in gallons. Asterisks denote leased equipment.

ard” mobilization delay of 3 hr at each facility, plus a 3-hr deployment delay at each of the spill locations. Not surprisingly, the optimal solution under this scenario is a “benign neglect” response, with expected damages totalling \$4.651 million per year, all in the form of unsatisfied demand. Clearly, if the system must have *some* “large” spill capability, then one should be willing either to accept a higher weight on damage costs, or to impose response time stipulations (or both). We first present the analysis of what happens if the damage weight is increased (and everything else remains the same).

2. The “benign neglect” policy of Run 1 remains optimal even if the damage weight is doubled. However, if the weight becomes 5.0, we begin to see *some* capability being stockpiled at Woods Hole (Run 2). Notice that this capability can respond only to spills up to 1 million gallons, leaving an expected damage of \$3.935 million per year, mainly in unsatisfied demand. Notice also that this capability averts only \$0.716 million per year in damages vis-à-vis the “benign neglect” policy, at an additional expected system cost of \$2.428 million per year. Finally, observe that the technology (or equipment type) chosen among the four candidates is type 3, probably because its low acquisition costs offset the benefits realized by a higher efficiency and operability (the latter are features of technology 4).

3. Raising the weight even further (to 10.0 in Run 3 and 15.0 in Run 4) produces some *expected* and some *unexpected* results. The fact that total system capability increases to a level capable of handling the maximum spill size (10 million gallons) is certainly predictable. One can observe this result in Run 3, where 10 million gallons of capability (again, of type 3) are stockpiled (again, in Woods Hole). Notice again the increased price required to reduce damages (\$12.259 million/year in system costs to reduce damages by \$2.188 million/year). What happens in Run 4, where the damage weight is raised to 15.0, is less predictable. Several interesting observations emerge from this scenario: first, two additional facilities are established, one in Provincetown and one in Penobscot, taking away a small fraction of capability from Woods Hole, with the total capability remaining constant at 10 million gallons. Second, while Woods Hole continues to be equipped with type 3 equipment, the two new facilities introduce technology 4, which is more expensive, yet more efficient and reliable. This situation can be explained by observing that the only way to further reduce damages from those of Run 3 is by reducing Z_6 (Z_5 is already zero). But Z_6 is the

damage due to delays and inefficiencies in the response. Delays can be reduced by “spreading out” the capability to three locations instead of one. Inefficiencies can be reduced by using equipment type 4 instead of 3. Indeed, this policy results in an *additional* damage reduction of \$0.171 million/year with an *additional* system cost of \$0.442 million/year.

4. Still, it is not immediately clear from Table V how the “mix” of equipment technologies shown in Run 4 is used. In fact, it is not even clear why *both* technologies are needed. We can shed more light on this issue only by looking at the entire response array r_{ijk} (not shown here). This investigation reveals the following facts: while all 3 facilities (using both types of technology) respond to the 10-million gallon spills, only Provincetown *and* Penobscot (which are equipped with technology 4 only) respond to the 1-million gallon spills (which are more frequent). Finally, *either* Provincetown (risk points 9, 11, 12, 14, 15, 16, 17 and 19) *or* Penobscot (risk points 2, 4, 6, 8), using technology 4 only, respond to the smallest spills (in the range of 100,000 gallons). In other words, equipment type 3 (stocked at Woods Hole) is acquired only to respond to the massive, 10-million gallon spills. Since this technology is used relatively rarely, it is preferable to type 4, which would cost more to acquire. By contrast, because of its higher efficiency and reliability, type 4 is more desirable for responding to the smaller but more frequent 1-million and 100,000-gallon spills, with no massive outlays.

5. Imposing a fixed cost of \$0.5 million per year to open each of the facilities (again assuming a damage weight of 15.0 and everything else equal) “shrinks” the optimal location set to only one facility, located again in Woods Hole (Run 5). If Woods Hole is excluded from the candidate location set (Run 6). Provincetown, which is second best, becomes the location of choice. And if Woods Hole, Provincetown and Penobscot are all excluded (Run 7), then New Haven becomes the best. Notice the small degradation of performance in Runs 6 and 7 vis-à-vis Run 5, and the continuing technology “mix” of types 3 and 4 in all three cases.

6. Runs 8, 9 and 10 return to the unit damage weight case, this time with response time stipulations and without fixed costs. Run 8 examines a (0.99, 18 hr) stipulation for all risk points, while Run 9 tightens the stipulation to (0.99, 15 hr). Finally, in run 10 we assume that mobilization *and* deployment delays are reduced by 3 hours everywhere, and then impose a (0.99, 12 hr) stipulation. Comparing Runs 8 and 10, we can see that the benefits of a total delay reduction

of the 6 hour amount to \$0.021 million/year in damages averted. These benefits should be compared to the costs (monitoring, administrative, etc.) that would be necessary to make such a reduction possible. Notice that in all these cases we fall back to the use of the "cheaper" technology of type 3, probably because of the "low" damage weight.

7. Finally, Runs 11 and 12 show how our model can be used to address "buy versus lease" decisions. Suppose that the strategic planner has the option to either acquire equipment of type 3 (and incur a cost structure as per Table III), or *lease* equipment of the same type from a private contractor at a flat rate of X dollars per gallon of capability *used* (no acquisition or transportation costs). If we assume that the contractor agrees to make this equipment available at the locations desired by the decisionmaker, and to be paid only if his equipment is used, the question is which of the two options—buy or lease—is the most desirable. The answer, of course, depends on the value of X . Under a (0.99, 18 hr) stipulation, run 11 shows that if $X = \$1/\text{gallon}$ (a very low price), then an "all-leased" capability is adopted. Notice that *for the first time, each of the two facilities selected serves a different*

set of risk points. This outcome can be explained by the observation that this option allows us for the first time to double the total capability available (to 20 million gallons), and this possibility eliminates the need for complementary response (with the exception of risk point 9). Run 12 shows that if $X = \$5/\text{gallon}$, the system can have both acquired and leased capability. Notice again the allocation pattern among risk points. For X sufficiently high, we obviously return to an "all-buy" policy, as suggested by Run 8.

3.2. "Small" Spill Analysis

Table VI is the equivalent of Table V for seven variants of the small spill problem. These cases illustrate the following results:

1. Run 1 displays the "benign-neglect" policy. Notice that by contrast to the "large" spill case, in this instance damage weight must be very low (close to zero) to make such a policy optimal. If the damage weight is set equal to one, the optimal policy would stockpile capability and leave no unsatisfied demand. This outcome is displayed in Run 2, which assumes zero mobilization and on-scene deployment delays, zero fixed costs and no response stipulations. Notice

Table VI
"Small" Spill Analysis for New England^a

Run No.	Damage Weight	Cost Components (as per Section 2)						Facilities Opened	Risk Points Served (as per Table I)	Response Capability	Equipment/Technology Type
		Z ₁	Z ₂	Z ₃	Z ₄	Z ₅	Z ₆				
1	0.0	—	—	—	—	5.090	—	—	—	—	—
2	1.0	—	0.028	0.003	1.500	—	1.067	Portland Portsmouth Boston Woods Hole Providence New Haven Penobscot	1, 3, 5, 18 All 1, 5, 7, 13, 18 5, 13, 18 1, 5, 10, 13, 18 1, 5, 13, 18 1, 3, 5, 18	900 10,000 100 100 1,000 100 100	2 4 2 1 2 2 2
3	1.0	0.2	0.027	0.006	1.500	—	1.129	Portsmouth Providence	All 1, 5, 10, 13, 18	10,000 1,000 1,000 100	4 2 2 1
4	1.0	0.1	0.020	0.008	1.475	—	2.075	Boston	All	10,000 1,000 100	4 2 1
5	1.0	0.2	0.012	0.002	1.903	—	1.043	Portsmouth Providence	All All	9,000 1,000	4 4
6	1.0	0.2	—	—	0.091	—	1.281	Portsmouth Providence	All 10	10,000 1,000	2* 2*
7	1.0	0.2	0.012	0.001	1.359	—	1.090	Portsmouth Providence	All All	9,000 900 1,000 100	4 2* 4 2*

^a All costs are per annum and in millions of dollars. All capabilities are in gallons. Asterisks denote leased equipment.

Table VII
Response Allocation Policy for Run 2 of Table VI^a

<i>i</i>	<i>e</i>	<i>j</i>	<i>k</i>	$r_{iej k}$	<i>i</i>	<i>e</i>	<i>j</i>	<i>k</i>	$r_{iej k}$										
1	2	1	2	900	4	1	5	3	100										
			3	900				13	2	100									
		3	5	2			1	100	5	2	1	3	1,000						
							2	900				5	3	1,000					
				3			3	3			900	10	1	1	100				
								3			900			2	1,000				
				18			4	1			3	7,800	6	2	1	3	100		
											3	10,000				5	3	100	
								5			3	3			7,700	13	2	2	700
												2			1,000			18	2
7	3	3	10,000		18	3		3	1,000										
		3	10,000					1	100										
10	3	3	3		9,900	8		2	1	1	100								
			3		7,700					2	100								
		13	3	3	100		3		3	3	100								
				3	100					5	3	100							
		18	3	3	100		18		3	3	100								
				3	100					3	2	100							
3	2	1	3	100	5	3	3	2	100										
			3	100				5	3	100									
		5	1	1			100	18	3	3	100								
				2			100			1	100								
		7	1	1			100	3	3	3	100								
				2			100			2	100								
		13	2	2			100	3	3	3	100								
				3			100			5	3	100							
18	3	3	100	18	3	3	100												
		3	100			3	3	100											

^a *i*, Equipment location index (see Table II); *e*, equipment type index (see Table III); *j*, risk point index (see Table I); *k*, spill volume index (1: 100 gallons, 2: 1,000 gallons, 3: 10,000 gallons). Only nonzero capabilities are displayed (in gallons).

that all but one facility are opened and that the response allocation pattern as well as the technology “mix” are nontrivial (later comments amplify on this issue).

2. Runs 3 and 4 differ from Run 2 in that all facilities involve a fixed opening cost of \$100,000/year, Run 4 also involving a mobilization delay of 6 hours at all facilities. Notice that although the total number of facilities predictably shrinks, we still observe a mix of three technologies stockpiled. Run 5 differs from Run 3 by allowing only equipment of type 4.

3. Runs 6 and 7 examine the “buy versus lease” decision for small oil spills, again under a fixed opening cost of \$100,000 per year everywhere. The basic options are either to acquire technology type 4, or to lease technology 2 at a flat rate of X dollars per volume of capability used. Run 6 assumes $X = \$1/\text{gallon}$ and results in an “all-lease” decision (notice the response allocation) while in Run 7 we have $X = \$20/\text{gallon}$, in which case the model commends a “mixed” policy.

4. Notice from the prior runs that the recommended allocation policies as well as the technological “mixes” are far more intricate for the small spill problem than for the large spill problem. Table VII further elaborates

on the results of Run 2 by displaying the entire response allocation array for that case (only nonzero $r_{iej k}$'s are shown). Trying to identify a pattern from that table, we can see that equipment type 4 (stockpiled at Portsmouth, NH) is used only for the largest of the small spills, despite the fact that even at those volumes it is more expensive to operate than types 1 and 2, which are used for smaller volumes. Notice that type 3, while preferable to type 4 in several large spill scenarios, is absent from all small spill solutions. Geographically, each facility typically responds to all volumes generated at risk points in its vicinity, and may also respond to more distant locations, but only to larger volumes at those locations. Indeed, only one location (usually, but not always, the closest facility) typically responds to each 100-gallon spill. By contrast, more locations (as many as four in this case) may respond to each 1,000-gallon spill, while even more facilities (as many as seven here) may be designated to respond to each 10,000-gallon spill.

4. Discussion and Conclusions

The previous example is *illustrative* and should be interpreted with caution. In that respect, less importance should be attached to the actual numerical

values of the outputs than to the trends and underlying problem structure that these numbers help reveal. Also, we should be cautious in generalizing these results to other geographical areas. For instance, an application of this model to the West Coast might very well lead to different conclusions, since, by contrast to New England, the wind in the West Coast typically blows *toward* the shore—which may increase the “damage potential” of large spills dramatically.

In spite of such caveats, we believe that the results of the New England application of this model provide strong evidence that, over the range of scenarios examined, it is generally more cost-effective to invest in “small-spill” response capability (i.e., $\leq 50,000$ gallons) than to acquire expensive large-scale equipment that would be used only rarely (and, if used, would probably avert far less damage than its cost could justify). Table V shows that New England would require more than \$12.5 million per year in system costs (most of which in equipment acquisition funds that would have to be disbursed with certainty) to reduce the *expected* value of damages from \$4.7 to about \$2.3 million per year. It might be worthwhile to investigate whether these same funds could be better utilized in a “revolving fund” (or insurance) scheme of *direct* compensation for damages, with minimal (i.e., “surveillance and monitoring”) response. A portion of these same funds could also be used to strengthen the response system for small spills, whose cleanup seems to be far more viable.

A secondary finding of the New England application concerns the siting patterns of response resources for these two categories of spills. We saw that small-spill capability is typically dispersed geographically, while large-spill capability is typically consolidated at one or a few “centers.” Departures from this pattern could occur in alternative scenarios (see Runs 3–7 vis-à-vis Run 2 of the small-spill analysis in Table VI). In an actual implementation of the model, a significant effort should be spent to determine which scenario is the most accurate representation of the real world.

With respect to an appropriate selection of inputs for this model, one of the input variables that this study made no attempt to investigate, but which could be crucial in an actual implementation, is W , the “damage weight.” If W is interpreted as the decision maker’s “willingness to pay” to respond to oil spills instead of directly bearing their damages, the New England analysis seems to indicate that if *both* categories of spills are to be responded to (with no unsatisfied demand), then the willingness to pay for large-spill response in New England ought to be at least one order of magnitude greater than the equivalent figure for small-spill response.

How do all these findings, which are, to a significant extent, area specific, compare with current practice? In the United States, the Coast Guard is already charged by the Federal Water Pollution Control Act and other laws with the responsibility of responding to small and large spill incidents. However, it happens that the mechanisms by which the Coast Guard responds to small and large spills are quite different. Indeed, large spills typically require the mobilization of the USCG’s National Strike Force (involving the use of heavy-duty, federally owned equipment). By contrast, regional and local resources (very often using private cleanup contractor equipment) typically respond to small spills. It is fair to say that so far there seems to be no evidence that public opinion, or the government, realizes the potential importance of small spills relative to large ones. Large spills typically capture the lion’s share of headlines, concern, and public debate regarding what should be done, and, probably as a result, influence the ultimate determination of society’s “willingness to pay” for a response system for these spills. In that respect, we are aware of no recent major strategic response study sponsored by the government that has focused on the *small-spill* end of the “response spectrum” (ours is an exception). To our knowledge, the most recent development in this context has been the USCG’s solicitation and sponsoring of a new response study, this time to deliver an “extraordinary spill” strategic plan for the Continental United States, Alaska and Puerto Rico. The USCG has defined “extraordinary spills” as being “sudden catastrophic discharges, prolonged major discharges, and otherwise unique incidents for which guidance beyond what is provided in the National Oil and Hazardous Substances Contingency Plan (NCP) is required” (USCG 1984). In light of the results outlined in Section 3, it might be worthwhile to investigate whether (and under what conditions) the patterns identified in the New England illustration are also manifested under other scenarios, in other geographical areas, or at the national level.

The purpose of the previous analysis and discussion has been to give a flavor of the potential of this model as a flexible tool for analysis of response options available to policymakers. Actually, quite a few additional insights on this problem can be obtained by appropriate uses of the model. Among other things, one can investigate the benefits of dispatching equipment by alternative transportation modes, the benefits of using chemical dispersants (or other cleanup technologies such as the “Sombrero” device for platform drilling blowouts), the incremental pollution costs of potential new offshore drilling activities, and many other possibilities.

We conclude with a few comments on directions for further improving this model. Considerable progress has been made in Ziogas and in Psaraftis and Ziogas in modeling the tactical/operational oil spill response decision process. It would be worthwhile to attempt to use these concepts to further refine the representation of spill response *tactics* in the *strategic* problem *without* a significant increase in computational effort. Another area for further study concerns the treatment of risk aversion. It would be interesting to try to incorporate some model of society's risk preference structure into this model in a more explicit way than via response time stipulations. The work of Fredrikson, although not associated with locational considerations, might prove useful in such an extension.

Acknowledgment

Support for the MIT Oil Spill Model has been provided by the National Oceanic and Atmospheric Administration's Office of Sea Grant, the U.S. Coast Guard, the U.S. Navy, the Commonwealth of Massachusetts, the Doherty Foundation, JBF Scientific Corporation, the Spill Control Association of America, Texaco and Petro-Canada. Opinions expressed in this paper are solely those of the authors. The computations of the paper have been supported by the Department of Ocean Engineering at MIT. We would like to thank several anonymous individuals for their comments on earlier versions of this paper and Dan Nyhart for his comments and overall assistance.

References

- ANTHONY, R. N. 1965. *Planning and Control Systems: A Framework for Analysis*. Harvard Business School, Cambridge, Mass.
- BAIRD, A. V., K. JOHNSON, R. CORWIN AND P. CSIK. 1982. Oil Spill Cleanup Economic and Regulatory Models: Volume 4. The Trajectory and Damage Assessment Models, MIT Sea Grant Report (June).
- BELARDO, S., K. R. KARWAN AND W. A. WALLACE. 1984a. Managing the Response to Disasters Using Microcomputers. *Interfaces* **14**, 29-39.
- BELARDO, S., J. HARRALD, W. A. WALLACE AND J. WARD. 1984b. A Partial Covering Approach to Siting Response Resources for Major Maritime Oil Spills. *Mgmt. Sci.* **30**, 1184-1196.
- CARTER, J. 1977. Oil Pollution of the Oceans. President's Message to Congress on Measures to Control the Problem (March).
- CHARNES, A., W. W. COOPER, J. HARRALD, K. R. KARWAN AND W. A. WALLACE. 1976. A Goal Interval Programming Model for Resource Allocation in a Marine Environmental Protection Program. *J. Environ. Econ. Mgmt.* **3**, 347-362.
- CHARNES, A., W. W. COOPER, K. R. KARWAN AND W. A. WALLACE. 1979. A Chance Constrained, Goal Programming Model to Evaluate Response Resources for Marine Pollution Disasters. *J. Environ. Econ. Mgmt.* **6**, 244-234.
- DEMIS, D. J. 1984. Oil Spill Management: The Damage Assessment Model and the Spatial Allocation of Cleanup Equipment. Unpublished S.M. thesis, MIT Department of Ocean Engineering (June).
- DEVANNEY, J. W., III AND R. J. STEWART. 1974. Bayesian Analysis of Oil Spill Statistics. *Marine Technol.* **11**, 365-382.
- FREDRIKSON, G. F. W. 1983. Cost/Benefit Analysis of Norwegian Oil Spill Preparedness. Report 82 02 22-2, Central Institute for Industrial Research, Oslo (October).
- PSARAFTIS, H. N., AND B. O. ZIOGAS. 1982. A Tactical Decision Algorithm for the Optimal Dispatching of Oil Spill Cleanup Equipment. *Mgmt. Sci.* **31**, 1475-1491.
- PSARAFTIS, H. N., J. D. NYHART, D. A. BETTS, 1983. First Experiences with the Massachusetts Institute of Technology Oil Spill Model. Proceedings, Oil Spill Conference, San Antonio, Tex., pp. 301-305.
- THARAKAN, G. G. 1982. The Complementary Locations on a Network. Unpublished Sc.D. dissertation, MIT Department of Ocean Engineering (May).
- TSC. 1979. Deployment Requirements for U.S. Coast Guard Pollution Response Equipment. Transportation Systems Center, Report to the U.S. Coast Guard.
- USCG. 1982. Polluting Incidents In and Around U.S. Waters: Calendar Year 1981 and 1982. U.S. Coast Guard, Commandant Instruction M16450.2E.
- USCG. 1984. Solicitation DTCG23-84/R/40039 (on "Extraordinary Spill Plan"), U.S. Coast Guard (July).
- WHITE, I. C., AND J. A. NICHOLS. 1983. The Cost of Oil Spills. Proceedings, Oil Spill Conference, San Antonio, Tex., pp. 521-524.
- ZIOGAS, B. O. 1982. Management of Oil Spill Cleanup. Unpublished Ph.D. dissertation, MIT Department of Ocean Engineering (February).