A Multiattribute Utility Analysis of Alternatives for the Disposition of Surplus Weapons-grade Plutonium

James S. Dyer
CBA 5.202 (B6500)
The University of Texas
at Austin
Austin, TX 78712-1175
(512) 471-5278
j.dyer@mail.utexas.edu

Thomas Edmunds
Lawrence Livermore
National Laboratory
7000 East Av. (L-195)
Livermore, CA 94550
(510) 422-5156

John C. Butler
CBA 5.202 (B6500)
The University of Texas
at Austin
Austin, TX 78712-1175
(512) 475-7797

Jianmin Jia
Department of Marketing
Chinese University of
Hong Kong
Shatin, N.T. Hong Kong

May 23, 1997

This report was prepared with the support of a cooperative agreement between ANRCP and U.S. Department of Energy, (Agreement No. DE-FC04-95AL85832) and Lawrence Livermore National Laboratory (LLNL) under contract W-7405-Eng-48. This work was conducted through the Amarillo National Resource Center for Plutonium and the Lawrence Livermore National Laboratory. However, any opinions, findings, conclusions, or recommendations expressed herein are those of the authors and do not necessarily reflect the views of DOE or LLNL.

A Multiattribute Utility Analysis of Alternatives for the Disposition of Surplus Weapons-grade Plutonium

Abstract

This paper outlines an application of multiattribute utility theory to the selection of a technology for the disposition of surplus weapons-grade plutonium. The analysis presented evaluated thirteen alternatives, examined the sensitivity of the recommendations to the weights and assumptions, and quantified the potential benefit of the simultaneous deployment of several technologies. The results were used by the Department of Energy to support its recommendation of an alternative.

Keywords: Multiattribute Utility Theory; Multi-criteria Decision Making; Decision Analysis in Public Policy

1. Introduction

The end of the Cold War and subsequent arms limitation and reduction agreements have led to a surplus of weapons-grade plutonium in the United States and Russia. In order to prevent the proliferation of nuclear weapons, steps must be taken to manage this plutonium in a responsible manner. The Office of Fissile Materials Disposition (OFMD) of the Department of Energy (DOE) announced a Record of Decision (ROD) regarding alternatives for disposition of surplus plutonium on January 14, 1997 [DOE-ROD 97]. The ROD recommended that an existing reactor option and an immobilization alternative be developed in parallel to provide an integrated system for disposing of the material.

This recommendation was based on an extensive two phase evaluation carried out by OFMD with the support of scientists and engineers from the national laboratories and the private sector. Phase I consisted of a screening of thirty seven candidate alternatives, while Phase II focused on the thirteen alternatives that survived the screening process. These appraisals took into account non-proliferation, economic, technical, schedule, environment, and health and safety issues.

At the request of OFMD, a team of analysts from the Amarillo National Resource Center for Plutonium (ANRCP) and Lawrence Livermore National Laboratory provided an independent evaluation of the alternatives for the disposition of surplus weapons-grade plutonium that were considered during the Phase II effort. The ANRCP is supported by a consortium of three universities, The University of Texas at Austin, Texas A&M University, and Texas Tech University, and funds a research program dedicated to the investigation of issues related to the storage and disposition of plutonium. This paper summarizes the results of the ANRCP study, and provides support for DOE's recommendation.

2. Background

The ANRCP team adopted multiattribute utility theory (MAU) as a means of assembling the results of detailed technical, economic, schedule, environment, and nonproliferation analyses. In the past, MAU has been applied to problems such as siting an electricity generation facility (Keeney, 1980), choosing among vendors for the evaluation of alternatives for the commercial generation of electricity by nuclear fusion (Dyer and Lorber, 1982), and selecting a nuclear waste clean up strategy (Keeney and von Winterfeldt, 1994). The MAU methodology has also been supported for use in similar situations by the National Research Council, an agency of the National Academy of Sciences.[1]

The ANRCP evaluation team became involved with OFMD in May 1995, which marked the approximate beginning of the Phase II effort. During the Summer of 1995, members of the ANRCP team met on numerous occasions with OFMD personnel, and with the three alternative teams, scientists from the National Laboratories (Lawrence Livermore, Los Alamos, Oak Ridge, and Sandia) and TRW who were organized according to the three major types of alternatives under consideration: reactors, immobilization, and borehole alternatives. The OFMD also arranged interviews for ANRCP personnel with representatives of the Department of State, National Security Council and the White House Office of Science and Technology Policy. Meetings were also held with representatives of TetraTech, a private consulting firm responsible for the Programmatic Environmental Impact Statement (PEIS) [DOE-PEIS 96].

---

[1] National Research Council, letter to Ben Rusche, DOE/OCRWM, dated October 10, 1985.

A major objective of the initial stages of this evaluation effort was to determine the goals and performance measures that would be used to evaluate the Phase II alternatives, and this was accomplished during the Fall 1995. The alternative teams and TetraTech provided a preliminary set of estimates for these performance measures for the plutonium disposition alternatives that were under consideration. The ANRCP team used these preliminary estimates to develop an evaluation of the alternatives, and provided these results to OFMD and to representatives of the alternative teams in the form of a pre-decisional draft delivered Spring 1996, and in summary form at the OFMD project meetings in Washington, D.C. on November 15, and December 15, 1995.

These preliminary performance estimates were refined during 1996, and this information was communicated to the ANRCP team. The evaluation of the alternatives was updated as new information became available. In the remainder of this paper we will describe the MAU model and outline the model building process, as well as indicate the additional insights provided by the model regarding the strategy for plutonium disposition that was announced in the ROD.

3. Building the MAU Model

3.1 The Alternatives

Table 1 provides a concise list of the thirteen disposition alternatives considered in the analysis. Five alternatives were identified which would use surplus plutonium to fabricate mixed oxide fuel (MOX) for nuclear reactors that generate electric power. The spent fuel from these reactors would ultimately be placed in a geologic repository. Six other alternatives would require the immobilization of the surplus plutonium materials in borosilicate glass, ceramics or metal alloy castings; additional radionuclides would be added to provide a radiation barrier to inhibit recovery and reuse. This material would be transferred to the federal waste management system. Two disposal alternatives involving the placement of plutonium in a borehole were considered to be reasonable, one of them requiring immobilization in an inert matrix and the other utilizing direct emplacement. For additional details regarding these alternatives, see the DOE technical summary report [DOE-TSR 96].

REACTOR ALTERNATIVES	
Existing Light Water Reactors, Existing Facilities	MOX fuel fabrication plant built in an existing building at a DOE site, MOX irradiated in existing privately-owned commercial reactors
Existing Light Water Reactors, Greenfield Facilities	A new co-located pit disassembly/conversion and MOX fabrication facility built at a DOE site, MOX irradiated in existing privately-owned commercial reactors
Partially Completed Light Water Reactors	Commercial LWRs on which construction had been halted would be completed and operated by DOE
Evolutionary Light Water Reactors	New LWRs would be built and operated by DOE
CANDU Reactors	MOX fuel fabricated at a U.S. facility would be transported to one or more Canadian commercial heavy water reactors and irradiated
IMMOBILIZATION ALTERNATIVES	
Vitrification Greenfield	Surplus plutonium would be mixed with glass and radioactive materials at a new facility to form homogeneous borosilicate glass logs.
Vitrification Can-in-Canister	Surplus plutonium would be mixed with non-radioactive glass and poured into small cans. These small cans would be placed in larger canisters, which are then filled with radioactive waste glass.
Vitrification Adjunct Melter	Surplus plutonium would be mixed with glass and radioactive materials in a supplemental melter facility to form homogeneous borosilicate glass logs.
Ceramic Greenfield	Surplus plutonium would be mixed with ceramic and radioactive materials at a new facility to form homogeneous ceramic disks. These disks would be placed in a canister.
Ceramic Can-in-Canister	Surplus plutonium would be mixed with non-radioactive ceramic materials to form sintered ceramic pellets. These pellets would be placed in larger canisters filled with radioactive waste glass.
Electrometallurgical Treatment	Surplus plutonium would be immobilized with radioactive glass-bonded zeolite.
DIRECT DISPOSAL ALTERNATIVES	
Deep Borehole (Immobilization)	Surplus plutonium would be immobilized with ceramic pellets and placed in a borehole.
Deep Borehole (Direct Emplacement)	Surplus plutonium would be converted to a suitable form and placed in a deep borehole.

Table 1 – Disposition Alternatives

3.2 Objectives and Measures

The next step in the construction of a MAU model is the development of a "hierarchy" of objectives, sub-objectives, and measures. This hierarchy organizes the primary evaluation criteria, or objectives, into subsets of related measures to aide in communicating the results of the analysis to the decision makers and other stakeholders.

A preliminary set of goals and measures proposed by a team from Lawrence Livermore National Laboratory (LLNL) was a useful reference point for this effort (Edmunds, Koopman and Myers, 1995). We have also reviewed measures proposed for previous studies involving technology choices (e.g., Keeney, Lathrop, and Sicherman, 1986; Keeney and von Winterfeldt, 1994; Merkhofer and Keeney, 1987), for previous studies concerned with the management and disposition of surplus plutonium [NAS 94, NAS 95], and a study to assess technologies for the supply of tritium (von Winterfeldt and Schweitzer 1997). The resulting goal hierarchy is displayed in Figure 1.

As shown in Figure 1, three major objectives were identified at the highest level of the hierarchy, Non-proliferation, Operational Effectiveness, and Environment, Safety, and Health (ES&H). Intuitively, these three goals reflect the objectives of OFMD to achieve the disposition as efficiently as possible and with minimum risk to the public and natural environment.

The Non-proliferation objective is further subdivided into five sub-objectives: Theft (minimize the opportunities for theft of the materials by unauthorized parties), Diversion (maximize the resistance of the disposition alternative to the diversion of the plutonium by the host nation during processing, and to provide an internationally verifiable and acceptable process), Irreversibility (maximize the difficulty of recovering the material after disposition is complete), International Cooperation (foster international cooperation with U. S. disarmament and nuclear non-proliferation goals), and Timeliness (minimize the time required for the disposition effort to begin and to complete the disposition mission).

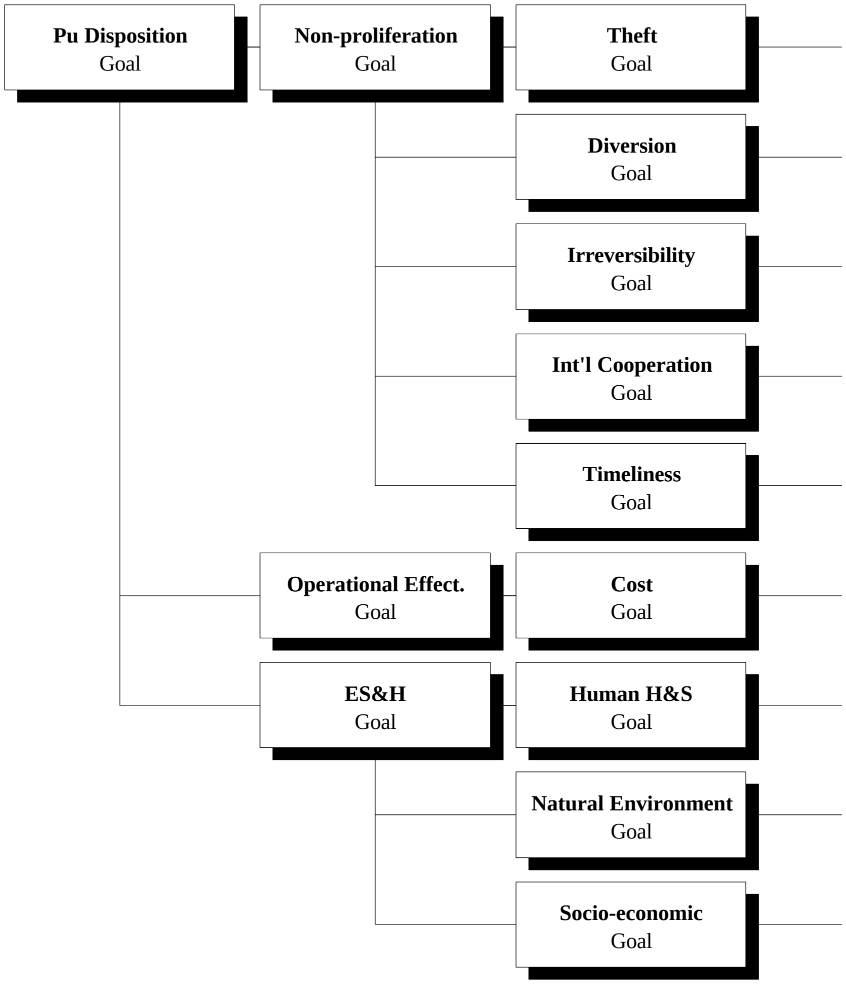

Figure 1 – High-level Objectives

The second objective shown in the Figure 1, Operational Effectiveness, includes the investment and life-cycle costs associated with an alternative. The third, Environment, Safety and Health, is made up of three sub-objectives: Human Health and Safety (minimize the incremental health impacts to the public and workers), Natural Environment (minimize the incremental impact on the natural habitat at the sites) and, Socio-economic (minimize the incremental adverse impacts to the human environment).

Figure 2 displays the measures used for the analysis. For a detailed description of these measures see the Dyer, et. al. (1996 and 1997).

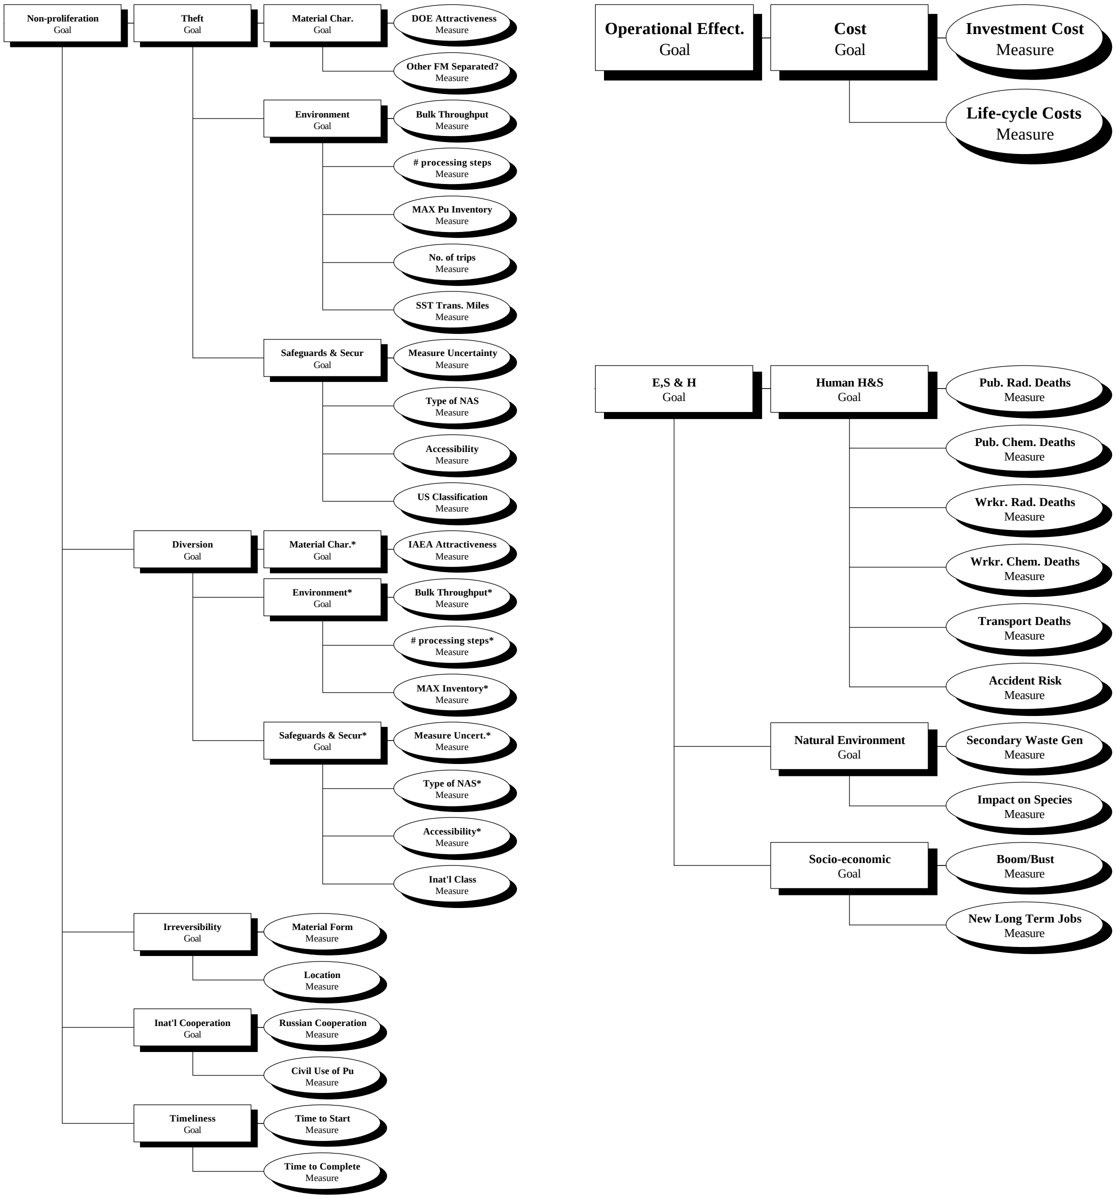

Figure 2 - Detailed Measures for Disposition Goals

In order to obtain performance estimates with respect to these measures, the ANRCP cooperated with the DOE in conducting a series of assessment meetings focusing on the major objectives. Members of the DOE safeguards and security (S&S) team played a major role in evaluating the performance of the alternatives on the Non-proliferation objective. Three two-day meetings were held with safeguards experts from Sandia, Lawrence Livermore and Los Alamos National Laboratories to define the measures and scales associated with S&S issues. The performance data for the Theft, Diversion and

Irreversibility sub-objectives were also provided by the S&S team. The Theft and Diversion data were provided by facility rather than by alternative. For example, the Existing Reactor, Existing Facilities alternative consists of four facilities: plutonium processing, MOX fuel fabrication, the reactor, and the repository. The S&S team provided data for each measure and each of these facilities. A “weak-link” approach was used to combine the scores at each facility into a single score for the alternative.

The weak-link approach assumes that a thief or the host country would strike at the most vulnerable facility. Therefore, each facility was scored using the assessed weights and utility functions, and the facility with the *worst* score was used to provide the performance measures for each alternative. A comparison with a weighted average of the facility scores was used to determine the impact of using the weak-link approach.

The data necessary to evaluate the International Cooperation sub-objective was provided by several smaller meetings with individuals DOE identified as experts in Russian technologies and policies. We also benefited greatly from conversations and meetings with the OFMD group.

The Operational Effectiveness data was provided by the three alternative teams. The three alternative team leaders met to maintain common standards and definitions across technologies. All cost data were processed using a common methodology and assumptions based on input provided by the alternative teams.

Environment, Safety, and Health data requirements were provided by the analysis necessary to develop the required Programmatic Environmental Impact Statement (PEIS) for the project, [DOE-PEIS 96]. Several meetings were held with personnel responsible for the PEIS to ensure the data were consistent with the objectives of the evaluation effort.

OFMD was concerned with selecting the best technology during this Phase II evaluation effort, and elected to postpone a siting decision until a more detailed study could be performed for the chosen technology. Thus, ES&H data were provided for each alternative at six candidate sites. Based on these preliminary estimates, we scored each of the six sites using the assessed weights and utility functions and used the data from the best performing site in the MAU model.

3.3 Aggregation of the Results

If stakeholder preferences are consistent with some special independence conditions, then a multi-attribute utility model $u(x_1, x_2, \ldots, x_n)$, where x_i represents the level of performance on measure *i*, can be decomposed into an additive, multiplicative, or other well-structured form that simplifies assessment. An additive multi-attribute utility model can be represented as follows:

$$u(x_1, x_2, \ldots, x_n) = \sum_{i=1}^{n} w_i u_i(x_i) \qquad (1)$$

where $u_i(\cdot)$ is a single-attribute utility function over measure *i* that is scaled from 0 to 1, w_i is the weight for measure *i* and $\sum_{i=1}^{n} w_i = 1$. If the decision maker's preference structure is not consistent with the additive model (1), then the following multiplicative model may be used, which is based on a weaker independence condition:

$$1 + ku(x_1, x_2, \ldots, x_n) = \prod_{i=1}^{n} \left[1 + kk_i u_i(x_i)\right] \qquad (2)$$

where $u_i(\cdot)$ is also a single-attribute utility function scaled from 0 to 1, the k_i's are positive scaling constants satisfying $0 \le k_i \le 1$, and k is an additional scaling constant that characterizes the interaction effect of different measures on preference. The value of *k* can be determined from one additional question similar to the questions used to determine the objective weights. As a special case when $\sum_{i=1}^{n} k_i = 1$, the multiplicative model (2) reduces to the additive model (1).

In this analysis of the thirteen alternatives for the disposition of plutonium, the additive model (1) was be used to aggregate the results of the evaluation effort. This model was chosen because the independence assumptions that justify the use of the additive model are reasonable for this analysis due to the relationships among the objectives and measures, and because the results of the analysis are easier to interpret when the additive model is used. In a later section we will discuss a sensitivity analysis where we tested the assumption of an additive model.

For a more detailed discussion of the assumptions underlying these two models, see Keeney and Raiffa (1976).

3.4 Tradeoffs and Weights

Each objective, sub-objective, and measure in the attribute hierarchy is given a weight. These weights reflect the value tradeoffs among objectives (or sub-objectives and measures within objectives), and are dependent on the ranges of the outcomes considered in the analysis.

The S&S team provided judgments regarding the measurement of Theft, Diversion and Irreversibility. These judgments were obtained in two different meetings that lasted two days each. After, the appropriate measures for these objectives and their associated utility functions were determined, which required the identification of the ranges over which performance levels of the alternatives might vary on each of the measures, we elicited tradeoffs among the measures. The assessment procedure continued until the group was able to reach a consensus opinion.

OFMD personnel provided the utility functions and tradeoffs for the International cooperation and Timeliness sub-objectives, as well as the judgments necessary to determine the weights within Non-proliferation. These assessments took place at several meetings with OFMD and representatives of other governmental agencies. At these same meetings, the weights and utility functions for Operational effectiveness were also elicited.

The majority of the tradeoffs presented above were made with respect to measures that were unique to the OFMD program. These tradeoffs considered measures related to the theft and diversion of weapons-usable plutonium, for example, or policy issues related to Russian cooperation and US non-proliferation policies. As a result, expert judgment was required to provide the baseline estimates of tradeoffs that give the information necessary to determine the weights on these measures.

In contrast, many different government programs have been subject to evaluations that have involved measures of environmental impacts, and information from these studies can be used to estimate reasonable tradeoffs within this category of measures. For example, Tengs, et al (1994) studied five hundred examples of life-saving interventions, typically US government actions and regulations, to determine the estimated cost per life year saved of these interventions. Previous DOE studies also exist which provide information regarding the tradeoffs between costs and statistical human lives, and between costs and other environmental impacts. Rather than asking for new expert judgments in this area where previous information is available regarding ES&H tradeoffs, this information from other DOE and government programs was used to determine the weights used in the baseline analysis.

Figure 3 displays the weights used in the analysis. Once again, it is important to emphasize that these weights do not represent measures of the relative importance of the corresponding measures. Instead, they are influenced by the ranges over which each of these measures varies, so they should be interpreted within this context.

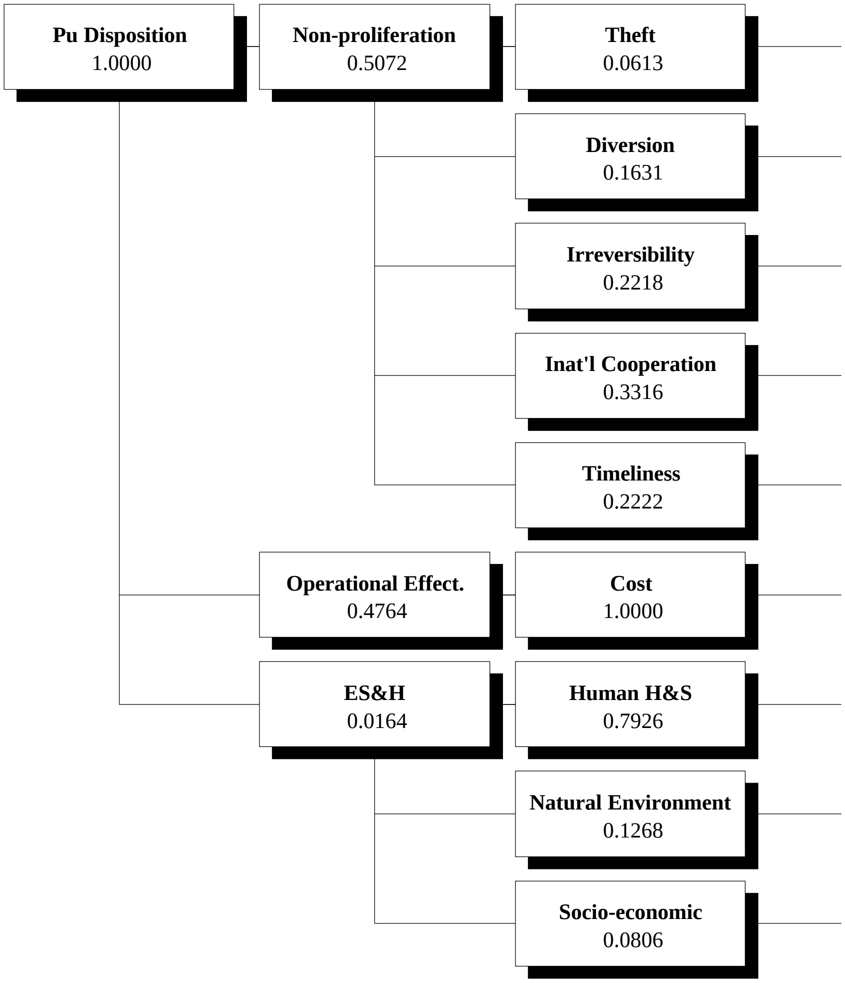

Figure 3 – High Level Weights

4. Evaluation of the Alternatives

4.1 Base Case Analysis

The component utility function scores were aggregated, using the additive multi-attribute utility function (1), within each of the three major objectives, and within each of the categories of objectives identified in Figure 2. During this aggregation, the weights are used to reflect the tradeoffs between measures, and are multiplied by the corresponding scores. This stage of the evaluation process is important and useful, since it provides scores for each alternative for the three major objectives of the plutonium disposition problem, and for the categories of objectives identified in Figures 1 and 2. At this stage it is possible to examine the relative strengths and weaknesses of the alternatives in more detail.

These results are presented in the form of *stacked bar graphs*, which provide a visual representation of the aggregated performance of each alternative on each major objective. In addition, these stacked bar graphs can be segmented to show the relative contributions of the individual sub-objectives and measures to the overall score for each alternative. Each segment represents the value of the performance of each alternative on each sub-objective or measure weighted by its relative importance captured through the tradeoff responses using the additive multiattribute utility model.

The Non-Proliferation objective highlighted in Figure 2 is the most complex of the three major high-level objectives of plutonium disposition. The Non-Proliferation objective consists of five sub-objectives: Theft, Diversion, Irreversibility, International Cooperation, and Timeliness. The rankings of the thirteen alternatives on the Non-proliferation objective are shown in Figure 4, which shows the relative contributions of performances on these five sub-objectives to the overall rankings.

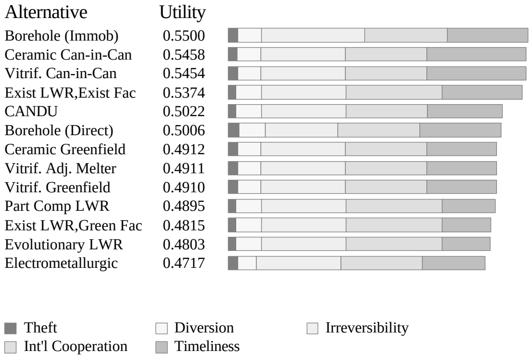

Figure 4 – Non-proliferation Ranking

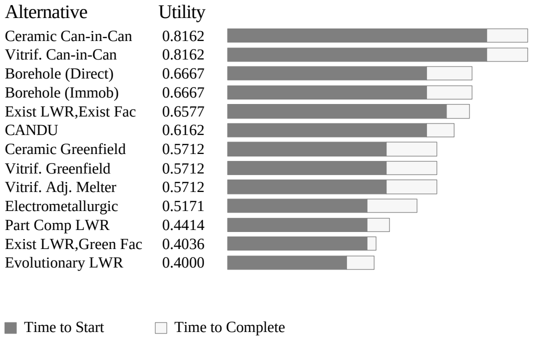

Figure 5 – Timeliness Ranking

Several observations can be made immediately based on the summary provided by Figure 4. It appears that Irreversibility, International Cooperation, Timeliness are the primary discriminators among the Non-proliferation sub-objective. The Timeliness objective ranking is displayed in Figure 5. The Can-in-Can alternatives can be started and completed relatively quickly and receive the highest scores. In Figure 6, it becomes apparent that all of the alternatives receive the same Irreversibility score, except for the borehole alternatives, because the material was all judged to be equivalent to spent fuel and would be placed in a repository. The borehole alternatives receive a borehole "bonus" due to the placement of the material in a deep borehole; the Borehole (Direct) alternative is penalized because the quality of the material would be more attractive for use in a weapons program if it were recovered. Finally Figure 7 displays the International Cooperation ranking. The experts expressed the belief that Russia was more likely to join a disposition effort that utilized reactor technology, but that this might be viewed unfavorably as a civil use of plutonium by certain interest groups. Thus, these two measures are negatively correlated.

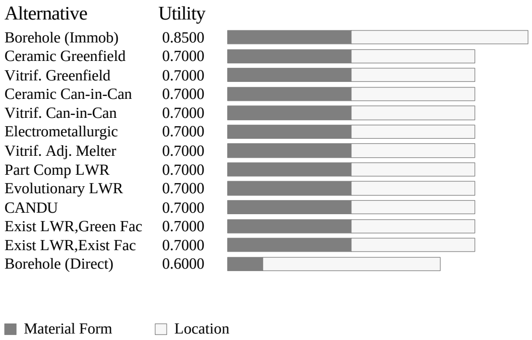

Figure 6 – Irreversibility

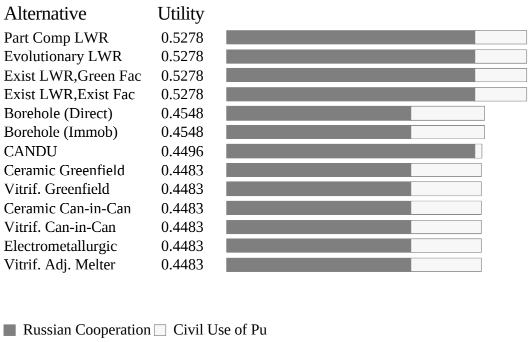

Figure 7 – International Cooperation Ranking

In terms of Non-proliferation, the Borehole (Immobilized) alternative is ranked first, and, based on inspection of the length of the component bars, this ranking is the result of superior performance on the sub-objective Irreversibility. The relatively poor performances of the Partially Completed LWR, the Existing LWR, Greenfield Facilities, and the Evolutionary LWR are due primarily to the Timeliness sub-objective as demonstrated in Figure 5. The Electrometallurgical alternative suffers from relatively poor estimates of performance on the Diversion and International Cooperation sub-objectives. These observations can be confirmed, of course, by reference to the underlying data, but Figure 4 may aid in identifying specific areas of performance that are "discriminators" among the alternatives.

Figure 8 presents the rankings of the alternatives on the Operational Effectiveness objective. The Vitrification Can-in-Can and Ceramic Can-in-Can alternatives perform very well on this objective, and are the highest ranked alternatives. The relatively high investment cost associated with the Evolutionary LWR alternative explains its low ranking on this objective.

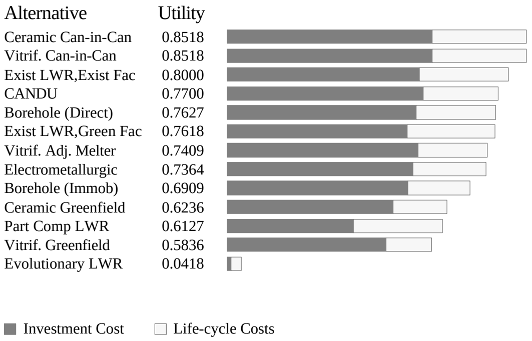

Figure 8 – Operational Effectiveness Ranking

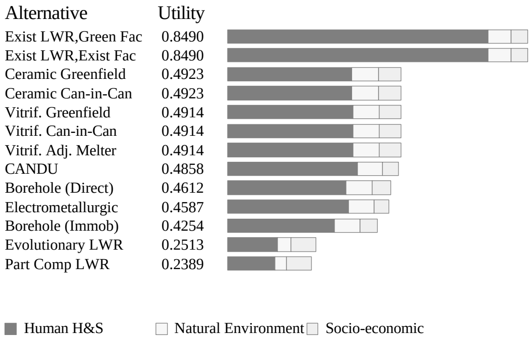

Figure 9 – ES&H Ranking

The rankings of the alternatives on the ES&H objective are shown in Figure 9. The use of either of the Existing LWR alternatives is expected to avoid incremental human fatalities due to the reduction in the need for mining and processing the uranium that would be displaced by the MOX fuel [DOE-PEIS 96 pgs. 4-977 – 4-980]. In contrast, the Evolutionary and Partially Completed LWR alternatives would require the operation of new reactors, and would result in forecasts of additional risks to human health and safety. The alternatives do not vary significantly with regard to the measures of impacts on the natural environment and socio-economic conditions. Additional details regarding the performances of the alternatives on the individual measures associated with ES&H are shown in Figure 9.

In the previous discussion illustrated with stacked bar graphs, the reactor, immobilization and borehole options were compared against one another on the major objectives. Comparisons within technology families may be useful in highlighting the relative strengths and weaknesses of relatively similar alternatives and variants. For the purposes of this discussion, the alternatives are divided into the reactor family and the immobilization and borehole family.

For ease of reference, the overall ranking of the reactor alternatives using the base case data is shown in Figure 10, and a comparison of the overall scores of the immobilization alternatives and the two borehole variants is shown in Figure 11. The weights used to aggregate the three primary objectives were assessed from OFMD personnel. Obviously, these "policy" weights are subject to debate as will be discussed in the section describing the sensitivity analyses.

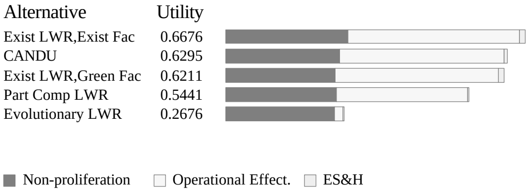

Figure 10 – Reactor Overall Ranking

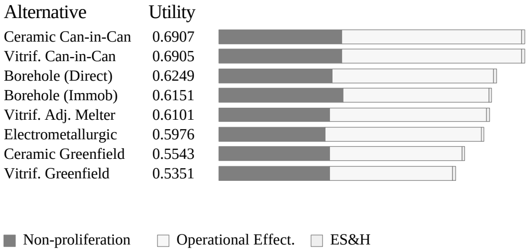

Figure 11 – Non-reactor Overall Ranking

Inspection of Figures 10 and 11 points out that while some alternatives score better than others, it may be unclear why these differences occur. In some cases, it is fairly easy to compare the underlying data of two alternatives and understand the differences. For example, Table 2 summarizes the differences between the Ceramic Can-in-Can and Vitrification Can-in-Can alternatives.

However, if the technologies are very different it is difficult to intuitively grasp the difference between the alternatives. For example, the Vitrification Can-in-Can alternative receives an overall utility score of 0.6905 while the Existing Reactor, Existing Facilities alternative scores 0.6676. But, what does this difference imply about the relative superiority of the vitrification alternative?

The overall utility score of each alternative is just a normalized combination of the scores on each measure. The choice of normalizing it to a range of (0,1) is arbitrary and done for mathematical convenience. This comment implies that one could choose to normalize the scores to any particular metric. Table 3 presents such a re-normalization for two measures: Life-cycle Costs, and Start Year.

Ceramic Can-in-Can	0.6907			
Vitrif. Can-in-Can	*0.6905*			
Difference	0.0002			
Measure	Ceramic Can-in-Can Level	*Vitrif. Can-in-Can* Level	Difference	% of Total Difference
MAX Inventory*	0.0400	5.0000	0.0010	463.30
*# processing steps**	*15*	*13*	*-0.0006*	*-261.50*
# processing steps	*15*	*13*	*-0.0003*	*-127.54*
New Long Term Jobs	2,402	2,300	0.0000	20.90
MAX Pu Inventory	2.00	5.00	0.0000	19.18
Wrkr. Rad. Deaths	*1.5*	*1.4*	*0.0000*	*-10.24*
Transport Deaths	*1.1910*	*1.1710*	*0.0000*	*-4.10*
Pub. Rad. Deaths	0.0000	0.0000	0.0000	0.00
Accident Risk	0.0000	0.0000	0.0000	0.00

Table 2 – Comparison of Vitrification Can-in-Can vs. Ceramic Can-in-Can

	Base Case MAU Model Score	*Life-cycle Costs*	*Life-cycle Costs Differential*	*Start Year*	*Start Year Differential*
Global Weight		*0.1733*		*0.0914*	
Ceramic Can-in-Can	0.6907	$7,140	$0	55.76	0.00
Vitrif. Can-in-Can	0.6905	$7,145	$5	55.79	0.03
Exist LWR, Exist Fac	0.6676	$7,674	$534	59.56	3.80
CANDU	0.6295	$8,554	$1,414	65.81	10.05
Borehole (Direct)	0.6249	$8,661	$1,521	66.57	10.81
Exist LWR, Green Fac	0.6211	$8,749	$1,608	67.19	11.43
Borehole (Immob)	0.6151	$8,886	$1,746	68.17	12.41
Vitrif. Adj. Melter	0.6101	$9,002	$1,862	68.99	13.23
Electrometallurgical	0.5976	$9,291	$2,151	71.05	15.29
Ceramic Greenfield	0.5543	$10,289	$3,149	78.15	22.39
Part Comp LWR	0.5441	$10,526	$3,386	79.83	24.07
Vitrif. Greenfield	0.5351	$10,733	$3,593	81.30	25.54
Evolutionary LWR	0.2676	$16,909	$9,768	125.20	69.44
Attribute Range	0.4231	$9,768		69.44	

Table 3 – Measure Equivalent Alternative Scores

The re-normalization procedure is based on the following question: how much of a particular measure would the decision maker be willing to sacrifice in exchange for an increase in the performance on all other measures from the current level to the best level? For example, one of the measures chosen for the re-normalization was Life-cycle Costs. Intuitively, the better the performance of an alternative on a measure, the less money one would have to sacrifice in exchange for increasing the performance of that measure to the best level. Inspection of Table 3 indicates this is the case. For example, the most preferred alternative Ceramic Can-in-Can requires the least amount of Life-cycle Costs to move its performance on all other measures from current to best.

Inspection of Table 3 provides some intuition as to how much more or less one alternative is preferred relative to another. For example, to answer the question originally posed, Table 3 indicates that the 0.0229 utility difference between Vitrification Can-in-Can and Existing LWR, Existing Facilities is equivalent to $529 million of Investment Cost or Life-cycle Cost or a 3.77 year delay in start time. Columns four and six of Table 3 show the difference between an alternative's measure equivalent score and that of Ceramic Can-in-Can, the most preferred alternative.

4.2 Sensitivity Analysis – changes in the weights

After the base case analysis was completed, the robustness of the alternative rankings was tested by varying the weights, the form of the multiattribute utility function, and the assumptions used in the base case. The first step was to consider the impact of varying the weights on the individual measures and goals in the MAU model. One-way, or single weight, sensitivity analyses were performed on the weights. Figures 12 and 13 display the sensitivity to the weight on Russian Cooperation for the reactor and non-reactor alternatives respectively. As the weight on Russian Cooperation increases, the reactor alternatives will be preferred to the non-reactors. The importance of this measure will be explored more fully in the section on deployment strategies.

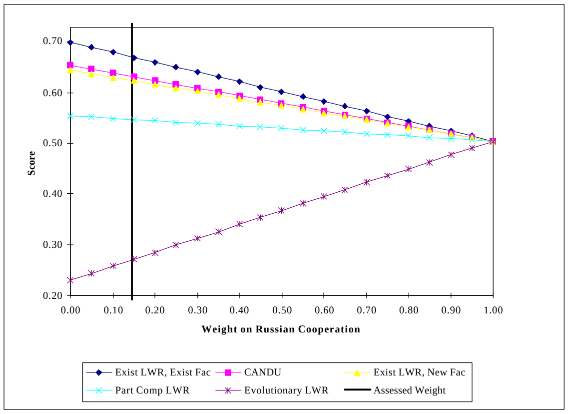

Figure 12 – Reactor Sensitivity to Weight on Russian Cooperation

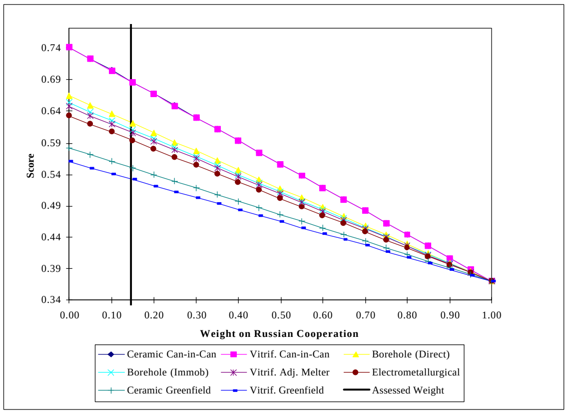

Figure 13 – Non-reactor Sensitivity to Weight on Russian Cooperation

Of particular importance were the weights at the highest level of the goal hierarchy on the objectives Non-proliferation, Operational Effectiveness, and ES&H. Many of the weights at lower levels in the hierarchy involved technical factors that limit the ability of non-experts to make the required judgments. These high-level weights represent policy judgments, and it is likely that these weights could be very different for different stake-holders or decision makers. To examine the impact of different weights sets on the decision, Figure 14 was created.

In Figure 14, the x and y-axes represent the weights for the Non-proliferation and Operational Effectiveness attributes respectively; the implied weight on the ES&H objective for any (x,y) pair is simply 1-x-y. Based on Figure 14, it appears that there are three alternatives in the set of the most preferred, and there are specific instances when each is preferred: when the weight on ES&H is relatively low, Ceramic Can-in-Can is preferred, when the weight on Non-proliferation is very high the Borehole (Immobilized) alternative is preferred, otherwise Existing Reactor, Existing Facility is preferred. Figure 15 represents a similar analysis except that the domain of weights is restricted so that they are consistent with the rank-ordering of the assessed weights: the weight on Non-proliferation is greater than the weight on Operational Effectiveness is greater than the weight on ES&H. As evident in Figures 14 and 15, in this case the choice the most preferred alternative reduces to a choice between three alternatives.

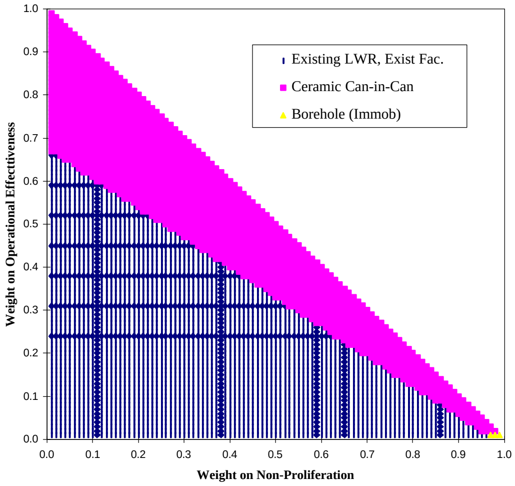

Figure 14 – Full Weight Range

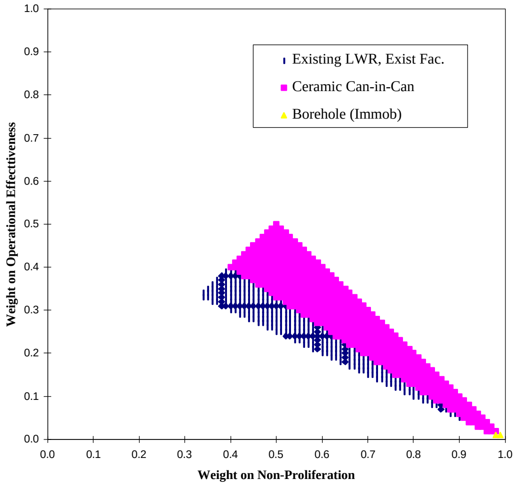

Figure 15 – Rank Order Weights

It is also useful to explore the results of changing all of the weights, perhaps simultaneously, in order to explore the robustness of the rankings of the alternatives in more detail. However, it would be extremely tedious to try to explore all reasonable combinations of values for the weights one at a time.

As an alternative to changing weights one at a time, weights were selected at random using a computer simulation program so that the results of many combinations of weights can be explored in an efficient manner. In addition, this simulation study provides a convenient means of testing the robustness of the MAU results to the use of the additive MAU model (1) rather than the multiplicative model (2).

The base case analysis assumes an additive MAU model (1). Although the assessment meetings with S&S and OFMD personnel indicated that this was a reasonable assumption, this simulation study allows for the possibility that some substructures of the MAU model are multiplicative, which implies some interaction effects between measures on the objectives. For example, though Safeguards and Security measures are always important, they may be considered relatively less important when the material in process is spent fuel rather than pits. In other words, a good score on Material characteristics might substitute or compensate for a lower score in Safeguards and Security. To incorporate this relationship, a multiplicative utility model must be used. The multiplicative model (2) includes the additive relationship (1) as a special case (Keeney and Raiffa, 1976). For a complete description of the simulation methodology, see Butler, Jia, and Dyer (1996).

Three types of simulation models were used to explore the sensitivity to various configurations of the weights. The first, random weights, made no assumption about the relative magnitude of the weights and is analogous to the 37-dimensional equivalent of Figure 14. The second was similar to Figure 15 in that the importance ordering of the weights was consistent with the assessed weights. Finally, each of the weights were treated as a response from a Dirichlet distribution (see DeGroot 1970 for a review of the Dirichlet distribution). The variance of this Dirichlet distribution can be easily varied. In each of these simulation approaches, the three sub-objectives of Theft and Diversion were assumed to be multiplicative.

Table 4 presents the results of the sensitivity analyses featuring 5000 randomly simulated sets of weights that preserve the same rank ordering as the assessed weights. The Ceramic Can-in-Can alternative is still superior according to its mean and mode with the Vitrification Can-in-Can and Existing Reactor, Existing Facilities alternatives close behind. Only two other alternatives, the Borehole (Direct) and Borehole (Immobilized) alternatives, were ever top ranked in any of the simulations. These results, and those from the other simulations, led to the conclusion that the base case rankings were robust with respect to the magnitude of the weights and the form of the utility function.

Alternative	*Best rank*	*Worst rank*	*Mean rank*	*Ranking Mode*	*5% Perc*	*25% Perc*	*50% Perc*	*75% Perc*	*95% Perc*
Exist LWR, Exist Fac	1	9	2.41	1	1	1	3	3	5
Part Comp LWR	3	12	11.71	12	10	12	12	12	12
Exist LWR, Green Fac	2	13	6.10	4	2	4	6	8	10
Evolutionary LWR	8	13	12.98	13	13	13	13	13	13
CANDU	2	13	8.19	11	4	6	8	11	12
Vitrif. Greenfield	6	12	10.50	11	8	10	11	11	12
Vitrif. Can-in-Can	1	5	2.36	2	1	2	2	3	4
Vitrif. Adj. Melter	4	10	6.24	6	5	6	6	7	8
Ceramic Greenfield	5	11	9.26	10	7	9	10	10	10
Ceramic Can-in-Can	1	6	1.78	1	1	1	2	2	3
Electrometallurgical	5	13	8.38	9	7	8	8	9	10
Borehole (Direct)	1	13	5.78	5	3	4	5	7	9
Borehole (Immob)	1	10	5.30	4	3	4	5	7	8

Table 4 – Results of Rank Order Weights Simulation

4.3 Sensitivity Analysis – changes in the assumptions

In addition to changes in the weights, it is also worthwhile to explore the sensitivity of the analysis to changes in some of the assumptions in the base case data and evaluation model. The first change that was explored is related to the notion of the "spent fuel" standard as the appropriate criterion for the end result of the plutonium disposition process. The definition of the spent fuel standard is that the material would be at least as difficult to retrieve and reuse for weapons manufacturing than retrieving and reusing plutonium spent fuel from commercial nuclear reactors (e.g., see the discussion of the spent fuel standard in [NAS 94]).

The sub-objective Irreversibility measures the degree to which the disposed material would be attractive to the host nation for potential retrieval and reuse, and rewards alternatives that go beyond the spent fuel standard and place the plutonium in a borehole. It was argued that no "bonus" value should be assigned to disposition alternatives that go beyond the spent fuel standard. In order to test the sensitivity of the results to adopting this point of view, the analysis was repeated after removing the Irreversibility sub-objective. This new model revealed that the Borehole (Immobilization) alternative moved down in rank, while the Borehole (Direct) alternative moved up. Nevertheless, both borehole alternatives were still ranked behind two immobilization alternatives and one reactor alternative.

The base case analysis is based on an aggregation of scores for the Non-Proliferation sub-objectives Theft and Diversion using the "weak link" concept. As discussed previously, the alternatives for disposal require that materials be processed in two or more facilities. The "weak link" approach uses the performance measures for the facility for each alternative that is identified as the lowest ranked of the facilities on the Theft and Diversion measures, respectively.

As an alternative, the Theft and Diversion measures can be based on a "weighted average" of the scores of the facilities associated with an alternative, where the "weight" is the proportion of processing steps required at the facility. The S&S team felt that the number of processing steps was an important measure of the need of a facility for safeguards against theft and for international assurances that diversion was not taking place. Therefore, the more processing steps at a facility, the higher that facility's weight should be when aggregating across facilities. A comparison of the overall results with the "weak link" rankings indicated minor variations; the Ceramic Can-in-Can and the Vitrification Can-in-Can alternatives switch ranks, as do the Borehole (Direct) and the CANDU.

The S&S team also requested a third model. Their idea was to build an absolute worst case based on the worst score for each measure across facilities. While this model represents no real facility or process, it does provide a worst case analysis. The rankings from this model were identical to the base case except that Existing LWR, Greenfield Facilities and Borehole (Direct) are virtually tied at sixth, rather than fifth and sixth as in the base case.

The current debate in Congress indicates a general agreement in both political parties that progress should be made toward a balanced budget within a relatively short time period. This focus suggests that initial investment costs may influence the choice among alternatives for plutonium disposition, with those alternatives that require smaller initial government funding commitments being favored. An alternative view is that initial investment costs should not matter if they are offset by lower costs or positive revenue streams in the future, so the only relevant measure of cost is properly discounted Life-cycle costs which includes the investment cost.

The alternatives most affected by the removal of Investment Cost as a separate measure are the Existing LWR, Greenfield and the Partially Completed LWR. In particular, the Partially Completed LWR moves from an overall ranking of twelve to a ranking of five when the investment costs are removed as a separate measure, although it is still ranked behind the two Existing LWR alternatives. Also, the CANDU alternative drops from fourth to seventh when its relatively low initial investment cost is no longer considered an advantage.

A final sensitivity analysis was performed on the base case assumption that all statistical lives had equivalent dollar value. Some previous government studies argued that workers were paid for the risks they encountered, and thus more money should be spent protecting the public. A model was constructed were the dollar value of statistical public fatalities was twice as large as that for workers. This change in the relative value of preventing fatalities had no affect on the base case analysis.

4.3 Sensitivity Analysis – Summary

This series of sensitivity analyses indicated that the ranking of the alternatives that was determined using the base case tradeoffs and assumptions is relatively insensitive to changes in these assumptions over reasonable ranges. Among the reactor alternatives, the Existing LWR, Existing Facilities and the CANDU alternatives are typically rated among the top two or three, and among the immobilization alternatives, the Vitrification and Ceramic Can-in-Can alternatives dominate the other alternatives.

However, the sensitivity analysis does provide some additional insights regarding a choice between a reactor and an immobilization alternative. According to the sensitivity analysis of the weights on the two measures associated with the International Cooperation sub-objective, the reactor alternatives become relatively more attractive if the weight on the measure Russian Cooperation is increased. Since one of the primary goals of the OFMD plutonium disposition effort is to prevent the proliferation of nuclear weapons, Russian cooperation could become a key issue, and become an extremely important consideration in the final choice of a US disposition strategy.

5. Hybrid Deployment Strategies

The base case analysis indicates that the four most desirable alternatives are, in order, Ceramic Can-in-Can, Vitrification Can-in-Can, Existing Reactor in Existing Facility, and CANDU. In the previous section, the results of the sensitivity analyses showed that these four alternatives are ranked at the top of the list for a wide range of possible weights that could be assigned to the measures. However, two important considerations were not fully captured by the analysis: the ability to influence the Russians to pursue a reciprocal disposition path and the risk of failure should a single technology be pursued. In this section we will discuss how parallel development of several technologies can better address these two concerns.

5.1 Influence on Russia and Other Nations

A key difference between the reactor and immobilization disposition alternatives is that irradiation of MOX fuel in reactors converts the material from a weapons grade to a reactor grade form of plutonium. It is more difficult, although not impossible, to fabricate weapons from reactor grade plutonium. The importance that Russia and other countries will place on this isotopic degradation is not known. Several studies have addressed this issue (see [NAS 95] and [DOE-NN 97]).

If Russian policy is to insist on isotopic degradation, Russia may not consider a U.S. disposition effort featuring immobilization as sufficient to meet the spent fuel standard for weapons grade plutonium. Further, the Russians may be more likely to join a reactor-based disposition program. However, if the Russian policy did not require degradation of fissile material, the U.S.'s most cost effective and timely course of action would be to pursue immobilization of surplus plutonium in ceramic or glass material due to the overall performance measures of those alternatives as highlighted in the previous section.

Of course, DOE could pursue a joint development approach featuring one reactor technology and one immobilization technology. This would lead to higher initial investment costs, but intuitively, such a "hybrid" should provide additional flexibility in light of the uncertainty about Russian policies, as well as other uncertainties cited in DOE-TSR 96. A "hybrid" alternative is defined here as a simultaneous deployment of two or more technologies; i.e. the investment in R&D, licensing and construction for two or more technologies.

In order to evaluate this strategy, the measure Russian Cooperation will be replaced with a probability distribution over the likelihood that that Russians will require degradation of fissile material. The weights on the other measures will be re-scaled so that the original ratios among the weights are maintained and the sum of the re-scaled weights is one.

This analysis is illustrated by the decision tree shown in Figure 16. In this diagram, decisions nodes are represented by squares with alternatives shown as paths through the tree, uncertainty is represented by circles, and endpoints or values received by proceeding down a path are represented by triangles. The logic represented by this tree is as follows: first, the U.S. must select a disposition strategy without knowing the future Russian stand on isotopic degradation; second, after the U. S. announces its decision, the official Russian policy will become known; finally, the U.S. will have to react to the Russian announcement. In Figure 16, the probability that Russia requires isotopic degradation is shown as 0.50 for illustrative purposes only. No attempt has been made to elicit an estimate of this probability from relevant experts. Instead, the objective of this analysis will be to identify the best development strategy for different values of the probability.

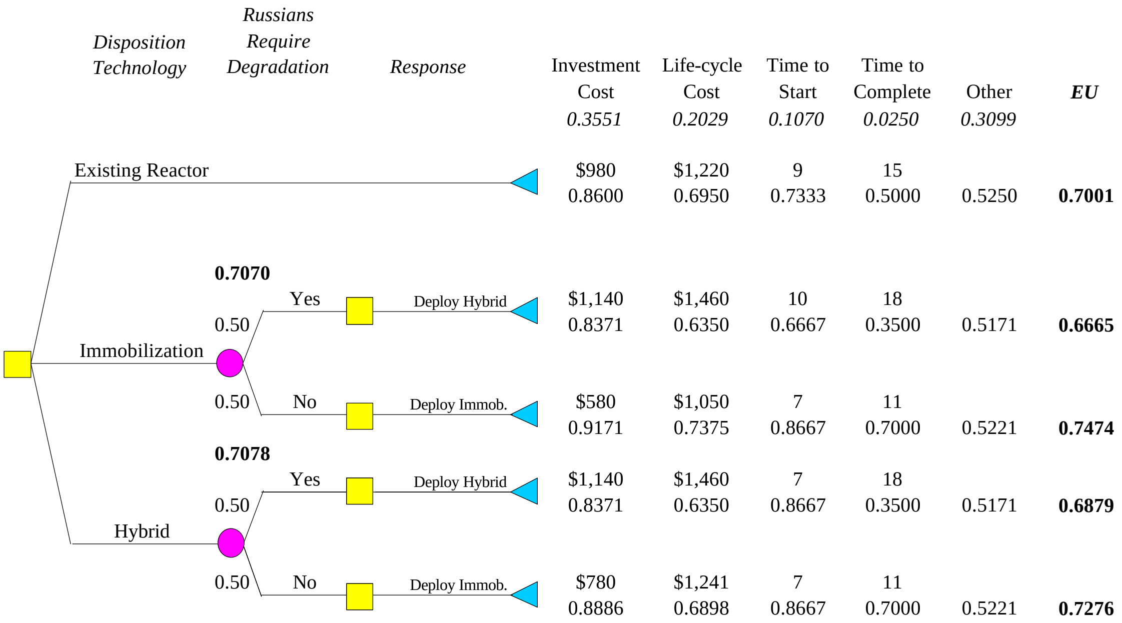

Figure 16 – Decision Tree for Existing Reactor in Existing Facility, Immobilization and Hybrid Disposition Alternatives

In this tree we have simplified the selection of a technology to a choice of three alternatives. From top to bottom in Figure 16 they consist of, Existing Reactor in Existing Facility, Immobilization, and the Hybrid. The Existing Reactor in Existing Facility was selected because it was the highest ranked reactor alternative. If the U.S. opted to follow this disposition technology, the Russian policy on degradation would be irrelevant; in either case, the reactor would not conflict with Russian policy. Therefore, there is no uncertainty associated with this alternative. With this technology, the U.S. program would expect the Investment Cost, Life-cycle Cost, Time to Start, and Time to Complete stated previously in this report. These measures were separated from the others in the objectives hierarchy because they are the most likely to be impacted by a hybrid deployment strategy. In addition, the thirty-two other measures (excluding Russian Cooperation) would contribute to the existing reactor's ability to achieve the disposition objectives. In Figure 16, these measures are all combined into a single expected utility score, 0.7003, using the methodology described in the previous sections.

The Vitrification Can-in-Can was arbitrarily selected as the immobilization technology because it was essentially tied with its ceramic counterpart as the most preferred immobilization alternative. If the U.S. were to deploy an immobilization (only) alternative, the announcement of Russian policy would be an important factor in valuing this choice. If the Russians require degradation, we have assumed that the Russians would not begin to dispose of their stockpiles due to their dissatisfaction with U.S. proliferation assurance. In response to this requirement, we have assumed that the U. S. would then begin to deploy the existing reactor alternative and use the immobilization plant to dispose of material that is not suitable for MOX fuel. The costs associated with the hybrid alternative would ultimately be incurred, but the reactor portion of the hybrid schedule is assumed to be delayed by 1 year. Accordingly, reactor operations with isotopic degradation of weapons grade plutonium would begin in 10 years, and the Russians would regard plutonium disposition to have begun at this point. If the Russians do not require degradation, the immobilization program would proceed with no additional cost overruns or schedule delays.

If the Russians require degradation, the utility of the immobilization alternative is 0.6666, while if it is not required, the utility is 0.7474. Using the assumed 50-50 probability, the expected utility of the immobilization alternative is 0.50*0.6665 + 0.50*0.6666 = 0.7070 as shown in Figure 16.

The final alternative is the hybrid deployment strategy. This alternative would simultaneously invest in R&D and licensing activities for both the immobilization and reactor technologies. This dual commitment leads to higher front-end investment costs, but there are no schedule delays in the future because the U.S. could deploy the reactor component of the hybrid without delay should isotopic degradation be required. Calculating the expected utility of the hybrid alternative would give 0.7078 per Figure 16.

If the U.S. were confident that the probability that the Russians would require degradation of fissile materials were 50%, our analysis would be complete and we would recommend the hybrid deployment because 0.7078 > 0.7070 > 0.7001. However, this probability was arbitrarily selected for illustrative purposes, so there is no basis for believing that it is correct. Therefore, a sensitivity analysis was performed over the complete range of possible probability estimates

Figure 17 indicates that given these three alternatives, each is preferred over a given range of the probability that degradation is required. If a decision maker believes this probability is less than 0.48, immobilization (only) should be the technology utilized. But, if this probability is between 0.48 and 0.69, the dual deployment hybrid is superior. Only if the probability is greater than 0.69 will the Existing Reactor, Existing Facility prove to be superior.

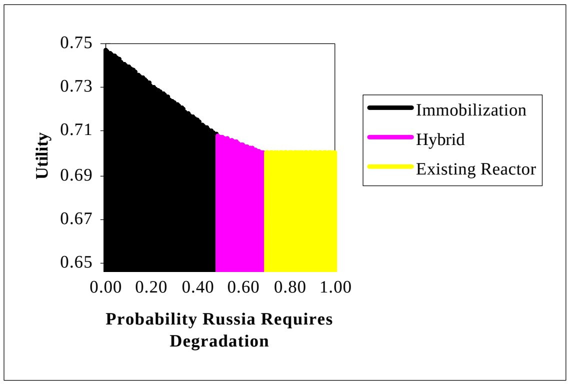

Figure 17 - Existing LWR Hybrid Evaluation as a Function of Probability

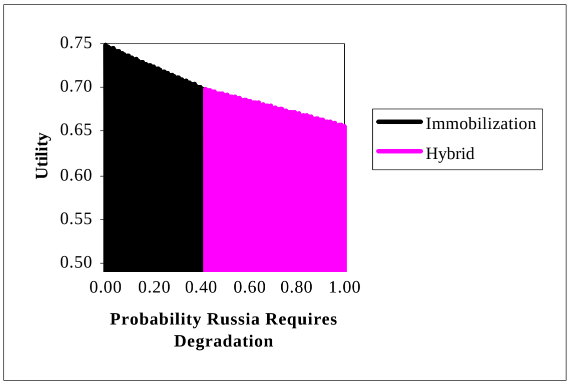

Figure 18 - CANDU Hybrid Evaluation as a Function of Probability

The other hybrid combination that was considered is similar except that the second highest reactor alternative, CANDU, was used. A tree similar to Figure 16 was developed and the probability profile in Figure 18 was generated. Based on Figure 18, when the CANDU reactor is considered, it is never optimal to implement the CANDU alternative alone. If the probability that degradation is required is between 0.0 and 0.41, the immobilization (only) alternative should be pursued. If this probability is greater than 0.41, the hybrid is superior.

Based on this analysis, if the probability that Russian policy requires degradation is estimated to be between 0.41 and 0.69, a hybrid development plan is superior, if either reactor alternative is considered. The reason for this is that the incremental cost associated with a hybrid is offset by the ability of the hybrid to quickly ramp up each technology component after the policy becomes known.

5.2 Other Advantages of the Hybrid

Other factors besides Russian insistence on isotopic degradation suggest the need for a hybrid strategy for plutonium disposition. For example, international concerns about the civil use of weapons plutonium in commercial reactors may delay the deployment schedule for the reactor alternative. Similarly, the R&D necessary to qualify the immobilized plutonium for disposal in a geologic repository may identify problems that require more time and budget to resolve than initially anticipated. Thus, the institutional and technical uncertainties associated with each alternative suggest the need for a “hybrid” approach in order to avoid delays in deployment of a disposition alternative. Therefore, it is logical that the ROD recommends proceeding with the parallel development of two of the highest ranked alternatives from this analysis, the Ceramic Can-in-Can immobilization alternative (or possible the Vitrification Can-in-Can variant) and the Existing LWR, Existing Facilities reactor alternative.

6. Summary and Conclusions

In conjunction with recent strategic arms reduction negotiations between the United States and the Russian Federation, the United States has identified approximately 50 metric tons of weapons-usable plutonium as surplus to national defense needs. The Department of Energy, Office of Fissile Materials Disposition, has been charged with selecting and developing technologies for transforming this plutonium into forms that are more difficult to use in weapons. Many disparate issues must be considered when evaluating candidate technologies for disposition including nonproliferation, cost, schedule, environmental, safety, and health factors. Three Department of Energy reports summarize these factors for the thirteen technology deployment alternatives that were considered for plutonium disposition ([DOE-SCR 95], [DOE-TSR 96], [DOE-PEIS 96]).

The decision analysis effort presented in this paper offered an approach for integrating all of these factors. In the early stages of the analysis, the model proved useful for highlighting data inconsistencies and focusing attention on discriminators among the technologies. The development of the objectives hierarchy and the tradeoff discussions were also useful exercises, independent of the modeling effort.

The preliminary results of the rankings provided by the model were presented on two occasions to the Department of Energy team while the alternatives were still being defined, and provoked considerable discussion. Some of the alternatives were refined to improve them in areas where relative weaknesses were identified. While the report outlining our recommendations was officially published after the DOE announcement concerning the selection of an alternative [DOE-ROD 97], we provided the decision makers with numerous drafts prior to the final decision as new information became available. The last drafts were qualitatively indistinguishable from the published version of our efforts, and aided DOE’s explanation of the final recommendation by identifying meaningful discriminators among the alternatives. Further, this analysis also presents the only explicit analysis of the benefits associated with a dual-deployment or hybrid strategy.

The model also demonstrates to the public that a broad range of factors were included in the evaluation, and the sensitivity analysis demonstrated that many of these factors are not significant discriminators among the alternatives. The complete report of the decision analysis effort is available to the public on the World Wide Web[2]. Overall, the effort made a positive contribution to the decision making process for this national security mission.

---

[2] http://www.pu.org/main/reports/reports.html (Amarillo National Resource Center for Plutonium web site). All reports at this site are in pdf format.

References

Butler, J, Jia, J. and Dyer, J (1997). Simulation Techniques for the Sensitivity Analysis of Multi-criteria Decision Models. *European Journal of Operational Research,* forthcoming.

DeGroot, M.H. (1970), *Optimal Statistical Decisions*. McGraw-Hill Book Company.

[DOE-LIPS 95] *Laboratory Integration Prioritization System Report.* United States Department of Energy. Draft 2.1. June 6, 1995.

[DOE-NN 97] *Nonproliferation and Arms Control Assessment of Weapons-Usable Fissile Materials Storage and Plutonium Disposition Alternatives, DOE/NN-0007.* United States Department of Energy; Office of Fissile Materials Disposition, January 13, 1997.

[DOE-PEIS 96] *Storage and Disposition of Weapons-Usable Fissile Materials Final Programmatic Environmental Impact Statement, DOE/EIS-0229.* United States Department of Energy; Office of Fissile Materials Disposition, December 1996.

[DOE-ROD 97] *Record of Decision for the Storage and Disposition of Weapons-Usable Fissile Materials Final Programmatic Environmental Impact Statement.* United States Department of Energy; Office of Fissile Materials Disposition, January 14, 1997.

[DOE-SCR 95] *Summary Report of the Screening Process to Determine Reasonable Alternatives for Storage and Disposition of Weapons-usable Fissile Materials, DOE/MD-0002.* United States Department of Energy; Office of Fissile Materials Disposition, March 17, 1995.

[DOE-TSR 96] *Technical Summary Report for Surplus Weapons-Usable Plutonium Disposition, DOE/MD-0003 Rev. 1.* United States Department of Energy; Office of Fissile Materials Disposition, October 31, 1996.

Dyer, J., Edmunds, T., Butler, J. C., and Jia, J. (1997). Evaluation of Alternatives for the Disposition of Surplus Weapons-usable Plutonium. Amarillo National Resource Center for Plutonium Technical Paper. http://www.pu.org/main/reports/reports.html

Dyer, J., Edmunds, T., Butler, J. C., and Jia, J. (1996). A Proposed Methodology for the Analysis and Selection of Alternatives for the Disposition of Surplus Plutonium. Amarillo National Resource Center for Plutonium Technical Paper. http://www.pu.org/main/reports/reports.html

Dyer, J. S. and Lorber, H. W. (1982). The Multi-attribute Evaluation of Program-Planning Contractors. *OMEGA, 6,* 673-678.

Edmunds, T., Koopman, R., and Myers, B.(1995). *Performance Measures for Evaluation of Plutonium Disposition Alternatives.* Draft. March 15, 1995.

Keeney, R. L. (1980). *Siting Energy Facilities.* New York: Wiley.

Keeney, R. L. and Raiffa, H. (1976). *Decisions with Multiple Objectives.* New York: Wiley.

Keeney, R. L. and von Winterfeldt, D. (1994). Managing Nuclear Waste from Power Plants. *Risk Analysis, 14* , 107-130.

[NAS 94] National Academy of Sciences (1994). *Management and Disposition of Excess Weapons Plutonium,* Washington D.C.: National Academy Press.

[NAS 95] National Academy of Sciences (1995). *Management and Disposition of Excess Weapons Plutonium - Reactor Related Options,* Washington D.C.: National Academy Press.

Tengs, T.O., Adams, M., Pliskin, J.S., Safran, D.G., Siegel, J.E., Weinstein, M.C., and Graham, J.D. (1994). Five-Hundred Life-Saving Interventions and Their Cost-Effectiveness. *Harvard School of Public Health Working Paper*.

von Winterfeldt, D. and Schweitzer, E. (1997). An Assessment of Tritium Supply Alternatives in Support of the U.S. Nuclear Weapons Stockpile. *Working Paper: Finalist, 1997 Franz Edelman Award Competition.*