The consequences of alternative environmental management goals: A ...
Jones, Donald W;Redus, Kenneth S;Bjornstad, David J
Environmental Modeling and Assessment; 0 Jan. 2000; 5, 1; SciTech Premium Collection
pg. 1

Environmental Modeling and Assessment 5 (2000) 1–17

The consequences of alternative environmental management goals: A non-linear programming analysis of nuclear weapons legacy clean-up at Oak Ridge National Laboratory

Donald W. Jones[a], Kenneth S. Redus[b] and David J. Bjornstad[a]

[a] *Energy Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831-6205, USA*
E-mail: jonesdw@ornl.gov
[b] *MACTEC, Inc., Oak Ridge, TN, USA*

Received December 1998; revised April 1999

Prioritization of projects within the U.S. Department of Energy's (DOE) Weapons Complex Clean-up Program, exemplified with data from the Oak Ridge National Laboratory, is quite sensitive to overall goals. Non-linear programming analysis of three alternative goals – mortgage reduction, terminal-period risk minimization, and current-period risk minimization – shows substantial differences in waste treated, risk reduced, and cost over a ten-year period.

Keywords: nuclear waste, prioritization, risk management

1. Introduction

The U.S. Department of Energy's (DOE) environmental program is one of the largest in the world, with over 130 sites across 30 states and territories. Within its environmental management program, DOE currently is spending some $5.5 billion per year cleaning up facilities and restoring land contaminated during the nuclear weapons production of the Cold War era. The effort is governed by U.S. environmental laws as enforced by various agencies with overlapping jurisdictions, funded by a less-than-enthusiastic U.S. Congress, and urged to higher levels of clean-up and greater expenditures by an array of parties, ranging from contractors and federal employees to state and local officials, all of whom have employment interests.

The program has emerged relatively quickly, and its dimensions and complexity are overwhelming. The engineering origins of both the problems and the proposed solutions have encouraged a bottom–up approach to thinking about the structure of the clean-up effort, but the cost and time implications of that conceptualization have not fared well in the post-Cold War U.S. federal budget politics. This paper looks closely at the cost, risk, and activity structure of the DOE clean-up program at one of the Department's sites, the Oak Ridge National Laboratory, and examines the implications of pursuing each of three alternative goals for the overall, national clean-up program. In the first part of this introduction, we offer some background on the emergence of this program, and in the second we preview the remainder of the paper.

1.1. The evolution of the U.S. Department of Energy's Weapons Complex Clean-up Program

DOE's Weapons Complex Clean-up Program was set in motion by a lawsuit instituted by the Legal Environmental Assistance Foundation (*Legal Environmental Assistance Foundation v. Hodel*, 586 F. Supp. 1163 (E.D., TN, 1984)), which made DOE subject to the provisions of the Resource Conservation and Recovery Act, dictating extensive cleanup. DOE created its Office of Environmental Management (DOE/EM) in 1989. In 1995, after several years of study, DOE/EM published its initial clean-up plan [8], which suggested a $230–360 billion effort over 75 years based on engineering analyses of waste groups and technologies. A subsequent baseline cost estimate of $227 billion was published a year later, with the difference attributed to assumptions about increased productivity, reduced scope, modification of compliance agreements and cost reductions from DOE privatization [11]. The lengthy approach to this environmental management program was radically revised by the so-called "Ten-Year Plan" introduced by Assistant Secretary of Energy Alvin Alm in July, published in December, of 1996 [10]. On the belief that a 75-year plan was too long to be credible, Alm's idea was to map out what could be accomplished within a ten-year period of concerted effort and budgetary support of the U.S. Congress, of around $170–200 billion, converting what was originally an engineering problem of finding and executing technological solutions to environmental problems into an economic problem in which ends far outstripped means and time available.

Within the Ten-Year Plan and its successor, the National 2006 Plan [12,13], prioritization of individual projects within the program became a critical issue, but while broad policy guidance was provided by DOE Headquarters, the details of prioritization were left to individual DOE field offices [10, pp. 29–30]. The prioritization guidelines included eliminating the most urgent risks, reducing "mortgage" costs to free up funds for further risk reduction, and maintaining a safe working environment for clean-up personnel as three of the highest considerations in a longer list. "Mortgage reduction" in particular remained an attractive, if somewhat imprecise, goal which was interpreted as paying to eliminate some waste groups early to avoid continuing storage, surveillance, and monitoring costs on them. The departmental and public discussion of the clean-up goals and prioritization did not crystallize the overall problem as one of maximizing efficiency in environmental risk-reduction, and the cost and risk-reduction implications of the alternative goals and related priorities remained unexplored.

1.2. Roadmap to the study

This study examines DOE's ten-year, weapons-complex clean-up program as an optimization problem in which program managers try to either maximize or minimize the "amount" of the goal they can accomplish in ten years with a given annual budget, by moving waste out of storage, treating it, and moving it to final disposal.[1] We study the influence of three alternative goals, two of which have been widely discussed but not formalized as official goals, the other of which, we believe, is new. The new goal we call "minimize terminal-period risk", which means that we want to finish the ten-year period with as little risk as is possible to carry over to the post-clean-up period.[2] We define this risk more fully in the text below. The two familiar goals are "minimize mortgage cost" (or "maximize mortgage-cost reduction") and minimize all risks during all ten operational periods.

We rely on production functions to characterize the costs of storage, treatment, and disposal. These functions define how much of each activity can be accomplished with specific quantities of labor and materials (the inputs to the production functions). Since the inputs have given prices (e.g., the going wage/salary structure for labor), these functions tell how much activity can be completed for a given budget. We also use production functions to characterize the relationship between wastes and risks in each of the three states and to describe how the application of inputs can be used to contain the risk from any given quantity of waste.

Data from the Oak Ridge 2006 Plan for the Oak Ridge X-10 facility (Oak Ridge National Laboratory (ORNL)) operationalize the model – i.e., solve the model numerically for each of the goals.[3] A solution for each scenario gives the time path over the ten-year operational period of waste (by type: Low-Level Waste (LLW), Mixed Low-Level Waste (MLLW), TRansUranics (TRU)) in storage, in treatment, and in disposal, the time paths of risks and expenditures, the terminal-period risk, and the cost per period of maintaining remaining wastes after the end of the ten-year effort.

Section 2 describes the structures of the models used to assess the different clean-up goals. Section 3 describes how the data from Oak Ridge are used to develop numerical values to be used for the various symbols of the model equations. Section 4 narrates the stories the optimization models reveal when we solve for the optimal waste management plan for each of the three goals. Section 5 summarizes and compares the results of these stories.

2. The optimization model of the ten-year clean-up effort

Section 2.1 outlines the structure of the problems: the initial situation, the choices available, the "manager's" constraints, and the general objectives. The technical narration of the model begins with the cost structure of the models in section 2.2. Section 2.3 describes the three alternative goals as separate objective functions. Each of the analyses relies on the same cost structure and investigates how the different goals induce different choices regarding the magnitudes of the instrument variables. The objective functions involve either risks or costs. The instrument variables are the quantities of MLLW, LLW, and TRU waste to be "treated" in each of the ten years of the operating period. The constraints always include a budget, sometimes allowable unit risks on waste groups during different operations, and sometimes lowest possible values of instrument variables.

---

[1] Tucker et al. [7] and Toland et al. [6] also approach prioritization of environmental management options with mathematical programming, the latter also focusing on DOE's nuclear waste clean-up program. Toland et al. analyze the least-cost technological approach for remediating a single waste group at a single site (Operational Unit #1 at DOE's Fernald Plant, Ohio, whereas we compare three alternative programmatic motivations for managing an entire site's waste groups on the assumption that technical choices reflected in site management plans already have been optimized.

[2] We realize that DOE has no intention of simply "walking away" from its environmental management responsibilities at the end of the 2006 Plan, but a reasonable interpretation of the skeletal structure of the Plan is that Congress and the taxpayers can be expected to have limited patience for an unlimited clean-up plan, so the closest reconciliation of the environmental and political realities is to make an extraordinary effort to get as much of the clean-up done in a relatively short period (hence the ten-year period), then move into a long-term management strategy more heavily oriented to surveillance and monitoring than to moving and burning dirt. The planning strategy embodied in our analysis emphasizes this structural feature of the 2006 Plan.

[3] The appendix discusses the initial form of these data, originally reported in the 2006 Plan for the DOE/Oak Ridge Complex, and the operations required to conform them to the conceptual structure of the optimization models. The appendix to Jones et al. [3] describes these data in full detail, to the project level, but the appendix to this paper offers a sufficient summary to permit the reader to see how the analytical results of our optimization modeling have emerged from the data that are being used to guide the Oak Ridge operationalization of the 2006 Plan.

2.1. The structure of the optimal clean-up problems

The initial situation. At time zero, there are 305 million cubic meters of MLLW, 60 billion cubic meters of LLW, and 42 million cubic meters of TRU waste under containment ("in storage" in the terminology of the model) at ORNL. Altogether, these wastes generate 216 thousand dimensionless units of risk as measured by the RDS/MEM (Risk Data Sheet/Management Evaluation Matrix) methodology. This is the full array of wastes eligible for treatment during the 2006 Plan.[4]

The choices. The "manager" behind the model's optimal program has three principal choices to make for each of ten time periods: the volumes of MLLW, LLW, and TRU waste to move from "storage" to treatment[5] and then to "disposal".[6] This optimal plan is determined in period zero and is implemented on schedule throughout the ten operational years. To implement these choices, the manager assigns certain amounts of a composite input which we call "labor" as a shorthand to various direct and supporting activities. The three states in which quantities of each waste may reside are tied together such that choices about treatment volumes imply choices about volumes remaining in "storage" and "disposal" at the end of the period in which the treatment occurs. Since maintaining the wastes in storage and disposal also requires resources in the form of the "labor" input, the quantities in storage and disposal put some limits on how much "labor" is available to conduct treatment in each period. Cost functions, developed from production functions, allow the manager to calculate how much "labor" will be required to deal with each waste type in each state. We do not impose a stochastic element on these cost–accomplishment relationships.

Constraints. While in some sense the manager's principal attention is directed toward treatment, management of the risk associated with each waste type in each state also requires application of the "labor" resource. As we specify the activity–cost relationships, the manager can identify separately the "labor" required to manage a particular volume of each waste type in each of the three states and that required to manage, or contain, the risk associated with the same waste type-volume-state combination. Cost functions derived from production functions also govern the relationship between waste volumes and "labor" required to contain risk to a particular level in each state. Thus, when the manager calculates an optimal treatment plan, he or she has to factor in how much "labor" will be required to keep the risks to desired levels for all wastes, states, and time periods for exactly that treatment schedule.

Because most risks are not subject to direct goals in two of the three plans, the manager would have little incentive, within the structure of the model framework, to worry about containing risk to any particular level. However, a myriad of orders, regulations, and laws govern many dimensions of allowable risks of various types, and we capture this feature of the legal environment with risk constraints. The risk per unit volume of each waste type, in each state and time period, cannot exceed a specified level. Except when an explicit minimization objective focuses on one or more of these sources of risk, these constraints are binding in the sense that the manager always lets the actual level of risk float up to its constraint level, although nothing except more pressing interests (financial exigencies) elsewhere makes him do this.

Budget and spending rules. The manager knows with certainty that \$203 million will be available for each time period in the ten-year plan. The price of a unit of "labor" also is known, so the quantity of those resources purchased can be calculated directly. The manager cannot shift funds between periods; "use 'em or lose 'em" and "no borrowing" are the rules. We work with "real" costs in 1997 prices, i.e., abstracting from possible inflation.[7]

The goals. The first goal we call terminal-period risk minimization. In this problem, the manager's goal is to choose a ten-year treatment profile for the three waste types that leaves the site with the least continuing risk to be managed during the indefinite, ensuing period of stewardship. The constraints on risk during each treatment period restricts the manager's ability to trade off more risk in earlier periods for lower risks in the terminal period. We offer more details about the continuing, or "terminal-period" risk in the subsection. The second goal we offer our manager we call "minimize current-period risks in all periods". This amounts to devoting whatever resources are needed to min-

---

[4] Identifying this array of wastes as all there is to be treated abstracts, of course, from the continual trickle or stream of new finds that come to the attention of managers. There is no inventory of such to-be-discovered, uncontained wastes, but clearly the dangers posed by some of these new discoveries may, in practice, pre-empt undertaking some pre-defined projects identified at the beginning of a ten-year planning period. We do not attempt to capture this additional dimension of complexity of DOE's environmental management task.

[5] There is, at present, no "treatment" for TRU waste. This state, applied to TRU, should be interpreted as preparing for shipment to the Waste Isolation Pilot Plant (WIPP) and shipment.

[6] "Disposal" covers a number of temporary or final destinations of treated wastes. We do not intend to imply any sort of free release of treated wastes, but rather some post-treatment disposition that will involve some surveillance, monitoring, and maintenance for an extended period, but at lower cost per unit volume per time period than prior to treatment.

[7] We also do not discount costs during the active ten-year portion of the clean-up program, although in the mortgage minimization program we do discount continuing stewardship costs in the post-clean-up period. There are two reasons for this choice to not discount: (1) discounting costs without defining a dollar-denominated benefit against which to compare discounted costs would result in simply pushing active treatment as far into the future as possible to "reduce costs", which we believe is nonsensical; (2) both Congress and DOE officials work with undiscounted budget figures when conducting their planning, and presenting them with discounted estimates of what funds they will have to raise (the one from the taxpayers, the other from the Congress) would be more perplexing than informative.

imize risks for all three waste types in storage, treatment (if any is conducted), and disposal (if any waste ever gets there) in each of the ten time periods. In this case, the risk constraints will not bind. It might seem likely that no treatment would ever get accomplished under such a goal, but for two factors: we use a minimum, non-zero treatment constraint in each time period for each waste type; and the structure of cost allocation between volume management and risk management for each waste type (in each state) exhausts all risk-reduction opportunities before the budget is exhausted. The final goal is mortgage minimization: treatment of such volumes of wastes during the ten-year operational period so to minimize the present discounted value of all costs: treatment costs during the ten-year plan period, plus storage and disposal costs during the first ten-years and well into the extended future. The risk constraints pose real limitations on minimizing this objective function, as risk per se is not "valued" according to this goal. Mortgage reduction has experienced many expressions in the DOE complex in the past several years, but any of them that are economically rational must be equivalent to this simple definition.

2.2. Technical description: Cost structure of the model

Every action the manager can take in this analysis costs him something, whether that is moving waste briskly out of storage, into treatment, then on to disposal, or just sitting and watching the waste in storage for ten years. The actions are the bases of the costs and are described with production functions.

We structure each activity in the model as a production process. Intuitively, more treatment of waste is produced by a larger quantity of the "labor" input. Possibly less intuitively but using the same logic, risk is "produced" by larger quantities of waste in any state, but it can be reduced (or contained) by applying "labor". Similarly with storage and disposal. Since each production function uses some units of the composite "labor" input, and we know how much a unit of the "labor" input costs, knowledge of a production function gives us complete knowledge of a cost function [16, pp. 28–34]. For this reason, we begin the description of the cost structure of the model with the production functions for each of the activities.

2.2.1. Production functions for the "states" through which waste can be moved

The production functions describe the amount of the variable, composite input k required to maintain a given volume of a particular waste group in any particular state, where the three states are, in sequential order, storage (S), treatment (T), and disposal (D).[8] The production function for the volume of waste of type i ($i =$ LLW, MLLW, TRU) stored in time t (subscript S indicates storage) is given by

$$S_{it} = A_{iS} k_{itS}^{\alpha_{iS}},$$

which can be inverted to solve for the amount of the variable input k required to store that volume of waste of type i, in time t:

$$k_{itS} = (S_{it}/A_{iS})^{1/\alpha_{iS}}.$$

The amount of waste of type i that can be treated in period t is determined by the amount of the variable input k devoted to treatment, and it is also affected by the volume in storage in the previous period. This latter part of the relationship captures the well-known effect that it is "easier" to take a given volume of waste out of storage and treat it when there is still a considerable amount in storage; subsequent amounts get more difficult to retrieve.[9] Thus, this formulation of treatment also includes retrieval. As with the quantity of the composite input k required for storage, the quantity of k required for treatment (subscript T indicates treatment) can be determined by inverting the production function, although we do not show that solution:

$$T_{it} = A_{iT} k_{itT}^{\alpha_{iT}} S_{i(t-1)}^{\beta_{iT}}.$$

The volume of waste of type i in disposal (subscript D indicates disposal) also requires some quantity of the composite input k, and as with the quantity of k devoted to the volume of waste in storage and treatment, the T equation can be inverted to solve for the number of units of k required to mind the volume of waste in disposal:

$$D_{it} = A_{iD} k_{itD}^{\alpha_{iD}}.$$

2.2.2. Production functions for risk

Maintaining risks associated with a given waste type in any particular state also requires the use of the composite input k. The theory underlying the risk "production" equations focuses on the risk per unit volume of waste. The total risk associated with a particular volume of waste would increase as the volume increases, but a theory of that total risk also needs to specify how the risk per unit of waste is affected by applying the composite input to risk containment. The model of risk containment holds that risk per unit volume of a particular type of waste, in a particular state (storage, treatment or disposal), can be held down or reduced by applying more units of k to a given volume

---

[8] For the relationships between physical calculations for engineering processes and the construction of economic production functions for those same processes, see [1,4], [5, chapter 2].

[9] Waste retrieval operations are normally multi-stage operations, conducted in batch processes. Typically, smaller quantities of material are taken out in subsequent stages, with the unit removal cost rising accordingly. Additionally, as the stock of a single waste group is drawn down out of, say, an underground storage tank, the mixing characteristics of the waste material – liquid, slurry, and sludge – will change, requiring more costly techniques for removal. In some storage formats, ease of access of equipment to the material may decrease as the quantity of material gets smaller. Each of these effects contributes to the inverse relationship between removal cost and remaining quantity of waste in storage. A smoothly increasing cost function is an approximation to what may in fact typically be step functions with many steps.

of waste. As the total volume of the waste increases, it gets more difficult to maintain the risk per unit volume at a given level; to do so would require applying some more units of k. Note that the negative exponent α on the k in the risk equations indicates that more units of k *reduce* risk per unit volume. The risk production equations shown below incorporate the model of risk per unit volume but show the relationship between total risk for an entire stock of waste of type i, the number of units of k used to contain it, and the total volume of the stock of waste. (The equation for risk per unit of waste has an exponent of β on the volume variable (S, T, or D), while the equation for total risk has an exponent of $\beta + 1$: the risk per unit volume equation has on the left-hand side the total risk, R, divided by the total volume (S, T, or D); multiplying both sides of the equation by the volume eliminates the volume variable from the left-hand side of the equation and adds the 1 to the exponent of the volume variable on the right-hand side.) The expression for total risk posed by the volume of waste of type i in storage at time t is, where the subscript R indicates constant terms and composite labor devoted to managing risk (as contrasted to waste volume):

$$R_{itS} = A_{iRS} k_{itRS}^{-\alpha_{iRS}} S_{it}^{\beta_{iRS}+1}.$$

The total risk posed by putting some part of the stored waste through "treatment" is shown by

$$R_{itT} = A_{iRT} k_{itRT}^{-\alpha_{iRT}} T_{it}^{\beta_{iRT}+1}.$$

Finally, the risk of treated waste held in disposal is

$$R_{itD} = A_{iRD} k_{itRD}^{-\alpha_{iRD}} D_{it}^{\beta_{iRD}+1}.$$

In each of these cases, the total risk and the risk per unit volume can be controlled by the choice of k, given the volume of the waste in that state. Consequently, in any period, the total risk of waste in all states is affected by the volumes in each state and the risk choices made for each state.

2.2.3. The relationships between waste in different states

In the initial period of the ten-year treatment program, we set the volume of waste of each type in storage to S_{i0}, and the volumes going through treatment and in disposal to zero: $T_{i0} = 0$ and $D_{i0} = 0$. In period 1, a particular quantity of waste of type i is removed from storage, treated, and put into disposal:

$$S_{i0} - T_{i1} = S_{i1} \quad \text{and} \quad D_{i0} + T_{i1} = D_{i1}.$$

Thus the choice of the volume of waste type i to be treated in time t determines both the amount remaining in storage and the amount in disposal. In general, these relations are

$$S_{i(t-1)} - T_t = S_{it} \quad \text{and} \quad D_{i(t-1)} + T_{it} = D_{it}$$

for every time period t.

These are the activity/production and cost relations and the relationships between states. These relationships form the building blocks of both goals to be maximized or minimized and the constraints under which those goals are to be optimized. We turn next to the mathematical statements of those goals.

2.3. Mathematical expression of the goals

Minimize terminal-period risk. Terminal-period risk minimization as a goal operates under the assumption – not entirely accurate, but certainly a satisfactory motivation for an analysis of waste management under the 2006 Plan – that waste will be treated only for the ten-year duration of this plan; after that, the waste volumes remaining in storage and disposal in time period 10 will remain there indefinitely, posing the same risk in each future period that they posed in period 10.[10] Consequently, the risk of the wastes in storage and disposal in period 10 is of unique interest: it is the risk that will continue "forever" at this site. The goal of DOE's EM program under this scenario is to minimize this continuing risk. Thus the objective function for this scenario is

$$\underset{\{T_{ijt}\}}{\text{minimize}}\, \overline{R}_{10} = \sum_{i} \sum_{j=S,D} R_{ij10}.$$

The manager of the facility picks values of X for each waste type and each time period (a total of 30 Ts) to minimize these terminal-period risks. Each X implies a value of k required for volume management, but nothing so far determines how many units of k will be allocated to risk management. This minimization is conducted, therefore, subject to constraints on the maximum allowable risk per unit volume for each waste type:

$$\lambda_{ijt}(R_{iS0}/S_{i0}) \geqslant (R/V)_{ijt} \quad \forall i, j, t,$$

where V_j represents the volume in state j – storage (S), treatment (T), or disposal (D): the allowable risk per unit for each waste type must stay below, or at least be no greater than, λ times the level of unit risk in storage in the initial period – period zero, not period 1. The values of λ chosen for these constraints can represent legal constraints or environmental management policy decisions, or a combination of both. With this constraint, the production equation for the unit risk in each waste type, on the right-hand side of the inequality, determines the number of units of k required to maintain the required unit risk level in each period. The values of λ could change over time, falling values, for instance, representing a mandate of continuously improving safety. Without the constraint on risk, no units of k would be allocated to risk in any activity except in storage and disposal in period 10.

---

[10] This is an abstraction, of course. Some risks will naturally decline over time; others may or will increase, either from greater likelihood of public encroachment or from deterioration of containment systems. To highlight the risk reduction produced by the ten-year clean-up program, we assume that the extent of risk reduction achieved within that period remains constant thereafter.

Next, the manager must buy, rent, lease, or somehow procure, the units of k with which to treat waste, manage waste volumes in storage and disposal, and manage risk in all states. Of course, the budget to support these acquisitions is not unlimited, and the minimization is subject additionally to a budget constraint in each period, B_t:

$$B_t \geqslant \sum_i \sum_{j=S,T,D} \sum_{a=V,R} p_k k_{ijat} \quad \forall t,$$

in which the subscript a represents the type of activity – either volume management or risk management. Thus, the sum of the expenditures on units of k has to stay within the budget for each period. There is no ability to trade budget among time periods.

The equations for the objective function, the risk constraints, and the budget constraints represent the manager's optimization problem for minimizing terminal-period risk. The production and risk equations presented earlier determine how many units of k are required to manage the volumes and risks.

Minimize all risks in each period. In the second scenario, the management goal is to minimize all risks in each and every period. Since it is possible that no treatment at all will be the risk-minimizing solution, we include a set of constraints that requires that at least a small volume of each waste type be treated in every period, as well as a constraint that at least a small portion of the budget be spent in each period.[11] Again, there is, of course, the budget constraint itself, without inter-period borrowing and lending. The objective function for this goal is

$$\underset{\{X_{ijat}\}}{\text{minimize}}\, \overline{R} = \sum_t \min\left(\sum_i \sum_j \sum_a R_{ijat}\right).$$

Minimize mortgage costs. The third scenario, maximizing mortgage reduction, is a particular form of overall cost minimization. The manager wants to conduct a treatment plan that will require the least amount of expenditure on storage, treatment, and disposal during the ten-year, active treatment period, *and* the discounted sum of tenth-period storage and disposal costs on all the wastes over some extended planning horizon. Essentially, this definition of reducing mortgage costs is finding a treatment plan that will maximize the difference between the costs of holding untreated waste in storage and holding treated waste in disposal, net of treatment costs. Treatment is the cost that must be incurred to get the reduced cost of holding waste in disposal instead of in storage – in a way, it is the price of buying lower future costs. It will take some number of years after the end of the ten-year treatment period for the reduced holding costs on the treated waste to repay the cost of treatment. The objective function for this scenario is

$$\begin{aligned} &\underset{\{X_{ijat}\}}{\text{minimize}}\, M \\ &\quad = \sum_{t=1}^{10} \sum_i \sum_j \sum_a C_{ijat} \\ &\qquad + \sum_{t=11}^{n} \sum_i \sum_j \sum_{a=S,D} C_{ija10}(1+i)^{-t}, \end{aligned}$$

where the C_{ijat} terms are costs of storage, disposal, and treatment. C_{ijat}, $t > 10$, represents the administrative and operational costs of what is called stewardship in the current DOE clean-up lexicon. In other words, the problem is to maximize the unit cost savings of moving waste from storage to disposal, net of treatment costs. The first term in the expression above is the minimization of costs in the ten-year treatment period, while the second term is the discounted costs of storage and disposal, as they were in the final period of the ten-year activity period, over some future period required to recover the treatment costs. The second term is just n times the storage and disposal risk incurred in period 10 (risk in both storage and disposal is calculated at the end of the operational period, after material has been drawn down from storage and added to disposal), where n is the number of years in the evaluation horizon, after the completion of the ten-year treatment period. The trade-off is between short-run cost savings in the first ten years and long-run cost savings over some future period. Constraints on unit risk levels are necessary to get any units of k applied to risk containment since no direct risk objectives exist in the goal. Again, of course, there is the period-by-period budget constraint. The reader will have noticed that we discount the second term, but not the first term. As we hinted above in the informal introduction to the stories of the model, it certainly would have been possible to discount both terms, but we have not discounted costs over the ten-year operational period for convenience of comparison to both DOE/OR budget plans and congressional authorizations. The discounting over ten years – i.e., looking from period zero, forward over the following ten operational years – would not have altered any planning actions in the two risk-related goals, simply because costs were not part of the goals. The current formulation thus gives a slight bias toward overstatement of operational-

[11] As we indicated in the introductory exposition of the model in section 2.3, some readers might find it odd that such a constraint would be binding and others might find it odd if it were not, so we offer some further guidance on how to interpret this constraint. First, the cost structure we derive for risk management/containment activities is not so expensive as to be able to absorb the entire budget without exhausting the opportunities for reducing risks. Second, moving waste from storage *into* treatment increases the risk per unit volume for each waste type, so there is a sort of "hump" to be gotten over at the very beginning of a ten-year program – the hump being the sharply higher costs of managing risk during treatment than during either storage or disposal. Third, the structure of treatment costs themselves has unit cost rise the smaller the stock available to be drawn from to treat, which we discuss further in the text below. Consequently, it would be possible that no treatment be chosen at all, to minimize all sources of risk in each time period individually in the absence of a constraint that requires at least a small amount to be treated regardless what it does to risk in the period. Since risk management by itself would not exhaust the budget, the full budget would not be spent without at least some treatment, possibly more than the constrained minimum.

Table 1
Coefficient values.

Parameter	Storage		Treatment		Disposal	
	Volume manag.	Risk manag.	Volume manag.	Risk manag.	Volume manag.	Risk manag.
$A_{\text{TRU, state, activity}}$	2.00E+7	7.33E−4	5.88E−2	7.27E−3	4.80E+5	1.08E−3
$A_{\text{LLW, state, activity}}$	6.07E+10	8.08E−8	1.05E−1	1.05E−6	4.82E+8	2.10E−7
$A_{\text{MLLW, state, activity}}$	1.80E+8	2.02E−5	2.05E−1	8.60E−5	2.30E+8	2.32E−5
$\alpha_{\text{state, activity}}$	0.115	0.100	0.400	0.500	0.990	0.250
$\beta_{\text{state, activity}}$	n.a.	0.100	0.900	0.075	n.a.	0.030

period expenditures relative to future cost savings, but as will be seen in the narrative of the results of this scenario, this bias is minimal.

3. Parameterization

We discuss how we used the 2006 Plan data to parameterize the production functions in the first section below and the characterization of "allowable" risks used in the constraint functions in the second section. Table 1 reports the parameter values used in the numerical investigations.

3.1. The production/cost functions

The 2006 Plan data do not offer all the information required to parameterize the production functions of the optimization models, but the economic theory of production, as embodied in the mathematical production functions themselves, supplements the information content of those data. The production (cost) parameters for which numerical estimates are needed are the constant terms, the A_{ij}, and the output elasticities, the α_{ij} and β_{ij} coefficients. The A coefficients scale the magnitudes of the right-hand-side variables (the inputs) to the left-hand-side variables (the outputs). There are substantial differences in the orders of magnitude among the composite labor variables (around 1,600 units in total, across all the uses, depending on the budget), the risk variables (in the thousands to tens-of-thousands range), and the waste volume variables (which range from the millions to the billions).

The α and β coefficients form the heart of the cost information because they identify the percent change induced in the right-hand-side variable by a one-percent change in one of the left-hand-side variables.[12] Accordingly, we chose the sums of the output elasticities (or the value of the single output elasticity in the cases of production functions with only one input) to characterize the degree of returns to scale we believed to characterize the activities. The value of each output elasticity is, under most circumstances, close to the associated input's share in the costs of the activity.[13] For example, with two inputs, output elasticities of 0.4 for input 1 and 0.6 for input 2 indicate that 40% of both the unit cost of the output and the total cost of all the units of output are contributed by the cost of input 1 and 60% by input 2.

From the 2006 Plan data base we can extract information on the costs expected to be associated with each waste type and can express those costs on a per-unit-volume basis. Those data do not give precise information on the costs in each separate state for each waste type, but at that point, we can supplement the Plan data with experiential engineering judgment on relative costs in each activity. We can anticipate that treatment (or "treatment", in the case of TRU) costs per unit volume, for each waste type are greater than either storage or disposal costs, and also that "disposal" costs are lower than "storage" costs (see [2] for the exposure experience in shutting down a chemical plant for maintenance, involving procedures we believe are roughly comparable to those involved in treatment in the DOE program); otherwise there would be little sense in the mortgage-reduction plan of moving waste out of its current state into some other state. The fact that the unit risk posed by "untreated" waste exceeds the unit risk posed by "treated" and "disposed" waste is *prima facie* evidence that the unit cost of "disposal" is lower than the unit cost of "storage" for each waste type. With the combination of the information yielded directly from the Plan data and the experiential engineering judgment on relative unit costs across states, we can make credible estimates of the A, α, and β coefficients for each of the volume management production functions.

To determine the corresponding parameter values for the risk production functions, we first estimated the proportions of total costs contributed by volume management and risk

---

[12] The sum of the α and β coefficients in any production function indicates the degree of returns to scale in that activity. If their sum is exactly 1.0, the activity experiences constant returns to scale: a given percent increase in *each* of the right-hand-side variables (the "inputs") causes exactly the same percent increase in the left-hand-side variable (the "output"). A sum of α and β coefficients greater than 1.0 implies increasing returns to scale, and a sum less than 1.0 decreasing returns to scale. If there are more than two inputs, and accordingly there are more coefficients than just α and β, then the sum of all the coefficients indicates the returns to scale. Similarly if there is only one input: the value of, say, α alone indicates the returns to scale.

[13] This relationship holds precisely when the sum of the output elasticities is 1.0. Naturally, when there is only one input and its output elasticity is greater or less than 1.0, its cost share cannot exceed or fall short of 1.0. In cases of non-constant returns to scale (either increasing or decreasing), the cost share of each input is closely related to the ratio of its input elasticity to the sum of all the input elasticities. Thus, the single input with an output elasticity of, say, 0.9 has a cost share of 0.9 divided by 0.9, or exactly 1.0, which is exactly what intuition indicates the cost share obviously should be, since there are not other inputs contributing to costs!

management. With this estimate, which we specified as a 90–10% split (90% to volume management, 10% to risk management), we were able to back-calculate values for each of the A, α, and β coefficients for all the production functions, in the manner described in the following paragraph. In the solutions to the optimization problems, the 90–10% split of costs between volume and risk management does not impose a hard-and-fast allocation of the composite labor inputs to the two types of activity, but it places a tendency for the relative costs to remain in the neighborhood of that ratio.[14]

The values of the A, α, and β coefficients for the storage production functions were determined as solutions to the production equations for the period-zero values of stocks and risks in storage, given the base budget and unit cost of the composite variable. Since no wastes are in either treatment or disposal in period zero, we used a slightly different strategy to back-calculate the coefficient values for volume and risk management for each waste type in those two states. For the volume production functions, we calculated the values of the constant terms using a hypothetical 10% of the initial stock in storage as a "typical" throughput in any period; for the risk production functions in treatment and disposal states, we used engineering judgment on relative risks per unit volume of each waste type, allowing for differential worker exposure, environmental hazard, and danger of the waste form across the three states to distribute unit risk per waste volume. That is, the risk production functions were parameterized on the basis of risk per unit volume of waste in each state.

3.2. The risk constraints

The values of the λ coefficients, which were used in the risk constraints to specify the maximum allowable unit risk level in each state are, for storage, treatment, and disposal, 1.00, 2.00, and 0.30. They represent compliance with legislation, regulations and DOE orders. These values were the same for each waste type and remained constant over the ten-year treatment period, although the cost of attaining these risk levels differs across the three waste types. Thus, the unit risk of material in storage could rise no higher than it was in the initial period; the unit risk of each waste type could rise to double its value in storage during treatment, reflecting the greater exposure of workers to the material and its more active handling; and, once disposed, the material had to remain at 30% of its initial unit risk level in storage.

4. The stories of different goals: Narrative of the results of the optimization models

The highlights of these results are the time paths of MLLW, LLW, and TRU waste volumes treated over the ten-year period, the time paths of risks over the same duration, and the value of "terminal-period risk", but many, more fine-grained numbers are yielded by the analyses at the same time, and we discuss these as they illuminate particular issues. Sections 4.1–4.3 report the stories yielded by solution of the model using each of the objective functions. Section 4.4 brings together and compares the optimal environmental management programs yielded by the three different goals. In each program, we begin with an annual budget of $203 million per year and a price per unit, per year, of the composite variable input k of $127 thousand. The total risk of all three waste types in storage in period zero is 216,000.

4.1. Program 1: Minimize terminal-period risk

The problem in this case is to find the treatment plan that will yield the lowest sum of risks in storage and disposal for all three waste types at the end of the tenth year. Unless all units of one or more waste groups are treated during the ten-year activity period, some units will remain in storage at the end of the plan period while others, having been treated, are in disposal. A budget of $2.03 billion for the entire period is insufficient to treat all wastes, but by the end of the ten years, 90.4% of the TRU, 76.7% of the MLLW, and 66.7% of the LLW have been treated and moved to disposal. The optimal treatment plan pulls out TRU early, maximizing the amount treated in the second period, and treating none in the tenth period. Relatively large percentages of the initial stock of MLLW also are treated in the earlier periods of the program, and as with TRU, none is treated in period 10. The reason for zero treatment of TRU and MLLW in period 10 is that their treatment costs draw money away from risk containment in storage and disposal, which are the focus of the optimization. The more resources available to contain risk in period 10, the lower will be the terminal-period (period 10) risk. With LLW, on the other hand, with its vast volume and low risk per unit, an optimal program consists of a uniform treatment of 6.67% of the initial-period stock in each period, including period 10.

Figure 1 shows the reduction in storage costs permitting the increase in treatment costs which lasts for the duration of the planning period. The budget constraint is binding in all periods except the last, when a small amount of budget goes unused. The unit risk constraints bind in each period on each waste type and in each state, as expected; i.e.,

---

[14] The consequences of this solution we adopted to fix the parameter values reveal themselves in the solution to problem 2, the minimization of all risks in each period: a smaller proportion of the total, ten-year budget is spent than may appear intuitively reasonable to some readers, because the 90–10% cost split causes an exhaustion of risk-reducing opportunities with expenditure of considerably less than the full budget. Intuitively, some readers might think that since the funds are there, no reasonable site manager is going to fail to spend them all, regardless what is accomplished with, say, the last twenty-five or thirty percent of the budget. This is probably correct, but this kind of judgmental decision is not captured in the optimization model: the optimization problem a site manager would be solving in making such a decision simply is not included in the objective functions we pose. We discuss the consequences of relaxing this 90–10% cost split in section 4.2, which presents the results of the minimize-all-risks problem.

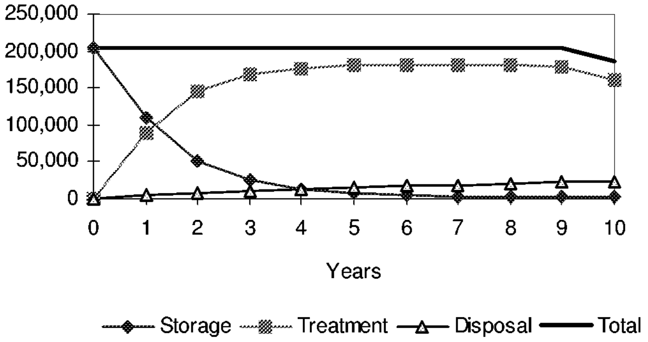

Figure 1. Expenditures for storage, treatment, and disposal ($1,000). Goal: Minimize terminal period risk.

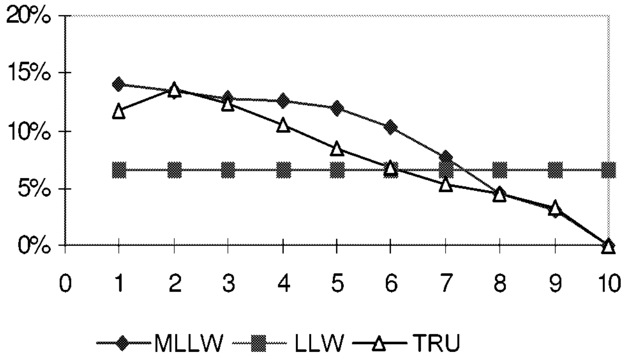

Figure 2. Percentage of each waste type treated per period (base budget). Goal: Minimize terminal period risk.

Table 2
Yearly treatment volumes, by waste type, terminal risk-reduction goal.

Period	MLLW (%)	LLW (%)	TRU (%)
1	14.06	6.67	11.81
2	13.39	6.67	13.53
3	12.85	6.67	12.29
4	12.62	6.67	10.42
5	12.00	6.67	8.50
6	10.34	6.67	6.78
7	7.70	6.66	5.43
8	4.45	6.67	4.55
9	3.01	6.66	3.35
10	0.00	6.67	0.00
Total	90.44	66.66	76.65

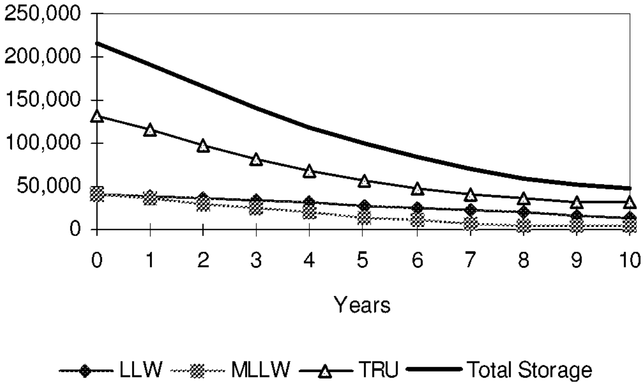

Figure 3. Storage risk. Goal: Minimize terminal period risk.

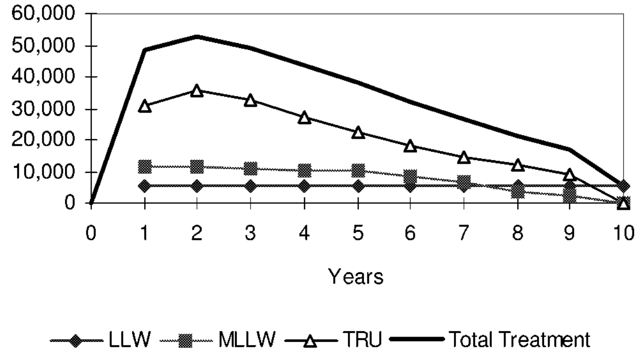

Figure 4. Treatment risk. Goal: Minimize terminal period risk.

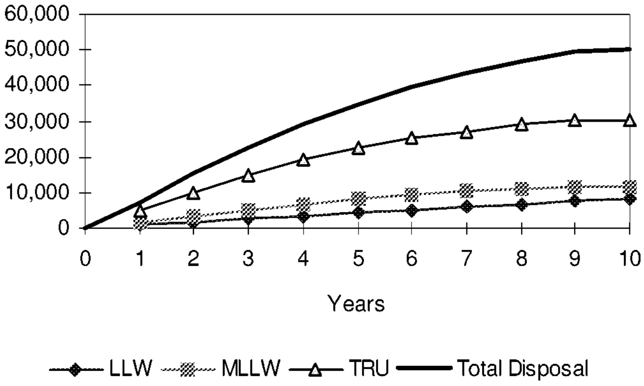

Figure 5. Disposal risk. Goal: Minimize terminal period risk.

without the risk constraints, no resources at all would be devoted to containing risk (in which case, risks would have gone to infinity). Figure 2 shows the percent of each waste type's initial stock that is treated in each period, and table 2 reports that information numerically.

Figures 3–5 show the progress of risks of each waste group in storage, treatment, and disposal. Risks in storage fall smoothly (figure 3), those in treatment rise sharply and stay about the same level throughout the treatment period (figure 4), and disposal risks rise as treated wastes are placed into that state. However, disposal risks increase by far less than storage risks fall, which is, of course, the point of treatment. (If we have overstated the extent to which storage costs fall when stored wastes are withdrawn for treatment, a smaller volume of waste could be treated than we have calculated, and terminal risks would remain higher.)

The base budget under this goal permits reduction of risk from its period-zero value of 216,000 to 99,000 (storage and disposal risk only in period 10; excluding treatment risks for LLW) in period 10. Increasing the budget by 50% in each period permits terminal risks to be reduced to just under 86,000. Figure 6 and table 3 report the annual treatment volumes accomplished with this 50% budget increment.

These two cases represent reductions of risks, by the terminal period, of 54% and 60%, respectively. The 50% increment in the budget purchases an incremental 6 percent-

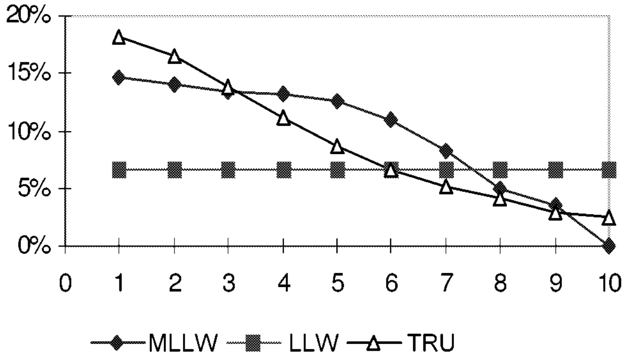

Figure 6. Percentage of each waste type treated per period (1.5 × budget). Goal: Minimize terminal period risk.

Table 3
Treatment volumes with budget increased by 50%, terminal risk-reduction goal.

Period	MLLW (%)	LLW (%)	TRU (%)
1	14.63	6.67	18.18
2	13.95	6.67	16.54
3	13.41	6.67	13.79
4	13.17	6.67	11.07
5	12.54	6.67	8.66
6	10.86	6.67	6.66
7	8.21	6.66	5.09
8	4.99	6.67	4.02
9	3.55	6.66	2.89
10	0.00	6.67	2.50
Total	95.30	66.67	89.40

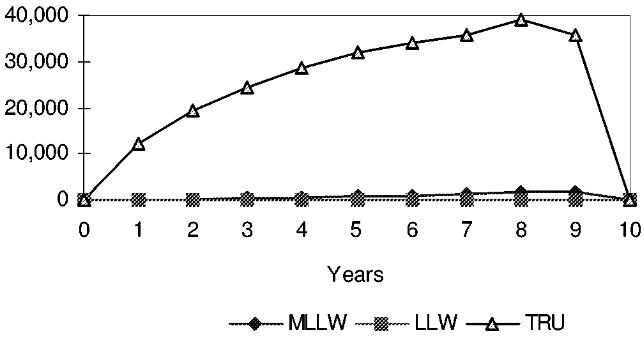

Figure 7. Cost per unit volume of waste types in treatment ($/million cubic meters). Goal: Minimize terminal period risk.

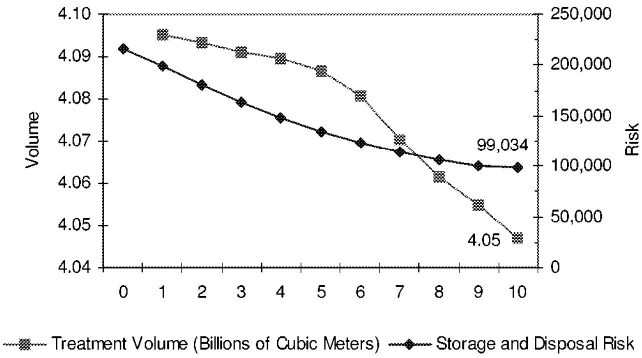

Figure 8. Progress of treatment and risk reduction for all waste types (base budget). Goal: Minimize terminal period risk.

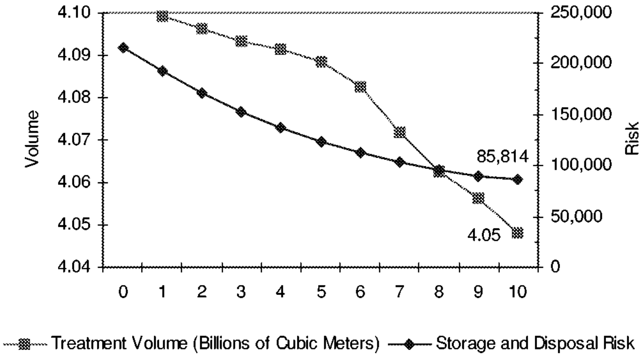

Figure 9. Progress of treatment and risk reduction for all waste types (1.5 × base budget). Goal: Minimize terminal period risk.

age points of terminal risk reduction. One of the principal reasons for this difficulty in "completing" the clean-up is that unit treatment costs rise as the volume of waste to be moved out of storage and through treatment falls. Thus, previous "successes" in treatment reduce the ability to conduct subsequent treatment. Figure 7 shows the progress of unit treatment costs over the course of a ten-year program, for the base budget of $2.03 billion. The unit cost of treatment for TRU nearly quadruples between years 1 and 8, despite increasing returns to scale (decreasing unit cost) of volume management in treatment. If a reader believes that our calculations of unit treatment costs, as shown in figure 7, are lower than treatment costs in fact are, the consequences of adjusting our cost structure accordingly are straightforward: less waste will be treated, and terminal-period risk will be reduced by less than we have simulated.

Figures 8 and 9 summarize the progress of treatment of all waste types and cumulative reduction of risk in storage and disposal, for the base and augmented budgets. In the first few periods, the program under the augmented budget treats a marginally larger total quantity of waste of all types, but by periods 4 and 5, the volumes are virtually the same as with the smaller budget. The cumulative risk for wastes in storage and disposal continue to drop off gradually, reflecting the smaller volumes of waste treated in the later periods.

Figure 10 shows how budget increases affect terminal risk and waste volumes treated. The horizontal axis shows increments of 10% of the base budget, up to 150% of the base budget; the left vertical axis measures total risk of all waste types in storage and disposal; and the right axis measures the percentages of the initial waste stocks treated. The risk axis runs from a maximum to 225,000 to zero because the total risk of wastes in storage before the ten-year treatment period begins is 216,000, and the minimum possible risk in this measuring system is zero. The base budget can deliver a terminal risk of around 99,000, while a budget 50% larger can deliver a terminal risk of around 85,000. The budget increments have no effect on the volume of LLW treated, primarily because it contributes so little to risk, which is the focus of the objective function (i.e., reducing risk is the goal, not treating waste per se). But

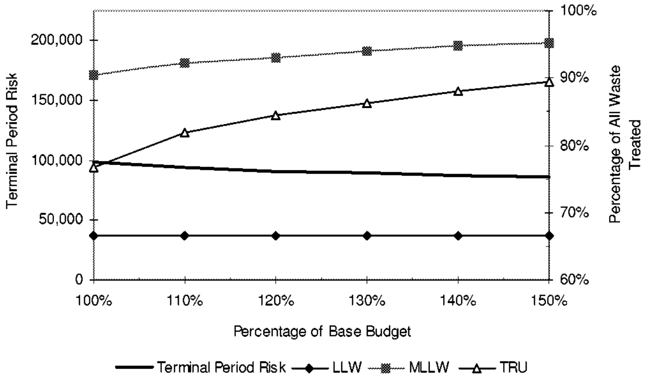

Figure 10. Impact of budget increases on terminal period risk and percentage of waste treated. Goal: Minimize terminal period risk.

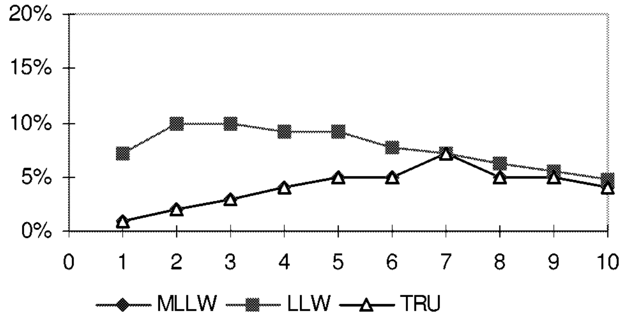

Figure 11. Percentage of each waste type treated per period. Goal: Minimize all period risk.

these increments increase the percent of TRU treated from around 78% to just under 90% and the volume of MLLW from around 91% to just over 95%, exactly because of their larger contribution to risks.

4.2. Program 2: Minimize current-period risk

The goal of this program is to minimize all risks at all points in time during the period of active clean-up: worker safety, public health, and environmental risks. These risks, for each waste group and for each state – storage, treatment, and disposal – are held to the minimum achievable levels. A major, systemic floor on the risk reduction that can be achieved is the large proportion of the composite, variable input that must go to volume management, as contrasted to risk management. Nonetheless, this program achieves much lower treatment proportions than does the terminal risk minimization program: 41.1% each of the TRU and MLLW and 76.7% of the LLW (somewhat more than under terminal risk minimization) are treated over the ten years. The goal of minimizing current-period risk rather than terminal-period risk makes it optimal to treat more of the less risky waste group, the LLW, because its treatment risks are low, and the higher treatment risks associated with TRU and MLLW can be avoided by treating the less risky waste. These annual treatment volumes are shown in figure 11 and table 4.

This program also reduces terminal-period risks of wastes in storage and disposal to 143,000, a reduction of nearly 34% from the initial level. The budget constraint never binds in this program, although a far larger proportion is spent than in the mortgage minimization program – nearly 88% as contrasted with 52%, as we show below in section 4.3.[15] Figure 12 shows the time path of expenditures for program 2, and figures 13–15 show the time path of risks in each state. The time profile of treatment risk (figure 14) is strikingly different from the corresponding profiles for terminal-risk minimization and mortgage-cost minimization (figures 4 and 21). Clear but less striking, cross-goal differences appear in the time profiles of storage and disposal risk (cf. figure 13 with figures 3 and 20, and figure 15 with figures 5 and 22).

Table 4
Yearly treatment volumes, by waste type, current-period risk minimization goal.

Period	MLLW (%)	LLW (%)	TRU (%)
1	1.00	7.14	1.00
2	2.00	10.00	2.00
3	3.00	10.00	3.00
4	4.00	9.09	4.00
5	5.00	9.09	5.00
6	5.00	7.69	5.00
7	7.14	7.14	7.14
8	5.00	6.25	5.00
9	5.00	5.56	5.00
10	4.00	4.76	4.00
Total	41.14	76.73	41.14

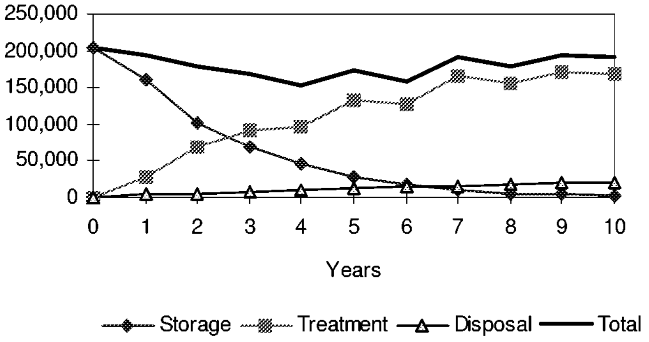

Figure 12. Expenditures for storage, treatment, and disposal ($1,000). Goal: Minimize all period risk.

---

[15] The constraints on minimum volumes of MLLW and TRU to be treated in each period are binding, only the volume of LLW treated exceeding the quantity specified in its treatment constraint, and that, in fact, by a considerable amount. If one considered these minimum-treatment-volume constraints to reflect the Site Treatment Plan and the overall management philosophy to be roughly consistent with the goal of this objective function, the unofficial wisdom that little of consequence would be done outside the Site Treatment Plan seems to have some foundation, even if it is not entirely correct. Thus, despite the fact that the budget constraint never binds in the solution to this program, the "extra" dollars do not go to additional risk reduction because the additional treatment that would involve would entail additional risk that the objective function tells the manager to avoid. Even less of the riskier wastes would be treated without the minimum-treatment-volume constraints.

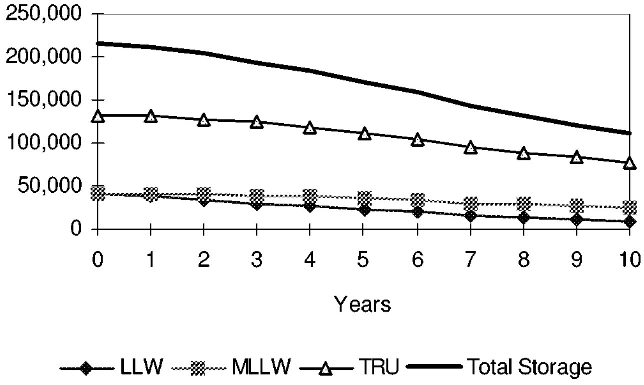

Figure 13. Storage risk. Goal: Minimize all period risk.

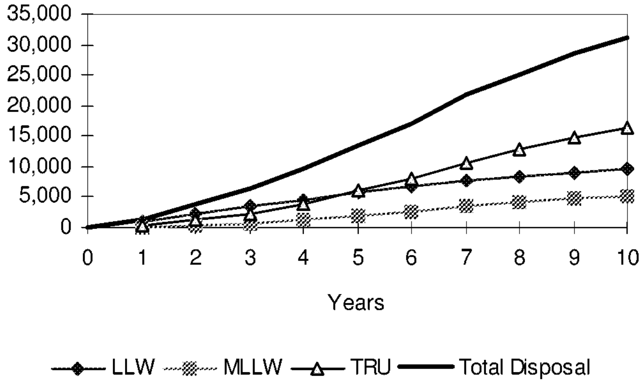

Figure 15. Disposal risk. Goal: Minimize all period risk.

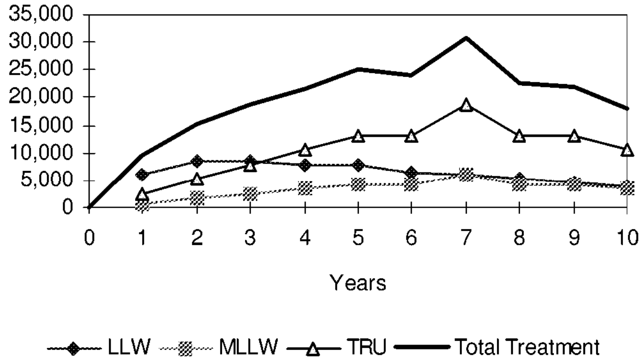

Figure 14. Treatment risk. Goal: Minimize all period risk.

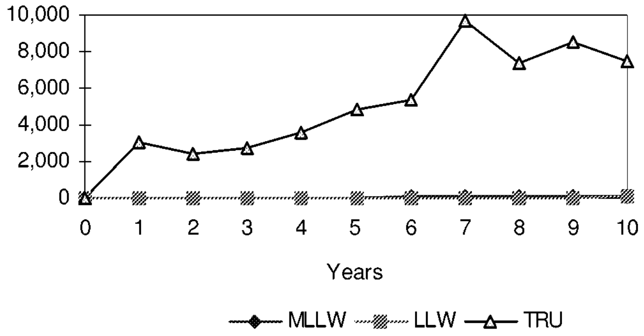

Figure 16. Cost per unit volume of waste types in treatment ($/million cubic meters). Goal: Minimize all period risk.

Figure 16 traces the path of treatment costs for each of the waste types, although only that for TRU is clearly identifiable on the scale of the graph. After period 1, the unit treatment cost of TRU falls until period 4; it rises gradually and continuously through period 6, and increases sharply in period 7, when the volume of TRU treated rises from its levels in periods 5 and 6. It falls back again in period 8, when the volume treated again falls back to near the level of periods 5 and 6. The unit cost is only about a quarter of what it is under the terminal-risk minimization goal (figure 7). Figure 17 presents the reduction in risk proceeding from total treatment volume in the current-period risk minimization program.

It may surprise some readers that as much waste gets treated in this scenario as we project, considering the possibly large costs of driving risks indefinitely lower. Recall, however, that in parameterizing the model, we used the 90–10% split of costs between volume management and risk management, and that while this starting value of the composite labor input for parameterization did not absolutely constrain the results to a 90–10% allocation of costs (or, equivalently, composite labor inputs) between volume and risk management in the solutions, it did impose a tendency for risk reduction to become asymptotic not far away from this general cost-ratio vicinity. That is, the additional physical reductions in risk tend to get very small when the amount of the composite labor input applied to the risk management activities exceeds the range of 160–200 units. If we re-parameterized the model to obtain values of the α and β coefficients using start values of, say 500–750 units of the composite labor input, keeping risk per unit volume at the maximum levels allowed by the constraint functions, in each of the programs, would be much more expensive than it is with the current specifications. In program 1, the terminal-period risk minimization problem, less waste could have been treated and the terminal-period risk would have been much higher. In the current problem of minimizing all risks in each period individually, the minimized level of risks would have been substantially higher, and less waste would have been treated.

4.3. Program 3: Maximize mortgage reduction

This problem is one of cost minimization, with a tradeoff between treatment costs during the ten years of possible clean-up activity and an indefinite future period over which taxpayers reap the benefit of lower Surveillance and Monitoring (S&M) costs on wastes that have been treated. This cost saving is literally purchased with treatment costs, and the optimal stopping condition on treatment is the equalization of additional treatment costs and the additional discounted present value of S&M cost savings.

The unit costs of treatment are quite high relative to unit costs of storage for untreated waste, so this goal would have DOE treat very little of these wastes. It is cheaper to keep the vast majority of them in untreated storage indefinitely than it is to pay the cost of treating them in return for future cost savings. Consequently, only small proportions

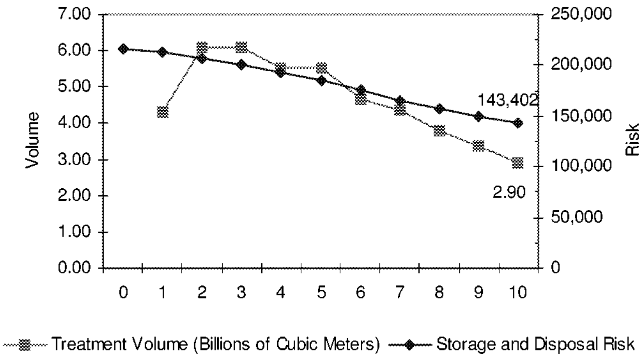

Figure 17. Progress of treatment and risk reduction for all waste types (base budget). Goal: Minimize all period risk.

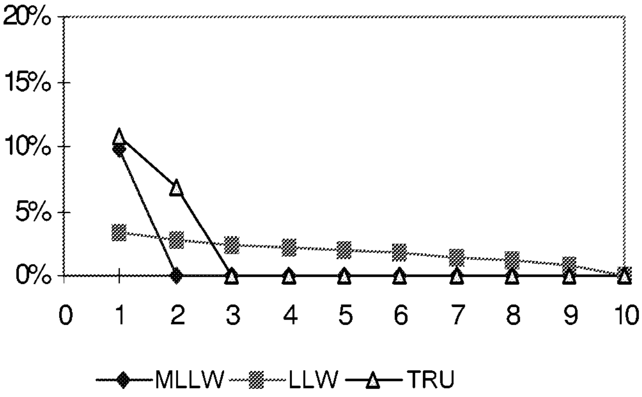

Figure 18. Percentage of each waste type treated per period. Goal: Maximize mortgage reduction.

Table 5
Yearly treatment volumes, by waste type, mortgage reduction goal.

Period	MLLW (%)	LLW (%)	TRU (%)
1	9.84	3.31	10.86
2	0.00	2.70	6.83
3	0.00	2.32	0.00
4	0.00	2.07	0.00
5	0.00	1.87	0.00
6	0.00	1.68	0.00
7	0.00	1.45	0.00
8	0.00	1.14	0.00
9	0.00	0.69	0.00
10	0.00	0.00	0.00
Total	9.84	17.22	17.69

of the initial stocks in storage are treated under a mortgage reduction plan: 9.8% of TRU, 18.4% of MLLW, and 17.2% of LLW. Using the GAO-recommended discount rate of 3% for capital projects, the cost of treatment is recovered by reduced storage costs in disposal in year 2013, the seventh year after the end of the treatment period. Figure 18 and table 5 show the annual treatment volumes.

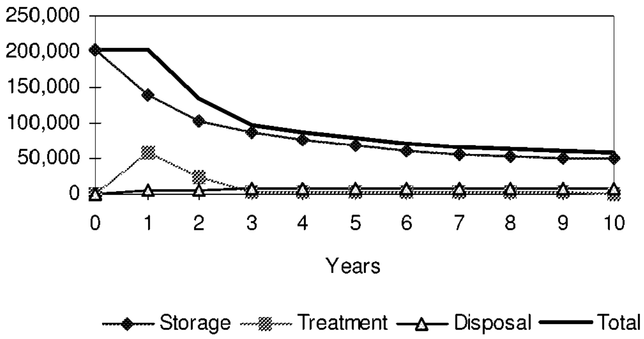

Figure 19. Expenditures for storage, treatment, and disposal ($1,000). Goal: Maximize mortgage reduction.

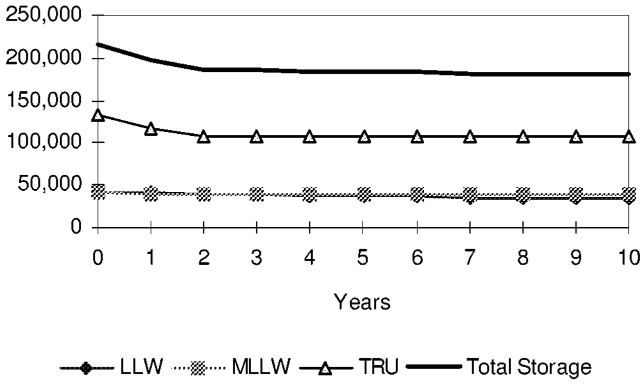

Figure 20. Storage risk. Goal: maximize mortgage reduction.

As a consequence of the small amount of treatment and the rapid reduction in storage costs with the withdrawal of a small volume, the budget constraint is never binding (although it comes close in the first period, with expenditures of $201 million out of a budget of $203 million). By the fourth period, slightly less than half of the budget is used, and by the seventh period, less than one-third of the budget is spent. Through the entire ten-year period, only 52% of the planned budget is spent. This is the reason we said above, in section 2.4, that we believed the magnitude of the bias to expenditures versus cost savings created by not discounting the ten-year operational expenses to be minimal. Discounting would have led to more treatment, but only by a small amount. Figure 19 shows the total expenditure pattern for the ten-year period, by type of expenditure.

Another consequence of the small treatment volumes is the relatively small reduction in risks achieved by the end of the treatment phase. The tenth-period risk corresponding to terminal risks in program 1 is 191,000, a reduction of 11.5% from the initial level. Figures 20–22 show the progress of risks in each state. Risk drops modestly in storage; spikes early, then falls to very low levels in treatment; and spikes sharply early in disposal, then rises very gradually, reflecting the small movement of wastes through treatment and into disposal. Figure 23 shows a fall in unit treatment costs for TRU from period 1 to period 2, but those costs are zero after period 2 because no TRU is treated. The peak unit cost is slightly higher than under the minimize-all-period-risk goal (figure 16), and appears substantially earlier, but otherwise has a much lower level throughout the rest of the plan period, essentially because so little is done with TRU under this goal.

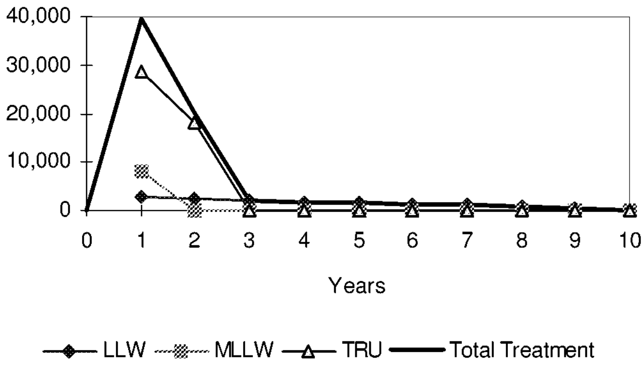

Figure 21. Treatment risk. Goal: Maximize mortgage reduction.

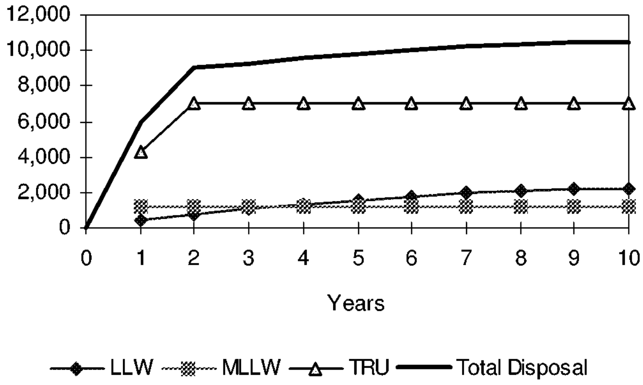

Figure 22. Disposal risk. Goal: Maximize mortgage reduction.

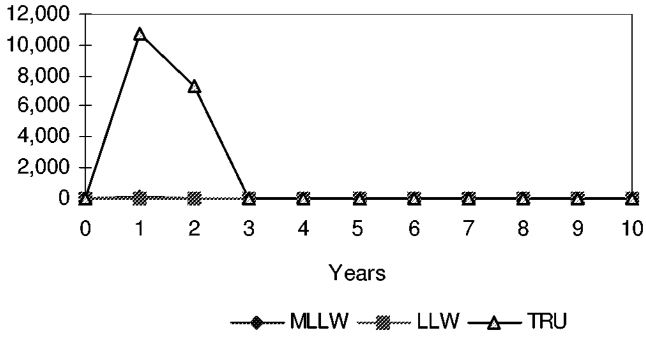

Figure 23. Cost per unit volume of waste types in treatment ($/million cubic meters). Goal: Maximize mortgage reduction.

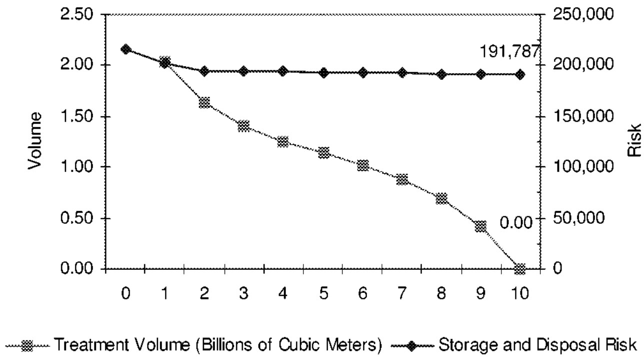

Figure 24. Progress of treatment and risk reduction for all waste types (base budget). Goal: Maximize mortgage reduction.

Figure 24 depicts the progress of risk reduction in storage and disposal, as it is reduced by volume treated over the ten-year period. Under the mortgage reduction goal, the rapid decrease in volumes of waste treated leaves risk reduction line nearly flat, emphasizing that this is not a risk-related goal.

5. Inter-program comparisons and conclusions

The three goals produce sharply divergent results regarding cost, risk, and clean-up. Table 6 summarizes the values of some of the variables of recurrent interest under these alternative sets of motivations. Perhaps most striking is the cost difference between the mortgage reduction and terminal risk minimization scenarios. The terminal risk goal is a genuine clean-up goal, and it is the most costly of the three goals. Also particularly interesting is the fact that the current-period risk minimization goal delivers less risk reduction over the course of the ten-year program than does the terminal risk minimization program. Clearly, current-period risk minimization trades off risks to the current generation against risk to future generations, pushing the incidence of risks off into the future.

The different overall goals impose strikingly different unit treatment costs, which are particularly visible for TRU in the diagrams for unit treatment cost. Under the goal of minimizing terminal-period risk, the unit cost of "treating" TRU rises to just under $40 thousand per million cubic meters, in period 8. Under the goal of reducing mortgage costs to the greatest extent possible, the maximum treatment cost for a million cubic meters of TRU is around $11 thousand, in period 1, and under the goal of minimizing risks in each period, the maximum cost of treating TRU is just under $10 thousand per million cubic meters in period 7. The reason for the treatment cost differences lies in the production function for treatment. The specification of the treatment *volume* function has increasing returns to scale in treatment (i.e., the more treated at one time, the lower the unit cost) but sharply decreasing returns to scale in management of unit risk; additionally, treatment costs are affected negatively by having a larger stock in storage to work with at the beginning of the period. Combining these three effects yields a substantially larger cost of treating the later units of TRU. In the goal of minimizing terminal-period risks, 90% of the initial TRU stock is treated during the ten-year period, while under the mortgage-cost-reduction and each-period-risk-minimization goals, 41% and 9.8% of the TRU stock is treated. Clearly, the much more extensive treatment of the TRU in terminal-risk minimization is quite costly.

The difficulty in pushing risk reduction much beyond 50% of its initial value points to a highly directed focus for R&D. Without improved treatment technologies, budget increments directed toward treating the last 10% of TRU and MLLW stocks seem largely futile. Reduced treatment costs also could yield a larger role for treatment, and consequent

Table 6
Summary statistics on accomplishments under alternative goals.

	Terminal risk minimization	Current-period risk minimization	Mortgage reduction
Budget spent (%)	100	88	52
Terminal-period storage and disposal risk (in units of 1000)	99	143	192
Reduction in initial risk (%)	55	34	12
TRU treated (%)	77	41	18
MLLW treated (%)	90	41	10
LLW treated (%)	67	77	17
Total waste treated (%)	67	77	17

Table 7
Costs of continuing stewardship if treatment were stopped at different times.

Time	Terminal period risk reduction	Optimal mortgage reduction	Total risk minimization in each period
0	2,908,125	2,908,125	2,908,125
1	1,637,135	2,063,096	2,368,292
2	841,779	1,557,367	1,570,746
3	516,102	1,347,518	1,100,552
4	380,424	1,192,640	793,278
5	327,800	1,074,904	577,743
6	313,812	984,959	456,426
7	317,599	917,816	368,398
8	329,069	870,997	345,424
9	342,682	845,008	338,576
10	353,589	845,008	341,252

risk reduction, in a mortgage reduction plan. If treatment costs are in fact believed to be substantially higher, relative to storage costs, than they are parameterized here, and if the fixed costs of storage are quite large, such that the early reductions in storage volumes do not sharply reduce storage costs, squeezing funds for treatment out of current storage budgets could be quite difficult. This is a situation that DOE personnel at field offices have noted. Such a conjunction of cost characteristics would imply that a simple reduction in the parameterized values of treatment costs used here might not result in substantially greater volumes of waste getting treated.

Rearranging the information from the three programs offers some additional insights into what the different objective functions are accomplishing. In table 7, we present what we call the continuing stewardship costs if all treatment were ceased in each period under consideration. For example, in the first row of each of the three goal columns, the Department of Energy would incur the present discounted value of $2.91 billion in future storage costs if it never proceeded with any treatment at all. Continuing down the "Terminal period risk reduction" column to the first period, we find that if treatment volumes were determined by this objective function for one period but thereafter were ceased altogether, the Department would face a present discounted value of $1.64 billion of future storage and disposal costs. Following this same goal but going through *two* years of the ten-year treatment plan and then stopping treatment forever would leave the Department with a future bill of $842 million, in present discounted value, for storage and disposal. Following these stewardship costs, so defined, down each column, we find the peculiarity under the goals of "Terminal period risk reduction" and "Total risk minimization in each period" that stewardship costs do not continue to fall throughout the entire ten-year treatment plan. They reach a minimum in period 6 with "Terminal period risk reduction" and in period 9 with "Total risk minimization in each period." They reach a minimum in period 9 of the "Optimal mortgage reduction" program and remain the same in the last period.

Because of the cost structure embedded in the production functions for storage and production, at some point, it becomes cheaper to leave waste in storage than to put it in disposal: quantities have gone down in storage and gone up in disposal to the points where increasing costs are kicking in in the latter and reducing costs even further in the former. The fact that this does not happen in the mortgage-minimization program is attributable to the cost-minimization in the objective function. That is another reason the latter program stops treating waste in the last period.

In none of the goals examined here has there been any role for the *value* of risk reduction. The expensive, terminal-risk-reduction program might be worth its extra costs in terms of the value of the additional risk reduction it offers, but without explicit information on those values incorporated into the plan, risks are simply reduced – or at least waste continues to be treated and placed in disposal – until the budget runs out. If adequate information were available on the costs of the initial risks and on the value of their reduction, a logical risk management strategy would be a cost–benefit approach in which incremental expenditures were equated to the incremental values of risk reduction. As the current plans are implemented by the optimization programs modeled here, the units of the composite variable input are allocated to risk management of the different waste groups, and across states (storage, treatment, and disposal), so as to equalize the incremental risk containment derived from an increment to the composite variable allocated to each waste group and state. This occurs simply via the optimization. However, there is no such clear, value rationale guiding the allocation of composite inputs between risk management and volume management. The default allocation rule used by the optimization model simply equalizes the impacts on the value of the objective function of additional units of the composite input in each waste group and state. As we have noted, nothing in the objective functions used here offers information on values – just quantities: quantities of risks and quantities of dollars. There is no information on which to draw that would tell us how much, say, an $x\%$ reduction in risk is worth.

Appendix: The DOE/Oak Ridge 2006-Plan data

We took cost, risk, and waste volume data from the Oak Ridge 2006 Plan. We identified and collected these data at the project level for 73 ORNL Remedial Action (RA), Decontamination and Decommissioning (D&D), and Waste Management (WM) projects involving Mixed Low-Level Waste (MLLW), Low-Level Waste (LLW), and TRansUranic (TRU) waste. Through a series of combining and allocative operations on the raw, project-level data, we arrived at initial-period stocks (volumes) of MLLW, LLW, and TRU waste and the risk associated with the initial stock of each waste type.[16] The relative costs and risks of working with each of the waste types we derived from a combination of the project-level data and widely available engineering knowledge.

We combined three data bases to create a unique and integrated data base of cost, risk, and waste volumes. The 2006 Plan provided the cost information. Using information contained in the spreadsheet of raw values from the Management Evaluation Matrix (MEM) and the MEM qualitative values, current risk was determined. The Environmental Management Enrichment Facility Program's Waste Generation Forecast Data Base provided waste volume information. We obtained cost, risk, and volume data from U.S. DOE/Oak Ridge Operations Office [14] and Lockheed Martin Energy Systems (LMES) sources responsible for the development of the 2006 Plan. We reviewed the data for consistency and completeness but did not attempt to identify root data sources and validate the estimates provided to us. Rounding error is present in our raw data sources and our computational approach. The bias of such error is insignificant since we are dealing with differences in orders of magnitude of no more than six while performing computations with floating-point accuracy of 10^{31}.

A.1. The cost data

The cost data relating to each project are budget projections of funds to be allocated to each project during the 2006 Plan, based on the DOE assumption of level funding for the ten-year period, the ORNL share of which is $203 million per year. These costs are in 1997 dollars; to the extent that allowances have been made in these budget projections for either real cost increases or decreases (i.e., not simply price-level change due to gradual inflation), these are expressed in terms of constant 1997 dollars.

We retain the constant-dollar approach in our use of the cost data. We generally (but not always) avoid discounting costs in our analysis. For this decision we offer several reasons. First, the 2006 Plan's budgeting process itself deals with undiscounted dollars because Congress allocates funds in undiscounted dollars; translating back and forth between, say $203 million in 2005 and its present discounted value of roughly $185 million to refer to the same cost of the same activities would accomplish little besides confusion. Second, given the planning parameter of level funding in real terms, discounting at OMB rates for only a ten-year period does not leave much scope for rearrangement of activities to achieve lower present values of costs. However, in our analysis of the mortgage-reduction goal, which involves spending money over the ten-year clean-up period to save recurring costs over some longer period in the future, we do discount the avoided future costs.

A.2. The risk data

The risk data in the 2006 Plans of each DOE field office are based on the methodology of the Activity Data Sheet/Risk Data Sheet (ADS/RDS) and the Management Evaluation Matrix (MEM) [9, Attachment 4], as modified by the project baseline summary (PBS) process early in 1997 [10, Attachment C]. The risk scores generated with this methodology are dimensionless numbers.[17] While the

---

[16] For "combining" operations, we collapsed "paper" projects such as a Record of Decision (ROD) into the "dirt-moving" projects they supported, reducing the apparent number of "projectized" projects, but keeping constant the total clean-up activities to be undertaken. For the "allocative" operations, since these paper projects frequently had been assigned sizeable risks, and the actual dirt-moving projects they supported, comparatively low risks, we distributed the risks assigned to the paper projects in the 2006 Plan data to the corresponding dirt-moving projects; similarly with the costs of the paper projects. Another allocative operation was the separation of different waste types in a single project into our three waste categories. Corresponding to this distribution of waste types, we allocated the original project risk and cost in proportion to the volume of each type of waste.

[17] That is, they do not have the dimensions of, say, excess cancer deaths thirty years hence, per unit of hazardous material.

full, RDS/MEM procedure includes components for various administrative risks (i.e., will we or will not we complete the job?) and technological/financial risks (will the proposed technology produce output which satisfies acceptance standards, or might we have to re-do part of all of the waste, at additional cost?). We have included only the portions of the risk scores attributed to worker safety, public health, and environmental safety. The measurement of risk with these tools has been the subject of some criticism, which we discuss more fully in the appendix to [3]. The risk data we use from Oak Ridge, plus the equivalent risk data from the other DOE facilities are what DOE is using to prioritize risks and sequence project actions throughout the entire Weapons Complex in the 2006 Plan. Rather than use these data blindly, however, we have studied the relative properties of our waste-specific risk scores to assess their correspondence to widely recognized orders of riskiness of the different waste types (e.g., TRU risk per unit of TRU waste > MLLW risk per unit of MLLW > LLW risk per unit of LLW waste) and their rough accordance with somewhat more judgmental, quantitative ratios of risks (e.g., TRU risk order-of-magnitude > LLW risk; MLLW risk several times > LLW risks).

This data source provides only initial-period risk data. One of the things the models do is reduce initial-period risk by treating certain volumes of specific waste groups each period and moving the treated material to some final (or long-term temporary) disposition state and location. The logical mechanisms of the RDS/MEM system which generated these risk data for the initial period (i.e., for the initial stocks of waste) do not continue to operate on the risk data once they are put into the models; the models provide their own logic regarding risk-reduction (or containment) per unit of effort (cost) in each of the three states of the model (storage, treatment, disposal). Thus any deficiencies in the construction of the initial-period risk data that are not reflected in their relative orders of magnitude do not continue to operate actively on these data while they are used in the models developed here.

A.3. The volume data

The volume data are relatively straightforward. At the project level, projects are identified as Remedial Action (RA), Decontamination and Decommissioning (D&D), or Waste Management (WM). Allocating the "volumeless" administrative projects across their associated "dirt-moving" projects, most of the identified projects at ORNL deal with two or more of the three waste types (MLLW, LLW, TRU). Our optimization models use quantities of the three waste types as data; the optimization is of the quantity of each waste type to be moved out of storage and treated, then moved to disposal, in each of the ten time periods. We do not attempt to derive some optimal sequencing of originally identified projects, but rather the quantities of the waste stocks dealt with. Consequently, we map the 73 originally identified projects onto three waste stocks existing at the initial time period ("time zero" in the terminology of the model). We describe the full details of how we accomplish that in the appendix to [3].

References

[1] H.B. Chenery, Engineering production functions, Quart. J. Econom. 68 (1948) 507–531.

[2] M. Hery, F. Diebold and G. Hecht, Exposure of contractors to chemical pollutants during the maintenance shut-down of a chemical plant, Risk Anal. 16 (1996) 645–655.

[3] D.W. Jones, K.S. Redus and D.J. Bjornstad, A non-linear programming model of alternative risk-related goals for DOE's weapons complex clean-up: a case study of the Oak Ridge site, Joint Institute for Energy and the Environment, Knoxville, TN (February 1998).

[4] J. Marsden, D. Pingry and A. Whinston, Engineering foundations of production functions, J. Econom. Theory 9 (1974) 124–140.

[5] V.L. Smith, *Investment and Production* (Harvard University Press, Cambridge, MA, 1961).

[6] R.J. Toland, J.M. Kloeber, Jr. and J.A. Jackson, A comparative analysis of hazardous waste remediation alternatives, Interfaces 28 (1998) 70–85.

[7] J.L. Tucker, D.B. Ridout and R.B. Shaw, Using linear programming to optimize rehabilitation and restoration of injured land: an application to US army training sites, J. Environ. Manag. 52 (1998) 173–182.

[8] U.S. DOE/EM, *Estimating the Cold War Mortgage; The 1995 Baseline Environmental Management Report*, DOE/EM-0232, Washington, DC (March 1995).

[9] U.S. DOE/EM, Addendum to the "Guidance for FY 1998 budget formulation", Washington, DC (31 January 1996).

[10] U.S. DOE/EM, *DOE Ten-Year Plan. Ten-Year Plan Guidance*, Mimeo, Washington, DC (20 December 1996).

[11] U.S. DOE/EM, *The 1996 Baseline Environmental Management Report*, DOE/EM-0290, Washington, DC (1996).

[12] U.S. DOE/EM, *Accelerating Cleanup: Focus on 2006*, Discussion draft (The National 2006 Plan), Washington, DC (July 1997).

[13] U.S. DOE/EM, *Accelerating Cleanup; Paths to Closure*, DOE/EM-1362, Washington, DC (1998).

[14] U.S. DOE/Oak Ridge Operations Office, *Environmental Management Ten Year Plan, Oak Ridge, Tennessee*, Oak Ridge, TN (31 July 1996).

[15] U.S. DOE/Oak Ridge Operations Office, *2006 Plan for the DOE/Oak Ridge Complex*, Mimeo, Oak Ridge, TN (September 1997).

[16] H.R. Varian, *Microeconomic Analysis*, 2nd ed. (Norton, New York, 1984).