25 Choosing a Tritium Supply Technology for Nuclear Weapons: Reflections on a Controversial Decision Analysis

Detlof von Winterfeldt

ABSTRACT. Nuclear weapons require the periodic replacement of tritium, a radioactive gas that decays at approximately 5.5 percent per year. Since 1989 the United States had no tritium supply facility, and, because of the decay of tritium, its inventory was expected to fall below the required reserve level in 2011. To decide how to fill this projected gap, the Department of Energy assessed several tritium supply alternatives, including several types of new reactors, an accelerator, and the use of commercial reactors. This paper describes the decision analysis process, conducted in the mid-1990s, to support the decision by the Secretary of Energy to choose among the options. This process involved two rounds of analysis, several surprises and many adjustments. In the end the decision analysis was successful in shaping both the intermediate decision by then-Secretary O'Leary and the final decision by Secretary Richardson.

The Problem

Tritium is a necessary component of all nuclear weapons. Because tritium decays at a rate of approximately 5.5 percent per year, it must be replaced periodically. Over the past 40 years, the Department of Energy (DOE) has built and operated fourteen reactors to produce tritium and other nuclear materials for weapons purposes. In 1988, then-President Bush shut down the last remaining tritium production facility, a heavy-water reactor in Savannah River, GA. Since that time, no tritium has been produced and the DOE had to rely on recycled tritium from existing weapons to replenish decaying tritium sources.

The strategic arms reduction treaties known as START I and START II resulted in reducing the number of weapons in the stockpile of the US and Russia and made it possible for the US to rely on a tritium recycling program for a limited time. However, when this analysis was conducted in the mid-1990s, it was clear that a new tritium production source would eventually be required. Under the START I treaty a new tritium source would be required by 2005 and under the START II agreement a new source would be needed by 2011 (Figure 25.1).

The main issue was what technology to use for producing tritium. Several new reactor types were considered as well as a large linear accelerator. Purchase of tritium from foreign sources (Canada and Russia have an oversupply of tritium) was considered not acceptable for national security reasons. Using an existing commercial reactor was originally thought to be in violation of national policy

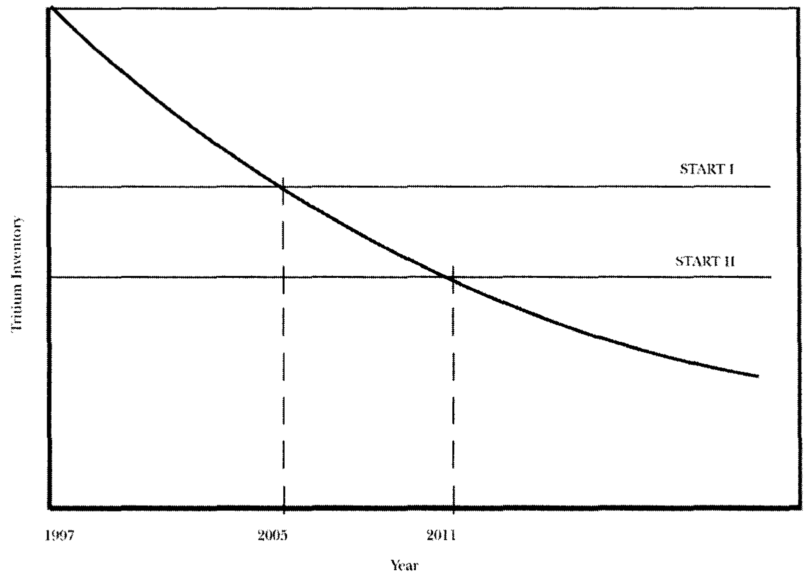

Figure 25.1. The U.S. tritium inventory without new production capabilities and levels to support START I and START II requirements.

and international agreements. Other issues concerned the siting of the tritium production facility and a choice of upgrading the existing tritium recycling facility at Savannah River versus building a new recycling facility co-located with the new production facility.

The stakes in this decision were high. First, the lack of a steady and reliable source of tritium was considered a threat to national security. Although most alternatives appeared to be able to meet the production requirements, significant uncertainties were associated with some. Second, the costs of the alternatives were very high. For example, building a new reactor or an accelerator can cost between $3 billion and $6 billion, and operating them can cost an additional $100 million to $250 million per year. Third, there were environmental concerns with most alternatives, including the production and disposal of radioactive waste and the risk of major accidents resulting in the release of radioactive material.

The Secretary of Energy at the time of this analysis was Hazel O'Leary. She had the ultimate responsibility of choosing among the tritium supply alternatives, sites and recycling facilities. Key participants in the decision-making process were the DOE Assistant Secretary of Defense Programs and the Under Secretary. Major stakeholders were the Department of Defense, Congress, public interest groups, and the private vendors of tritium production. The main concern of the Department of Defense and its Nuclear Weapons Council was to ensure timely tritium production in the required amounts. Congress and its appropriations committees were concerned with cost. Several public interest groups expressed strong opposition to building new nuclear reactors and concern about environmental impacts and employment. The private vendors lobbied for their own designs for tritium production.

Getting Started

In 1993, I was asked by the staff of Fluor Daniel, Inc., a major subcontractor to the DOE, to give several one-week short courses on using decision analysis as a tool to improve the key decisions that had to be made in DOE's Office of Reconfiguration. This office, which was part of DOE's defense program, was created to manage the major transitions of the weapons complex following the demise of the Soviet Union.

The short courses were received well and the director of the Office of Reconfiguration, Fluor Daniel and other subcontractors appeared to be committed to use decision analysis ideas and tools to address a variety of problems. However, in early 1994, the Office of Reconfiguration was split into three parts: a tritium production office, a plutonium disposition office, and a nuclear stockpile stewardship office. Both the tritium production office and the plutonium disposition office continued to use decision analysis in their major decisions (for the plutonium dispositions study, see Chapter 24 in this volume). I am not aware of any decision analysis activities regarding stockpile stewardship.

When the Office of Tritium Production was created, I was asked to lead a decision analysis of the choice of a tritium production technology, site, and recycling facility. I realized that this was a substantial job and would require help from several decision analysts as well as significant portion of my own time. I therefore asked for a leave of absence from the University of Southern California and recruited Ralph Keeney, Richard John, Robin Dillon (who worked at Fluor Daniel, Inc. at the time), and Robin Keller to work on this project. In the process, I also launched Decision Insights, Inc., which was incorporated on July 1, 1994, serving primarily this effort and several activities to support DOE's environmental management program.

One of the most important tasks in getting a decision analysis started is to identify the decision maker or decision makers involved in the problem and to establish appropriate reporting mechanisms. It was clear to me that the key decision makers were (in order from lower to higher): the deputy director and the director of the tritium production office, the assistant secretary for defense programs, and Secretary O'Leary. Reporting to a subcontractor, and possibly through an ill-understood chain-of-command, seemed to invite problems. I therefore requested to report directly to the director and deputy director of the Office of Tritium Production, thus bypassing the staff of Fluor Daniel, Inc., who were formally my contractors. This decision was absolutely crucial, as it turned out later, because it gave me open and continuous access to the decision makers, without filtering our analyses through layers of management.

Another important part of this analysis was to divide the work between different contractors. Fluor Daniel, Inc. was the main technical contractor and TetraTech, Inc. was the main environmental consulting company. Once my own reporting mechanism was established, I became, in effect, the lead decision analyst working with the directors of the office to oversee and guide the activities of the two major subcontractors. This also was an extremely important development as decisions made later about the direction of the decision analysis had substantial implications for the two contractors.

Physical proximity to the decision makers also turned out to be important. I set up an office in the Fluor Daniel complex in Alexandria, Virginia, with open access to office space in DOE's headquarters. As a result I was in virtually constant and direct personal contact with the key decision makers over the next eighteen months. This was important because many key decisions about the analysis occurred during casual meetings at the offices of Fluor Daniel or at DOE's headquarters.

When the analysis started in April, 1994, about seventy engineers, cost specialists, and environmental analysts worked on various aspects of the tritium project. Fluor Daniel, Inc. was charged with developing the engineering and cost data and estimates for the tritium supply alternatives. TetraTech, Inc., was charged with conducting the environmental analyses and developing the programmatic environmental impact statement (PEIS). A guideline to the contractors was that all estimates were to be based on existing reports and data. As a result, a major part of the early effort was to collect and catalogue several hundred boxes of existing information on the tritium supply alternatives. One of the major challenges was to organize the collection and assessment of information to support the decision by the Secretary of Energy and to draft a record of decision, the legal document presenting the Secretary's decision.

The project lasted for about eighteen months with an overall cost of approximately $20 million. The decision analysis effort cost about $1 million, not counting the work contributed by Fluor Daniel and TetraTech staff to this effort. Figure 25.2 shows a timeline of the project.

The Multiattribute Utility Analysis

The initial formulation of this decision problem included six tritium supply alternatives (a heavy-water reactor, two advanced light-water reactor designs, a modular high-temperature, gas-cooled reactor, and two version of a linear accelerator). The DOE wanted to evaluate these alternatives against numerous conflicting objectives, related to production assurance, cost, and environmental impacts. This concern with multiple objectives suggested the use of a multiattribute utility analysis (Keeney and Raiffa 1976; von Winterfeldt and Edwards 1986). This analysis, conducted in the second half of 1994, consisted of defining the objectives, collecting data to estimate how well the alternatives met them, and obtaining trade-off judgments from DOE officials to evaluate the alternatives in terms of production assurance, cost, and environmental impacts. In addition, we held several expert workshops to quantify the uncertainty in the data and estimates.

The analysis started in a fairly conventional way. Because it was not clear whether production technologies, site alternatives, and recycling options had some

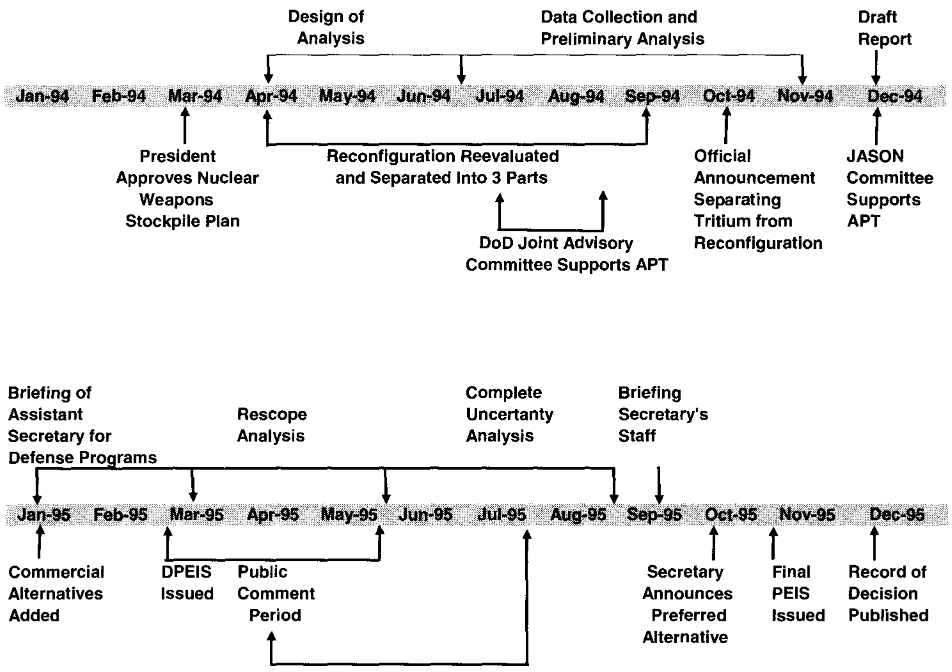

Figure 25.2. Timeline of the tritium decision analysis process.

interactions with respect to the objectives, we decided to analyze all possible combinations – fifty-four altogether:

1. Six production technologies
   a. Large advanced light-water reactor (LAWR)
   b. Small advanced light-water reactor (SALWR)
   c. Heavy-water reactor (HWR)
   d. Modular high-temperature gas-cooled reactor (MHTGR)
   e. Small accelerator (SAPT – designed to meet START II requirements)
   f. Large accelerator (LAPT – designed to meet START I requirements)
2. Five sites
   a. Idaho National Engineering Laboratory (INEL)
   b. Nevada Test Site (NTS)
   c. Oak Ridge Reservation (ORR)
   d. Pantex Plan (PANTEX)
   e. Savannah River Site (SRS)
3. Two recycling alternatives
   a. New recycling facility at any of the five sites
   b. Upgrade existing recycling facility at SRS

Although this generated sixty possible combinations, at Savannah River only upgrading the recycling facility made sense, reducing the set of alternatives to fifty-four.

There were many discussions about whether we could separate the technology choice from the site evaluation and either or both from the decision about recycling. At the early stage of this analysis, the argument prevailed that there could be significant interactions leading to this exhaustive list of combinations.

The objectives for the multiattribute utility analysis were obtained from reviewing the content of a previous programmatic environmental impact statement, previous studies of both technology choice and facility siting decisions and interviews with environmental and technical specialists of the DOE and its contractors. The three major objectives were:

1. Ensure production of tritium.
2. Protect the environment, health, and safety.
3. Reduce costs.

These were broken down as follows: Production assurance had three subobjectives: meet schedule; provide sufficient capacity; Ensure availability. The schedule objective came from the concern that some of the alternatives may not meet the 2011 deadline for producing tritium. The capacity objective was to address concerns, primarily with the accelerator, about the annual amount of tritium that could be produced with the technology. The availability concern addressed issues about the continuity of production throughout the many years of the planned life cycle of the facility.

The environmental objectives were broken down into seventeen objectives ranging from worker risks due to normal operations, to public risks due to major disasters, to pollution and socioeconomic impacts. Costs were broken down into the total life-cycle cost of designing, constructing and operating the tritium supply facility, the cost of the recycling facility, and the possible revenues from electricity production.

Measures were defined for each of the twenty-three objectives. For example, the schedule objective was measured in terms of the year and month at which tritium from the production facility would become available. Worker risks were measured in terms of the expected number of worker cancers and other fatalities due to the operation of the facility over 40 years. Cost was expressed as total life-cycle costs, measured in discounted 1995 dollars.

At this point, the analysis was set up to provide consequence estimates of fifty-four alternatives on twenty-three objectives – a daunting task. Fortunately, many of the cells in this table of alternatives by objectives had identical or similar consequences. For example, the production assurance assessments did not depend on the site chosen. In addition, we had many experienced assessment teams and could build on a draft environmental impact statement to pull together the requisite estimates. The technical, cost, and schedule estimates were provided by the staff of Fluor Daniel, Inc. The environmental estimates were provided by the staff of TetraTech, Inc. In all cases, we organized the elicitation in small group workshops. In some cases (schedule, capacity, availability, and cost), we obtained probabilistic estimates from the experts.

Table 25.1. Aggregate comparisons of six tritium production alternatives

	Production assurance	Environment, health and safety	Total life-cycle cost ($M)
HWR	99.98%	71	$5,551
Large ALWR	99.99%	58	$5,940
Small ALWR	99.97%	65	$4,444
MHTGR	99.89%	72	$6,529
Nominal APT	97.29%	82	$5,943
Full APT	97.39%	82	$6,907

Note: Production assurance assumes no technical or institutional problems; the production assurance for the nominal APT is for START II requirements.

To aggregate the three subobjectives of production assurance into one measure, we developed a model that simulated the production cycle of each technology over 40 years (see below for more detail). This production simulation model gave us an estimate of the expected percentage of tritium production at or above the required amount over the entire 40-year life cycle of production. This percentage ranged from 99.99 percent for the light-water reactors to 97.29 percent for the nominal accelerator.

We aggregated the other sets of subobjectives using standard multiattribute weighting techniques (swing weights and pricing out) involving interviews with four DOE officials. For the environment, health, and safety objectives we used an aggregate utility measure ranging from 0 (worst) to 100 (best). For the aggregate cost we used discounted (at 3.8 percent) total life-cycle cost. We decided early on to assign weights only to the subobjectives under the three main objectives (production assurance; cost; and environment, health and safety), instead of weighting these three subobjectives as well. This was done primarily to leave the decision makers some room for discussion and for deliberating possible tradeoffs among the three major objectives.

As the multiattribute utility analysis progressed, it became clear that only some minor interactions occurred between sites and technologies. As a result, we were able to report the results separately for technologies, sites, and recycling alternatives. The results are shown in Tables 25.1 and 25.2.

Regarding technologies, all alternatives show a high degree of production assurance with the large ALWR being best and the accelerator being worst. From an environmental point of view, the accelerator is the best alternative because it has no severe accident risks and produces no or little waste. The large ALWR is worst for exactly the opposite reasons. Regarding cost, the small ALWR is best, the accelerator is worst. As a result, the decision comes down to a tradeoff among the three main objectives. If a large weight is placed in environmental concerns (60 percent or more), the accelerator is the preferred alternative.

Regarding sites and recycling options (Table 25.2), the best environmental site was the Nevada Test Site; the worst was the Savannah River Site. A new recycling plant co-located with the new production facility was always better from

Table 25.2. Aggregate comparisons of site and recycling alternatives (for the small ALWR)

	Environment, health, and safety	Life-cycle cost ($M)
Nevada Test Site		
New recycling	69	$5,178
Upgrade at SRS	65	$4,444
Idaho National Eng. Lab		
New recycling	59	$5,198
Upgrade at SRS	56	$4,184
Savannah River Site		
New recycling	n/a	n/a
Upgrade at SRS	47	$4,158
Oak Ridge Reservation		
New recycling	53	$4,688
Upgrade at SRS	50	$4,098
PANTEX Site		
New recycling	58	$4,940
Upgrade at SRS	54	$4,250

an environmental perspective. In terms of costs, the Oak Ridge Site was least expensive; the INEL Site was most expensive. Recycling at Savannah River was always cheaper than building a new recycling plant at a different site. These results are shown in Table 25.2 for the small ALWR, but the results are similar for all other technologies. Overall, the finding was that the choice of a site or recycling facility made less difference in terms on environmental impacts or costs than the choice of a production technology.

With this material in hand, we briefed the assistant secretary for defense programs in January 1995. The results were noted by the assistant secretary and his staff, but many issues were raised about the risks and uncertainties surrounding cost and production assurance. As a result, we were asked to start a second phase of analysis, which focused on these risks. Before describing the findings of this analysis, a short interlude about the politics of this decision process is needed.

Interlude about the Politics of Decision Making

While we were conducting the multiattribute utility analysis, two events occurred that indicated a strong support for the accelerator by key decision makers. In the summer of 1994, the Joint Advisory Committee of the DOE and the Department of Defense issued a statement supporting the use of the accelerator for tritium production and in November of 1994 the JASON committee, a scientific advisory committee to the DOE, also supported the accelerator. Both committees considered the accelerator to be a sound technology for producing tritium that was environmentally clean and enjoyed public support. Our analysis team received several messages from DOE staff suggesting that the accelerator may be the preferred alternative, even though our analysis was not completed yet.

This created a potential dilemma. On one hand, I could see how the accelerator could be defended using a very high weight on environmental objectives. At the same time, I questioned whether the relatively minor environmental benefits of the accelerator could outweigh the moderate concerns about production assurance and the significantly higher costs. In the briefing with the assistant secretary I therefore pushed the issue of this tradeoff, asking, in essence, why the accelerator was so strongly preferred. His arguments were that the accelerator was environmentally friendly, would face little opposition or delays, and, although costly, the costs could be reduced through alternative designs and research and development. The analysis left room for this argument, but I pointed out that these benefits came at a significantly increased costs.

Meanwhile, there were other – implied or unspoken – concerns. No nuclear reactor had been built in the United States for 20 years and it seemed unlikely that the Clinton administration wanted to do anything to revive a nuclear power program. In addition, the accelerator was a potential research and development bonanza for the national laboratories, because many technical issues needed to be resolved prior to building it. In contrast, all the reactor options would be developed and constructed by the private sector with little laboratory assistance.

None of these political agenda items were made explicit in the course of this analysis, but it is hard to believe that they did not have an influence on the decision making process. Our briefing on the relative advantages of the accelerator versus a reactor was clearly not a strong endorsement of the accelerator. We also had identified many uncertainties and risks of all technologies, including risks of some of the reactors not receiving licenses and risks of schedule and cost overruns. Whatever the reasons, the result of the briefing with the assistant secretary was to continue our work and conduct a more-detailed risk analysis on production assurance and cost.

As we began to rescope the analysis to include a more thorough risk analysis, five new alternatives were proposed, three that made use of commercial reactors and two that represented advanced designs of the heavy-water reactor and the modular high-temperature gas-cooled reactor. The commercial reactors had been considered in the early analysis stage, but they had been eliminated because it was unclear whether a utility would want to sell a commercial reactor to the DOE or provide tritium production services. With the emerging deregulation of the utility industry, the commercial options appeared more attractive in 1995 and, with their lower costs, had support in the DOE. However, there were significant policy and international treaty issues that needed to be resolved before a commercial option could be viable.

In March 1995, the DOE issued the draft programmatic environmental impact statement (DPEIS) with an analysis of some thirty environmental impacts. The DPEIS results made it clear that there were only three environmental discriminators among the tritium production alternatives: spent fuel production, low-level waste production, and severe accident risks. About the same time, the cost and the uncertainties associated with it became a major issue. To resolve this issue, the secretary issued an independent study by Putnam, Hayes, and Bartlett (1995) that provided cost ranges for all alternatives, but no best estimate or probability distribution. These ranges were substantial, covering, in some cases, $5–10 billion.

In mid-1995, it was clear that the unresolved issues were the uncertainties associated with the schedule aspects of production assurance and with cost. In addition, issues remained about the political viability of the commercial option. In the remainder of the analysis, we therefore focused on quantifying these uncertainties and documenting the results from the environmental impact study. Thus, what began as a fairly standard multiattribute utility analysis turned into three related efforts: A risk assessment of production assurance, a risk assessment of costs, and a simple environmental analysis. In the end, there was no need for making tough trade-offs. The results spoke for themselves.

The Risk Analysis

We began the risk analyses in early 1995, shortly after briefing the assistant secretary. At this point, we were asked to consider ten tritium production alternatives:

1. Large HWR.
2. Small AHWR.
3. Steam-cycle MHTGR.
4. Direct-cycle MHTGR.
5. Large ALWR.
6. Small ALWR.
7. Accelerator production of tritium (full APT).
8. Purchase an operating commercial light-water reactor (CLWR).
9. Purchase a partially completed CLWR.
10. Purchase irradiation services from an operating CLWR.

Production Assurance Analysis

The three production assurance subobjectives contribute to the major objective of providing enough tritium in time to maintain the tritium inventory at the START II level. To analyze how well the ten tritium supply alternatives met this overall objective, we developed a dynamic model that simulated production patterns over 40 years. This model had three uncertain inputs: the schedule, the production capacity, and the annual availability (see Figure 25.3). To quantify the uncertainties about these inputs, we conducted several formal expert elicitation processes (Hora and Iman 1989; Keeney and von Winterfeldt 1991). These expert elicitations included explicit considerations of possible technical and institutional problems with the implementation of the technologies.

SCHEDULE RISK ANALYSIS. For each alternative, Fluor Daniel, Inc. (1995a) developed a base-case schedule divided into the several components ranging from

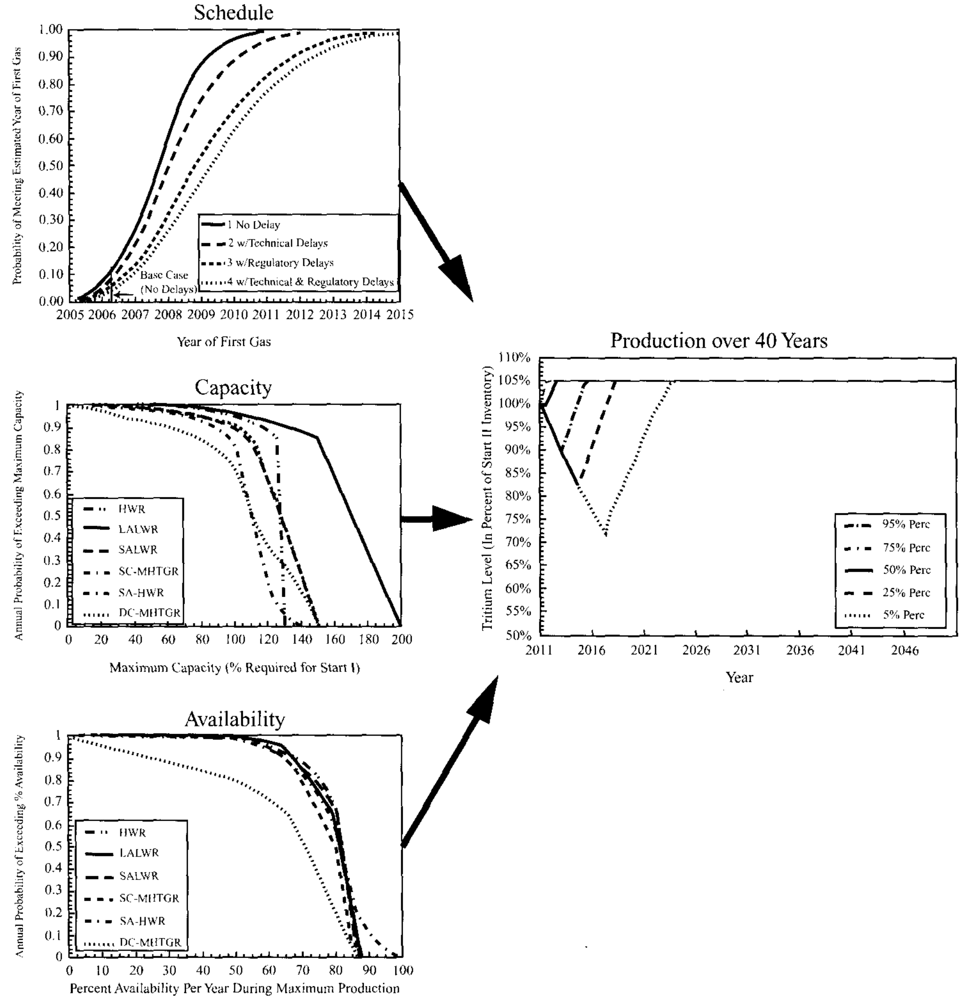

Figure 25.3. Production risk analysis model.

conceptual design to pouring first concrete, to start up and operations. The total length of the schedule was determined with consideration of the critical path relationships between the schedule components.

To assess the schedule uncertainties, we convened three panels of schedule specialists, with 10 technical staff members of the DOE office of reconfiguration, its contractors and consultants in each panel. Participants first made their three probability estimates independently, and, subsequently, we presented all estimates for discussion. During the discussion, participants could revise their initial estimates, but they were not required to reach consensus. After calculating each expert's overall schedule distribution we averaged the distributions across experts, giving each expert equal weight.

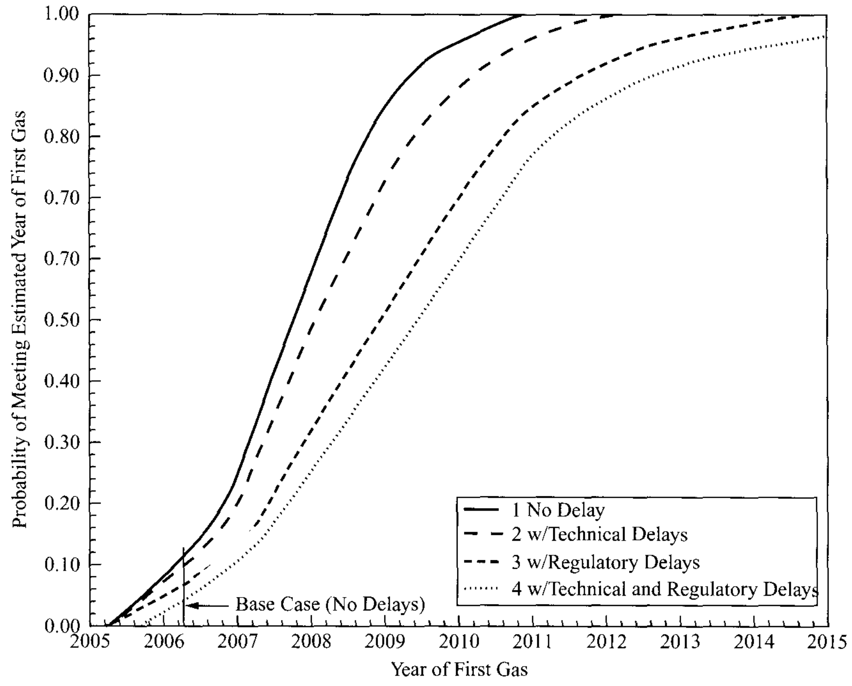

Figure 25.4. The results of the schedule risk analysis for the small ALWR.

In addition to this assessment, we asked the experts to identify technical and institutional problems that could cause delays of greater than one year. We also asked them to provide probabilities for these problems and estimate their schedule impacts. We folded these assessments of technical and regulatory delays into the simulation.

Figure 25.4 shows the resulting cumulative probability distributions over the year of delivering the first tritium gas for the small ALWR. The schedule provides for a high probability of meeting the 2011 start date (0.98 without delays and 0.78 with both technical and regulatory delays). The technical delays were due primarily to the possible failure of the new passive safety system of this reactor. The regulatory delays resulted from possible complication in the licensing process. Table 25.3 provides a summary of the schedule risk analysis, with the first column showing the base-case schedule estimate by Fluor Daniel, Inc, and the other columns showing the probabilities of meeting the 2011 start-up date with and without considering delays.

Two observations about this table are important. All base-case schedule estimates are within the 2011 date. However, the schedule risk analysis shows that these estimates may be highly optimistic in some cases. For example, the heavy-water reactors and the gas-cooled reactors have a fairly low probability of meeting the 2011 start date, when considering technical and institutional delays. The light-water reactors and the accelerator provide the highest assurance of meeting

Table 25.3. Summary of the schedule uncertainty analysis

Tritium supply alternative	First tritium (Base case)	Probability of Meeting 2011 Without delays	Probability of Meeting 2011 With delays
Large HWR	2009	0.82	0.4
Small advanced HWR	2006	0.84	<0.40
Steam-cycle MHTGR	2009	0.60	0.22
Direct-cycle MHTGR	2010	0.14	0.14
Large advanced LWR	2007	0.92	0.78
Small advanced ALWR	2006	0.98	0.78
Accelerator	2008	0.92	0.76
Purchase operating LWR	2005	>0.99	see note
Purchase partially complete LWR	2006	>0.99	see note
Purchase irradiation services	2004	>0.99	see note

Note: Commercial production requires congressional approval.

the 2011 date. The commercial options can be ready in a few years and provide the highest assurance of meeting the schedule. However, there is the possibility of institutional delays or even a chance that these options are institutionally infeasible.

CAPACITY ANALYSIS. We defined production capacity as the maximum amount of tritium that could be produced in one year, assuming 75 percent availability for tritium production. The results were expressed as percentages of the START I requirement. Most tritium alternatives had no problems with meeting or exceeding the availability requirement, but there were substantial uncertainties.

Fluor Daniel, Inc. (1995b) prepared a summary of the available information on production uncertainties. This summary listed the factors that were likely to increase or reduce the capacity of the ten alternatives. From this information, eleven DOE and contractor staff members provided estimates of the probability that the production capacity would exceed 50, 75, 100, 125, or 150 percent of the START I goal. Participants first wrote down their initial responses for each alternative. Subsequently, they compared and discussed all responses for that alternative and revised their estimates, if they wished to do so. Much of the discussion focused on the plausible upper bounds of the production capacities, and the group came to a consensus on these bounds.

We fit probability distributions to the averages of the individual probability estimates (Figure 25.5). All reactors have a very high probability of meeting the capacity required by START I. Only the direct-cycle MHTGR and the small AHWR have probabilities of meeting the START I goal below 0.90.

AVAILABILITY ANALYSIS. All ten tritium supply alternatives were designed to operate 75 percent of the time. To assess whether they could meet or exceed this goal, we defined availability as the percentage of time that the production facility would be capable of producing tritium during any one year of maximum production.

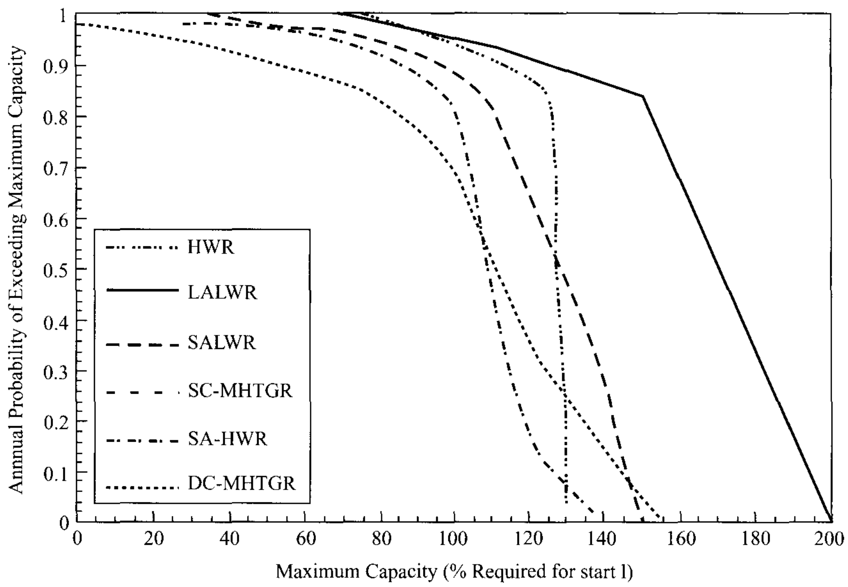

Figure 25.5. Results of the capacity risk analysis.

This availability is referred to as "maximum availability" as opposed to the actual availabilities that one may see during the course of a 40-year history of production.

Fluor Daniel, Inc., provided a summary of qualitative factors that influence the availability of each alternative. We elicited availability estimates for the new reactor options using a process that was similar to that used for capacity estimates. Nine DOE and contractor staff members provided probability estimates that a tritium supply alternative could exceed 50, 65, 75, 85, and 95 percent availability. After discussion and revision, we averaged the probability estimates across individuals and fit a probability distribution (Figure 25.6).

PRODUCTION ASSURANCE SIMULATION. To determine the production behavior of each of the ten alternatives, we conducted a dynamic simulation using the following steps:

1. We sampled the start year of tritium production once from the schedule distribution.
2. We sampled the maximum capacity once from the capacity distribution.
3. We sampled the maximum availability once every 40 years after the start date.
4. From steps 1–3 we calculated the amount of tritium produced in any given year.
5. We added this amount to the amount remaining in the tritium inventory after 1 year of decay.

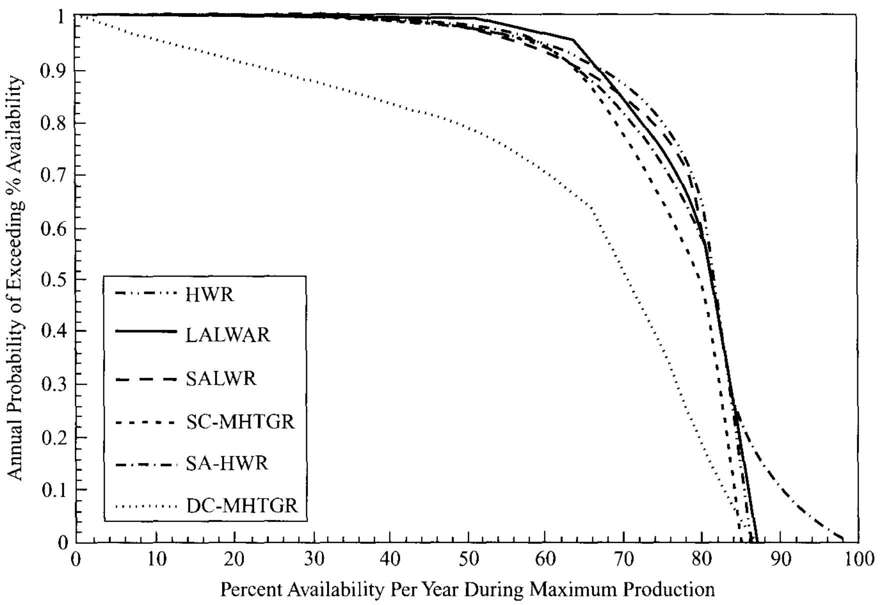

Figure 25.6. Results of the availability risk analysis.

Each iteration of steps 1 through 5 of the simulation provides an output that shows the tritium inventory over time. Repeating the process several thousand times provided a probability distribution over the available tritium inventory at any given point in time. This distribution can be compared with the required inventory levels. For example, the production simulations for the small ALWR shows that it is very likely that production starts on time and that shortfalls are rare (Figure 25.7). In contrast, most production simulations of the steam-cycle MHTGR start with a delay, creating shortfalls (Figure 25.8). However, once production starts, the required tritium inventory level is reached in a few years.

The production simulations showed that the MHTGRs and the HWRs had significant production problems because of start-up delays. All the other alternatives achieved START II production levels close to the required start date and easily maintained or exceeded this level throughout the 40 years of production. In addition, we ran several simulations that assumed a five-year shutdown of tritium production, for example, because of a major accident. Our purpose in these simulations was to determine how long it would take to make up for the five-year decay and return to the START II inventory level. These simulations showed that all alternatives were able to replenish the five-year decay in less than five years.

Cost Risk Analysis

Once the production assurance issues had been settled, cost became a major issue in this analysis (for a more detailed discussion of the cost analysis, see Dillon et al. 2002). First, the costs of all tritium supply alternatives were substantial, ranging from about $1 billion to $10 billion. Second, several alternatives could create

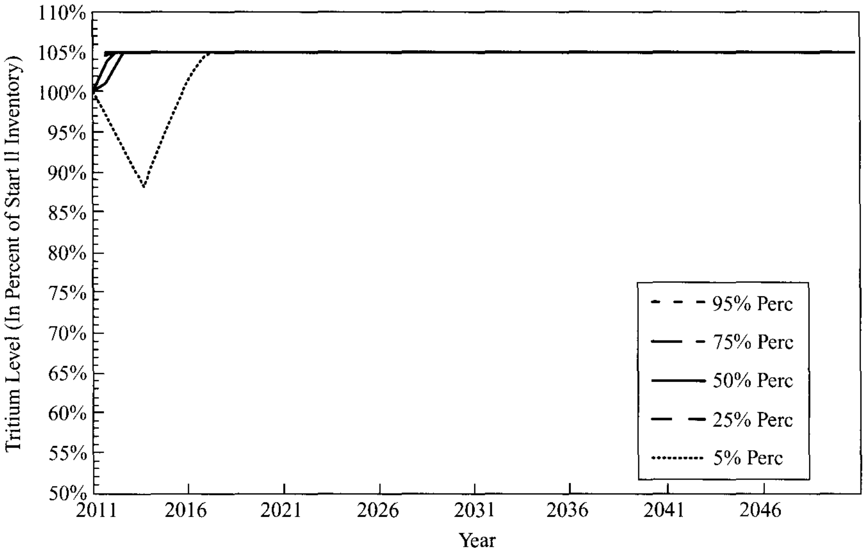

Figure 25.7. Production simulation for the small ALWR.

revenue through electricity generation, and the experts did not agree on how to account for revenues in the analysis. Third, there was substantial uncertainty in cost and revenue estimates, and some cost specialists estimated that the original engineering estimates could be exceeded several times over. Fourth, the discount rate became an issue, with arguments ranging from not discounting at all to using a discount rate as high as seven percent per year.

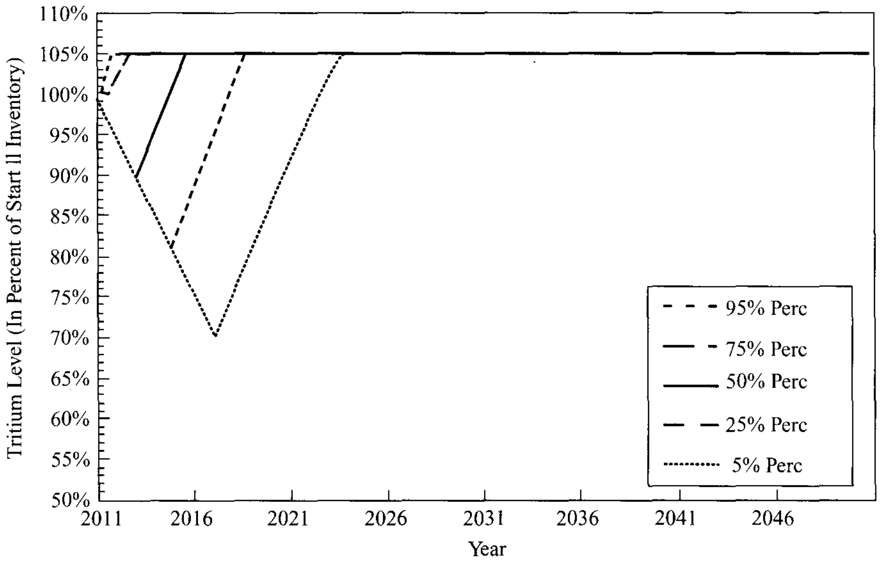

Figure 25.8. Production simulation for the steam-cycle MHTGR.

Table 25.4. The base-case cost estimates for ten tritium supply alternatives (costs were discounted at 4.9% annually)

Tritium supply alternative	**Total life-cycle cost (1995 $m)**
Large HWR	$4,354
Small advanced HWR	$2,703
Steam-cycle MHTGR	$4,113
Direct-cycle MHTGR	$2,364
Large advanced LWR	$1,678
Small advanced LWR	$1,212
Accelerator	$3,603
Purchase operating LWR	$30
Purchase partially complete LWR	$675
Purchase irradiation services	$959

Partly because of these issues, we developed the cost analysis in several steps. Fluor Daniel, Inc. (1995a) developed base-case costs using data from previous DOE and commercial studies. It divided the costs for each alternative into several cost categories (e.g., construction, operations and maintenance, fuel and electricity cost) and estimated base-case costs for each category. In addition, Fluor Daniel, Inc., estimated base-case revenue estimates and cost and revenue profiles over time. They used cost and revenue profiles to calculate the net present value in 1995 dollars using a discount rate of 4.9 based on guidance by the Office of Management and Budget. The base-case costs for the new facilities range from a low of $1,212 million for the small ALWR to a high of $4,354 million for the large HWR (Table 25.4). The commercial options are the least expensive, and, because of the revenues, the option to purchase an operating reactor is close to a financial break-even point ($30 million).

Recognizing the uncertainties in these cost estimates, DOE issued a cost study by Putnam, Hayes, and Bartlett (1995) that provided a range of high and low cost estimates for each alternative. These ranges covered in some cases $5–10 billion for a single alternative. However, Putnam, Hayes, and Bartlett did not provide any guidance about the likelihood of the costs within these ranges. Oak Ridge National Laboratory (1995) conducted an additional study to quantify the cost uncertainties, which provided base-case cost estimates plus a contingency cost that depended on the degree of experience with the technologies.

In general, the three cost studies were fairly consistent in that the base-case estimate was usually close to the low end of the cost range of Putnam, Hayes, and Bartlett's study. Oak Ridge National Laboratories' estimate fell in the lower third of the Putnam, Hayes and Bartlett study. However, there also were inconsistencies between studies, and it therefore became important to provide an integration of the results. In particular, DOE considered it useful to obtain probability distributions over the range of costs developed by Putnam, Hayes, and Bartlett.

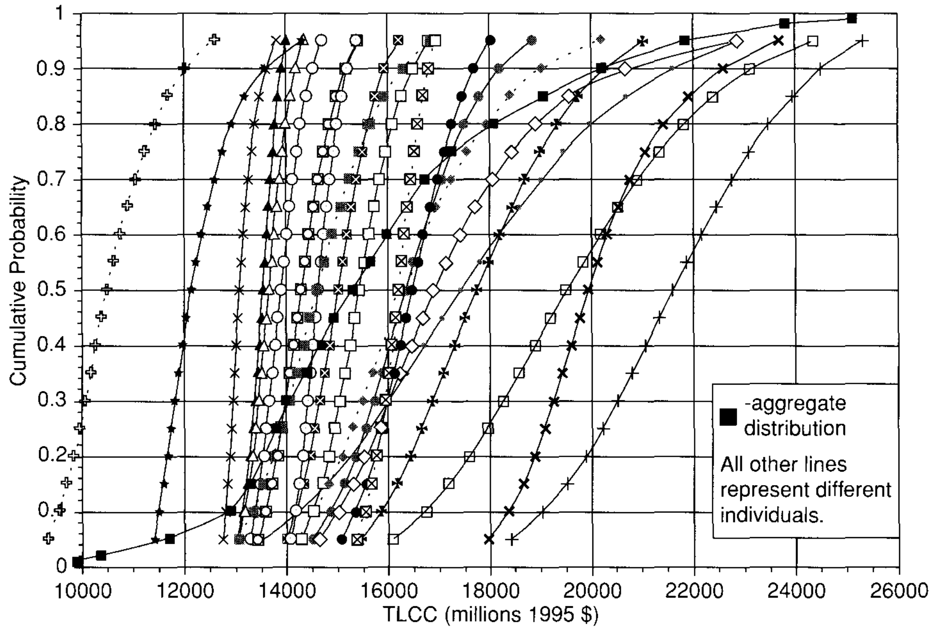

Figure 25.9. Individual and aggregate probability distributions for the total undiscounted life-cycle cost for the accelerator.

To resolve the differences in cost estimates and to assess these probability distributions, we conducted a major expert judgment exercise involving twenty-two cost and technical specialists from within and outside of DOE who met during two two-day meetings. During the first meeting, we introduced the specialists to the base-case cost estimates and presented them with the results of the previous studies on cost uncertainty. In addition, we trained the participants in probability assessment and they practiced cost probability assessment with two of the ten tritium supply alternatives.

Between the first and the second meeting, participants made independent cost estimates for each of the cost categories for each alternative. For each category, they provided a low (5th fractile), median (50th fractile) and high (95th fractile) of their probability distributions. They also provided comments or data sources and logic to justify their estimates and showed the calculations they used to generate estimates.

For major cost categories (construction, operation and maintenance, fuel cost, electricity cost) and for electricity revenues, they provided plausible minimum and maximum estimates in addition to the three fractiles.

Prior to the second meeting we calculated the overall cost distribution for each expert separately. The results of the individual elicitations are shown in Figure 25.9. These distributions of twenty-two experts are nothing short of astonishing. First, almost all experts show a great deal of overconfidence (indicated by very tight distributions). Second, the medians cover an extremely wide range from around $10 billion to $22 billion.

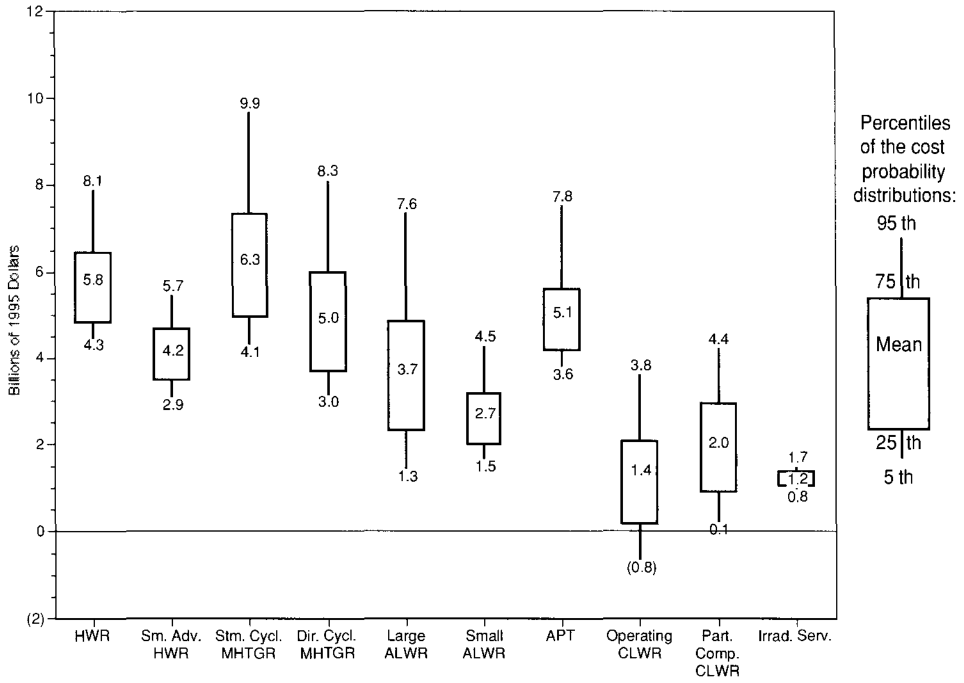

Figure 25.10. Box plots of the costs of ten tritium supply alternatives (with revenue) in billions of 1995 dollars, using a discount rate of 4.9 percent.

We started the second meeting by showing the experts the distributions in Figure 25.9. Each expert was given an envelope that included his or her symbol identifying his or her probability distributions without revealing it to other experts. Interestingly, the experts did not seem to be disturbed by the large variability of their judgments. In most cases, they attributed the variability to making different assumptions about cost growth phenomena. We discussed some of the more contentious estimates and subsequently encouraged participants to revise their estimates, but few did. We then calculated the revised cost distributions for each individual and averaged the distributions. The results are shown in the form of box plots in Figure 25.10.

Although the box plots show a large overlap in the cost estimates, there also are some important insights:

1. The commercial options have the lowest costs and the lower uncertainty. This is especially true for purchasing irradiation services.
2. Of the new construction options, the small ALWR has the lowest cost and lowest uncertainty. In fact, its cost distribution is dominated (higher cumulative probabilities for higher costs) by all other new construction alternatives.
3. The accelerator is one of the highest-cost options. Only the gas-cooled reactor options have higher costs.

Table 25.5. The results of the environmental analysis

Tritium supply alternative	Spent fuel per year (cubic yards)	Solid low-level waste per year (cubic yards)	Annual cancer risk from accidents
Large HWR	7	5,200	5.10E-05
Small advanced HWR	<7	<5,200	<5.10E-05
Steam-cycle MHTGR	80	1,300	1.00E-05
Direct-cycle MHTGR	82	~1,300	<1.00E-05
Large advanced LWR	55	710	2.60E-07
Small advanced LWR	36	660	2.30E-06
Accelerator	0	57	2.80E-11
Purchase operating LWR	40	160	no add. risk
Purchase partially complete LWR	~55	~710	~2.60E-07
Purchase irradiation services	<40	160	no add. risk

Environmental Analysis

In the initial round of analysis we evaluated seventeen environmental impacts, but we soon realized that only three impacts truly make a difference when choosing a tritium production technology: spent fuel production, low-level radioactive waste production, and risks from severe accident. The results are summarized in Table 25.5.

Spent fuel was measured by the cubic yards of radioactive spent fuel rods produced during reactor operations in one year. The new reactors generate spent fuel amounts ranging from 7 cubic yards to 80 cubic yards. The options to purchase an operating reactor or to purchase irradiation services would create up to 40 cubic yards of additional spent fuel (if only one reactor were used) because of shorter refueling cycles. If there were no change to the refueling cycles, no additional spent fuel would be generated. The option to purchase an incomplete reactor would create amounts of spent fuel comparable to those of the large ALWR. The APT does not generate any spent fuel.

For the new facility alternatives, the HWR creates by far the most low-level radioactive waste (5,200 cubic yards), followed by the other new reactors. The APT generates the least amount of low-level radioactive waste. The options to purchase an operating commercial reactor or to purchase irradiation services would create 160 cubic yards of additional low-level radioactive waste due to the use of additional fuel rods and due to handling additional radioactive materials. The option of purchasing an incomplete reactor would produce amounts of low-level radioactive wastes that are similar to those produced by the large ALWR, the small HWR, and the direct-cycle MHTGR. The low-level waste estimate for the APT applies to its helium target only and it is larger for the lithium target. The environmental impacts for purchasing a partially complete CLWR are similar to those of the large ALWR. The amount of additional spent fuel for the option to purchase irradiation services can be as high as 40 cubic yards per year, depending on the number of reactors used and their fuel cycle.

Cancer risks due to severe accidents can affect a population living within a 50-mile radius of a facility. For purposes of comparison, we used the DOE's Savannah River Site (SRS) in South Carolina because it has a relatively large population within 50 miles. The annual cancer risks from a severe accident of the new reactor technologies are very low, ranging from 1.0×10^{-5} to 2.6×10^{-7}. The APT has the lowest annual cancer risk (2.8×10^{-11}). The options to purchase an operating reactor or to purchase irradiation services would pose no significant additional severe accident risks because of adding tritium production. The option to purchase an incomplete commercial reactor would have severe accident risks that are comparable to those of a large ALWR.

In summary, the APT and the commercial options to purchase an operating reactor or to purchase irradiation services generate no additional spent fuel, have the lowest amounts of additional low-level radioactive waste, and have the lowest cancer risks from severe accidents. The new reactor alternatives and the completion of a partially complete commercial reactor produce spent fuel and low level radioactive waste, and they present a very small additional cancer risk from a severe accident.

The Decision

The analysis provided several key insights: First, the HWRs and the MHTGRs have significant problems with meeting the schedule and no cost or environmental advantages. The remaining reactor alternatives (ALWRs and commercial options) have similar production assurance and environmental impacts, but the commercial options are clearly less expensive. Although the commercial options look like a clear winner, they face several institutional issues. These include whether a U.S. utility would be willing to sell a reactor to the DOE and whether production in commercial reactors is compatible with international law and national policy. The accelerator provides a reasonable (though not very high) degree of production assurance and, although it is more expensive than the commercial options and some of the new facilities, it is also environmentally sound and it does not present unresolved institutional issues.

At this point, we faced two tasks: How to present these findings to the Secretary of Energy and how to use these findings in writing a formal "Record of Decision," once the Secretary had made her decision. After several briefings with close associates of Secretary O'Leary, we created the diagram shown in Figure 25.11, which is a qualitative summary of a consequence table representation. This diagram went through many discussions and wording changes, but in the end it conveyed the key messages of the analysis.

First, regarding the heavy water reactors and the gas-cooled reactors, it shows that they have schedule problems, which are not compensated by lower cost or environmental impacts. Therefore, the conclusion was not to go forward without these options.

Second, when looking at the choice between either building a new advanced light water reactor or using an existing one, the diagram illustrates a clear

Alternatives	Production Assurance	Cost	Environment	Recommended Path Forward
L-HWR	Schedule Problems	High	LL Wast	NO GO
SA-HWR	Schedule Problems	Medium	LL Wast	
SC-MHTGR	Schedule Problems	High	Spent Fuel	
DC-MHTGR	Schedule Problems	Medium	Spent Fuel	
L-ALWR	No Problems	Medium	Spent Fuel	COMMERCIAL
S-ALWR	No Problems	Medium	Spent Fuel	
O-CLWR	No Problems	Low	Spent Fuel	
PC-CLWR	No Problems	Low	Spent Fuel	
IS-CLWR	No Problems	Low	Spent Fuel	
APT	No Problems	Medium	No Waste	APT

Figure 25.11. High-level summary of the analysis and recommendations to the secretary of energy.

dominance relationship. All reactors have equally high production assurance and produce additional spent fuel, but the commercial options are much cheaper. Thus, when restricting the choice to these five options, it is clear that the commercial options are should be preferred.

Third, the accelerator has a clear advantage in terms on environmental issues. Although the diagram suggests that there are no production assurance problems, some DOE staff members and contractors thought that it had not as high a production assurance as the ALWRs or the commercial options. In addition, although the costs were labeled medium in the diagram, they really were among the highest of the ten options. Thus, the accelerator would make sense only if one places a very high weight on the environmental issues.

When these briefings occurred, two factions had developed within the DOE, one favoring the accelerator and one favoring the commercial option. Although the analysis informed the debate between these two factions, the ultimate decision appeared to be left to political considerations. The Secretary of Energy, faced with the debate about the pros and cons of the two options, decided to pursue both for some time. The diagram shown in Figure 25.12 represents this decision.

The idea of the dual-track decision was to pursue the commercial options to determine whether the institutional issues can be resolved; simultaneously, DOE would pursue the accelerator to make sure that it can meet production goals and time tables. If the DOE can resolve the institutional issues of the commercial options and if it eventually chooses one of the commercial reactor options, the cost savings to the government would be several billions of dollars. If the institutional problems cannot be overcome, the accelerator will provide an environmentally clean alternative.

The DOE published the results of the production assurance and cost analysis described in this article in a technical reference report for tritium supply and recycling (DOE 1995c). It published the environmental analysis, as required by the

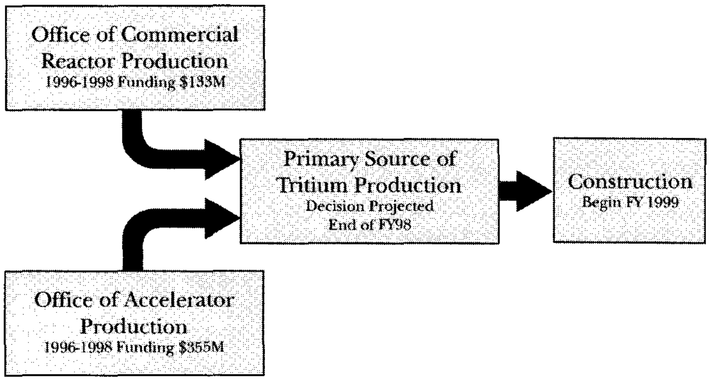

Figure 25.12. The dual-track implementation decision.

National Environmental Policy Act, in the programmatic environmental impact statement (DOE, 1995b). Meanwhile the deputy director of the tritium production office and I drafted the document which became the basis of the formal justification of the Secretary's decision. On December 5, 1995, the Secretary of Energy published the "Record of Decision: Tritium Supply and Recycling Programmatic Environmental Impact Statement" in the Federal Register (DOE 1995a). This decision confirmed and justified the dual-track strategy.

As a result of the Secretary's decision, the DOE discontinued the office of reconfiguration and created two new offices: The Office of Commercial Reactor Production (1996–1998 funding: $133 million) and the Office of Accelerator Production (1996–1998 funding: $355 million), which operated from the end of 1996 to the end of 1998.

The Aftermath and Some Thoughts about the Role of the Decision Analysis

Secretary O'Leary was followed by Secretary Pena to lead the Department of Energy. During Secretary Pena's tenure, Congress affirmed that the commercial option was an acceptable option for producing tritium. As a result, the major obstacle to this option was removed. Shortly after being appointed as Secretary of Energy in 1998, Secretary Richardson decided to cancel the accelerator program and to move ahead with the option to purchase irradiation services from a nuclear power plant. The designated facility for this purpose is the Hartsville nuclear plant owned and operated by the Tennessee Valley Authority.

There is no record that the decision by Secretary Richardson was informed by the decision analysis, though the analysis results are obviously consistent with this decision. Throughout the decision analysis, the justification of the accelerator as a viable option had caused some stress. Although clearly a viable option, it made sense only if DOE was willing to spend additional billions of dollars for what some perceived as relatively minor benefits.

This analysis, after many ups and downs, ended well. But it was not an easy process. The main problems that occurred during the eighteen months of analysis can be summarized as follows:

1. The first round of analysis involved a rather complex representation of many alternatives and many objectives and it did not get to the key issues.
2. Some key decision makers announced their preference for the accelerator early in the analysis.
3. The decision frame, especially the alternatives, was frequently shifting, with new alternatives entering and others being redesigned.
4. In the final stages of the process, decisions were made on a political basis and it is unclear to what extent the analysis informed the political discussions.

The first problem will be familiar to many decision analysts. Almost all analysts go through an early phase, in which the structure of the analysis gets overly complicated. It takes significant skill to include only the important features of the problem and build simple models that capture the essence of the problem, while not ignoring other important features. In hindsight, it is clear that the real issue of the tritium analysis was the choice of the production technology and that the key concerns were the production risks and costs associated with each alternative. Going through an elaborate multiattribute utility analysis was useful because it led to identifying these issues. However, we could have focused on the risk analysis much earlier than we did, had we not committed to a complete analysis of all alternatives early on.

The announcement by key decision makers of a preference for the accelerator came as a bit of a shock to me and the other decision analysts. Quite frankly, I seriously thought about resigning from this analysis because it seemed inappropriate to spend a significant amount of effort simply to justify a decision that may already have been made. However, I was reassured by several considerations and statements by the decision makers. First, the accelerator could indeed be a possible "winner," if a large weight was placed on environmental issues. Therefore, the accelerator would not be a clear "loser." Second, the advocates of the accelerator assured me that they would not interfere with the analysis process and that they would be open to any presentations of the "facts" about all options. Third, and perhaps most importantly, there were advocates of other options, especially of the small ALWR and of the commercial option, in the DOE. Ultimately, I could see how the analysis would inform the debate among these different factions.

The shifting decision frame was more of a nuisance than a serious problem. Decisions about which alternative to include and which to exclude sometimes happened in fairly casual conversations in the hallway. Often, these decisions had significant consequences for the analysis process, and, more importantly, for the subcontractors working on this project. It was useful to work in very close contact with the decision makers and to be able to influence some of these decisions to assure that the new frames were manageable.

It was somewhat disappointing that the analysis was not represented more vigorously in the final stages of decision making. Nevertheless, the eventual decision was consistent with the analysis. Clearly, the analysis showed that the heavy-water reactors and the gas-cooled reactors were nonstarters. It also favored the commercial option over the new light-water reactors. DOE would later be called to task for the rejection of the new light-water reactors by some of the vendors of these reactors, but the analysis was used to successfully defend the decision.

The only point of contention between the analysis and the decision to pursue the dual-track strategy was the wisdom of pursuing the accelerator and spending, as it turned out, another $355 million on that option. Accelerator advocates argued for its value by pointing to the potential for producing medical isotopes and other uses. They also continued to revise the cost estimates downward. But in the end, the investment in the accelerator seemed like a very large investment for relatively minor benefits.

Acknowledgment

This chapter was adapted from von Winterfeldt and Schweitzer (1998) with permission by the Institute for Operations Research and the Management Sciences (INFORMS).

REFERENCES

Department of Energy. (1995a). Record of decision: Tritium supply and recycling programmatic environmental impact statement. *Federal Register*, December 12, 1995.

Department of Energy. (1995b). Final Programmatic Environmental Impact Statement for Tritium Supply and Recycling, DOE/EIS-0161 .Washington, DC: Office of Reconfiguration, United States Department of Energy.

Department of Energy. (1995c). Technical Reference Report for Tritium Supply and Recycling, DOE/DP-0134. Washington, DC: Office of Reconfiguration, United States Department of Energy.

Dillon, R. Dillon, R., John, R., and Winterfeldt, D. (2002). Assessment of cost uncertainties for large technology projects: A methodology and an application. *Interfaces, 32*, 52–66.

Fluor Daniel, Inc. (1995a). Schedule Data Package for Tritium Supply and Recycling. Irvine, CA: Fluor Daniel, Inc.

Fluor Daniel, Inc. (1995b). Technical Data Package for Tritium Supply and Recycling. Irvine, CA: Fluor Daniel, Inc.

Fluor Daniel, Inc. (1995c). Cost Data Package for Tritium Supply and Recycling. Irvine, CA: Fluor Daniel, Inc.

Hora, S. C., and Iman, R. L. (1989). Expert opinion in risk analysis: The NUREG 1150 methodology. *Nuclear Science and Engineering, 102*, 323–331.

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

Keeney, R.L., and von Winterfeldt, D. (1991). Eliciting probabilities from experts in complex technical problems. *IEEE Transactions in Engineering Management, 38,* 191–201.

Oak Ridge National Laboratory. (1995). Cost Analysis of Multipurpose Reactor Options. Oak Ridge, TN: Oak Ridge National Laboratory.

Putnam, Hayes, and Bartlett, Inc. (1995). DOE Tritium Production Options: PHB Final Report on Cost Analysis. Washington, DC: Putnam, Hayes, and Bartlett.

von Winterfeldt, D., and Edwards, W. (1986). *Decision Analysis and Behavioral Research.* New York: Cambridge University Press.

von Winterfeldt, D., and Schweitzer, E. (1998). An assessment of tritium supply alternatives in support of the U.S. nuclear weapons stockpile. *Interfaces,* 28, 92–112.