May 2004

Allocating contractor risks in the Hanford waste cleanup

Jeffrey M. Keisler
University of Massachusetts Boston, jeff.keisler@umb.edu

William A. Buehring
Argonne National Laboratory

Peter D. McLaughlin
Pacific Northwest National Laboratory

Mark A. Robershotte
Pacific Northwest National Laboratory

Ronald G. Whitfield
Argonne National Laboratory

Interfaces
Vol. 34, No. 3, May–June 2004, pp. 180–190
ISSN 0092-2102 | EISSN 1526-551X | 04 | 3403 | 0180
DOI 10.1287/inte.1040.0078

Allocating Vendor Risks in the Hanford Waste Cleanup

J. M. Keisler
Department of Management Sciences and Information Systems, College of Management,
University of Massachusetts Boston, 100 Morrissey Boulevard, M/5-230, Boston, Massachusetts 02125,
jeff.keisler@umb.edu

W. A. Buehring
Decision and Information Sciences Division, Argonne National Laboratory, 9700 South Cass Avenue,
Building 900, Argonne, Illinois 60439-4832, wabuehring@anl.gov

P. D. McLaughlin, M. A. Robershotte
Pacific Northwest National Laboratory, Richland, Washington 99352
{peter.mclaughlin@pnl.gov, mark.robershotte@pnl.gov}

R. G. Whitfield
Decision and Information Sciences Division, Argonne National Laboratory, 9700 South Cass Avenue,
Building 900, Argonne, Illinois 60439-4832, rgwhitfield@anl.gov

Organizations may view outsourcing as a way to manage risk. We developed a decision-analytic approach to determine which risks the buyer can share or shift to vendors and which ones it should bear. We found that allocating risks incorrectly could increase costs dramatically. Between 1995 and 1998, we used this approach to develop the request for proposals (RFP) for the US Department of Energy's (DOE's) privatization initiative for the Hanford tank waste remediation system (TWRS). In the model, we used an assessment protocol to predict how vendors would react to proposed risk allocations in terms of their actions and their pricing. We considered the impact of allocating each major risk to potential vendors, to the DOE, or to both and identified the risk allocation that would minimize the DOE's total cost—its direct payments to vendors plus the costs of any residual risks it accepted. Allocating inappropriate risks to the vendor would have increased costs because the vendor would add a large risk premium to its bids, while allocating inappropriate risks to the DOE also would have increased costs because the vendor would not take adequate risk-reduction measures. With the improved risk allocation, the RFPs resulted in bids that were acceptable to the DOE.

Key words: decision analysis: risks; government: agencies.
History: This paper was refereed.

In the mid 1990s, the US Department of Energy (DOE) was planning the cleanup of its nuclear waste tanks at the Hanford site in Washington State. It faced well-documented problems, because of the volume, the makeup, and the poor characterization of the waste held in the underground storage tanks. The DOE's task was complicated by many legal obligations and constraints, and by the political sensitivity of the problem arising from a history of delays, disagreements, and disclosures at the site. Various stakeholders were concerned about spending, environmental issues, and radioactivity. The anticipated cost was in the tens of billions of dollars.

To minimize costs and risks, in 1994 the DOE undertook what it called a privatization initiative under which contractors (vendors) would have as much decision-making authority and responsibility as possible. The DOE, in partnership with the Pacific Northwest National Laboratory (PNNL), established a large privatization team near the Hanford site to support this initiative. Its first major task was to contract with one or more vendors for the development and demonstration phase of the cleanup. The DOE had to draft and issue a request for proposals (RFP), evaluate and select bids, and negotiate contract terms.

Comments from potential vendors on a 1995 draft RFP (TWRS Privatization 1995) showed that risk allocation was a major concern: the privatization initiative depended on the contractor's being able to obtain financing, and the financial market probably would not accept certain risks. Other risks were potentially acceptable but would require a high risk premium. The potential vendors were under no obligation at this point to tell the DOE how high the risk premium really had to be, however, and the DOE wanted to avoid being locked into bad contract terms.

From late 1995 to early 1998, the PNNL contracted with coauthors (Keisler, Buehring, and Whitfield) all at that time from Argonne National Laboratory's (ANL) Decision and Information Sciences Division to support the risk-management group (headed by coauthors McLaughlin and Robershotte) of its privatization team. Thus, the PNNL was the direct client for this work, using it to write portions of the RFP that the DOE ultimately issued (TWRS Privatization 1996). ANL proposed the basic structure for this work and conducted most of the assessments and analysis, while the PNNL advised on the approach, participated in the assessments, managed interactions with the privatization team, and integrated the results with the rest of the privatization plan. We developed a decision-analytic approach to help the DOE determine how it should allocate risks within the RFP, that is, to what extent should the DOE accept risks and to what extent should it pass risks on to the vendor?

We defined alternative allocations and assessed the risks resulting from them to quantify their cost implications, while accounting for the different incentives and risk-bearing capacity of the DOE and its potential vendors, leading to a successful RFP and further efforts to institutionalize our approach.

The problem of how to allocate risks is prevalent in all sorts of contracting situations, and using RFPs that allocate risks poorly can cause inefficient outcomes. Our conceptual approach applies in many of those situations.

Conceptual Approach

Allocating a risk to the party better able to control it motivates that party to find the best way to manage that risk. Each party has expectations about how the other will behave given the incentives implied by the risk allocation, and those expectations ought to inform contractual negotiations. It is also reasonable to think that the larger party (the government in this case) would be better able to deal with large variances in financial outcomes than would smaller parties (vendors). This implies that, other things remaining equal, the contractors would demand a larger premium to bear risks than would the government. To allocate risks efficiently, the decision makers must take note of all these factors and balance them.

Viewing this as a decision-analysis problem, with the DOE as the decision maker, we wanted to choose the risk allocation that would maximize the DOE's expected utility. It faced environmental uncertainties and uncertainties about how vendors would respond to the proposal. Inspired by game theory, or at least the asymmetric prescriptive-descriptive approach Raiffa (1982) outlined, we built a prescriptive model for the DOE's decisions and a descriptive model for the vendor's decisions. From the DOE's point of view, the problem could be viewed as a stylized decision tree with nodes representing its decisions and nodes representing uncertainties (Figure 1).

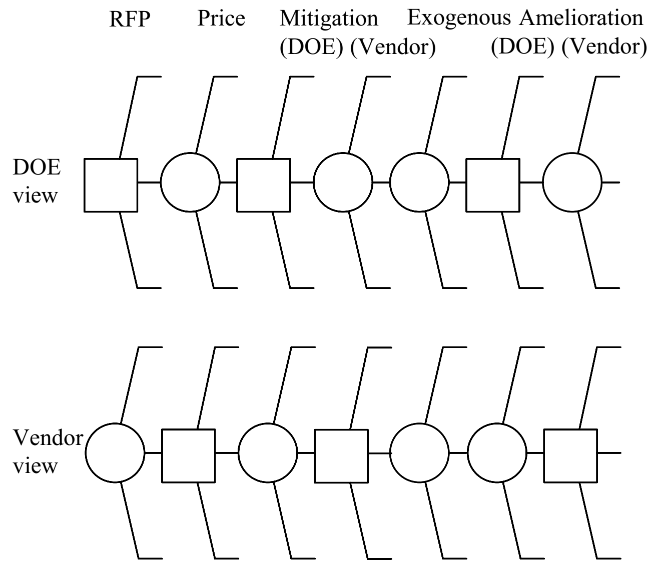

Figure 1: Stylized decision trees were used to represent the cash flows resulting from a contract. For the DOE, the first node is the DOE's decision about how to allocate risks in the RFP. The second node is the potential vendor's decision (an uncertainty from the DOE's perspective) about what price to bid (to be predicted using a vendor-response model). The third node represents the DOE's decisions taken to mitigate risks. The fourth node consists of the vendor's decisions taken to mitigate risks (independent of the previous node). Following this are uncertain exogenous events, after which the DOE acts to ameliorate the outcomes for the events that have occurred, after which the vendor does the same (independent of the previous node). The DOE's end-point values are determined by the contractual payments it makes to the vendor along with the costs incurred due to the risks it bears and the events that occur. For the vendor, the first node is an uncertainty about how the DOE decided to issue the RFP. The second node is the vendor's decision about what price to bid. The third node is a decision node about its steps to mitigate risks. The fourth node is an uncertainty about the DOE's steps to mitigate risks (independent of the vendor's choices about the same), and the fifth node consists of uncertain exogenous events. The fifth node consists of the vendor's decisions to ameliorate negative outcomes, and the sixth node is the DOE's actions (about which the vendor is uncertain) to ameliorate risks (again independent of the previous node).

From the vendor's point of view, the DOE has already chosen its risk allocation. We simplified the vendor's decision tree and assumed that the vendor's goal is to maximize its chance of getting the contract, that is, to bid as low a price as possible, subject to the constraint that it still be possible to finance the project. That is, it passes on the risks it faces to its financiers and it is these financiers' risk attitudes that we modeled.

The financiers consist of equity and debt providers. Equity providers consider the predicted variance of cash flows and demand an expected return on equity based on that risk. Debt providers look at the likelihood that they will be repaid without problems and demand not only that an interest rate be based on that measure of risk but also that the project have a percentage of equity investment that depends on the level of risk. The prediction of cash flows (that will be used to make debt payments and the amount of residual cash flow left to give a return on equity) is based on the vendor's anticipated fixed and variable costs, the contracted price for waste treated, and the possible consequences of risks as allocated in the RFP.

Assessment of Individual Risks

We started with a list of 100 risks. We combined some of the risks and narrowed the list substantially. The final list contained risks for which the potential impact was sizable, the range of possible allocations was substantial, and the right allocation was not obvious. We examined nine major risks, some requiring separate assessment of subsidiary risks:

(1) Interest rate (interest rates may change between the time of the bid and financing),

(2) Inflation (the vendor faces exposure to inflation risk in the draft RFP),

(3) Change in law (applicable laws and regulations may change),

(4) Permitting (the vendor may not be able to obtain necessary permits),

(5) Waste envelope C (probably the most technically difficult waste to process),

(6) Appropriation (Congress may not appropriate funds adequate for planned activities),

(7) Decontamination and decommissioning (D and D may be more costly than expected),

(8) Other uncontrollable circumstances (lawsuit, sabotage, earthquake, tornado), and

(9) Not-to-exceed (NTE) price risk (the contract pricing arrangement could be a literal price limit, as in the draft RFP, or could be a target price that allows selected adjustments).

We also assessed a background risk level for the DOE and the vendor, that is, the residual risk each would bear even if all nine risks resolved to their base cases.

We characterized each risk through discussions with the PNNL privatization staff experts. We used a common form to organize the assessments. For each risk, we defined three alternative allocations: the risk was to be borne primarily by the DOE, to be shared, or to be borne primarily by the vendor. We constructed these allocations using both qualitative and quantitative definitions to match reasonable alternative phrasings the DOE might include in the RFP. This meant, for example, that even for risks allocated primarily to the vendor, the DOE could still bear some tangible risk. Similarly, shared risk does not necessitate a completely symmetric risk-bearing arrangement. The specific vendor was not yet known, so we did our assessments for a representative potential vendor. We intended the definitions to be precise enough to translate without ambiguity to specific wording within the RFP. When the DOE later used the model to evaluate specific bids, we conducted different assessments for different vendors.

We then defined the possible outcomes. After obtaining verbal descriptions of the mitigation steps each party would likely take under each potential risk allocation, we assessed probabilities of each outcome (implicitly conditioned on the likely mitigating steps taken by both parties). For this step, we tried to take the vendor's view as well as the DOE's view. We assumed that vendors would assign the same probabilities to each outcome as would the DOE, unless there were special reasons not to (for example, the DOE might be far more trustworthy regarding a certain risk than the vendor believed, or the vendor might be more concerned about different risks and events than would the DOE). We also obtained verbal descriptions of the steps each party would be likely to take to ameliorate the negative outcomes of risks under each allocation.

Finally, we obtained verbal descriptions of what the consequences would be for each side for each outcome of the risk, assuming both sides had taken the predicted steps. From this information, we assessed point estimates or probability distributions over the financial consequences for each side. The consequences were defined in terms of the net present value (NPV) of cash flows computed at the current risk-free rate, and expressed as increments from a base case in which the project proceeds without difficulty. In some cases, we noted that outcomes of one risk would determine the relevance of other risks; for example, if the project stopped early, later risks would not affect cash flows.

The following summary of one such assessment, namely the assessment for D and D risk, illustrates the process of defining the alternate allocations and the risk-management scenarios and the resulting parameters that we used in the model (Table 1). It took us approximately one half day to assess this risk. Keisler and Buehring (1996) give the details of all the risk assessments.

Results of Interviews to Characterize One of the Risks (D and D)

Risk Name: Decontamination and decommissioning (D and D).

Description: D and D may be more costly than expected, particularly if the vendor is not provided with an incentive, through allocation of risk to the vendor, to keep these costs low.

Consequences: The consequence is given in terms of the probability distribution over D and D costs, assumed to occur 10 years after production starts.

Risk allocation	DOE takes risk	Shared risk	Vendor takes risk
Description	DOE is responsible for D and D.	Vendor sets aside limited funds, which may be refunded if D and D costs are less than expected.	Vendor is responsible for D and D; DOE pays excess if vendor reaches the hypothetical limit of its ability to pay.
Probability distribution for cost to DOE	$5 M : 10% −$50 M : 40% −$150 M : 50%	$2 M:30% −$40 M : 60% −$140 M : 10%	0 : 75% −$50 M : 25%
Probability distribution for cost to vendor	No cost	$3 M : 30% −$10 M : 70%	$5 M : 50% −$50 M : 25% −$100 M : 25%

Table 1: We summarized the risk-assessment interviews in the form of quantitative descriptions of the risks used as inputs to the simulation model. The description for the D and D risk shown here is typical.

Low cost is approximately $5 million less net D and D cost than expected. Medium cost is approximately $50 million more net D and D cost than expected. High cost is approximately $150 million more net D and D cost than expected.

Risk Allocation (Representative Characteristics)
The DOE bears the risk: The DOE states up front that it is responsible, and the vendor allocates no money for D and D.

Shared risk: Responsibility is defined, and the vendor establishes a medium-sized sinking fund and gets money back if D and D is under budget; the DOE covers overage.

The vendor bears the risk: The vendor pays for D and D and must establish a large sinking fund from the start; it gets the unused money back.

Mitigation and Prevention Strategies
The DOE bears the risk: The DOE establishes strict standards the vendor must follow in planning and operation and have close cooperation with environmental regulators.

Shared risk: The DOE monitors the vendor somewhat in planning; the vendor is responsible for operation with checkpoints. The DOE, the vendor, and the regulators together coordinate some of these activities.

The vendor bears the risk: The DOE intervenes only in extreme cases; all the vendor's actions take into account their impact on future D and D costs, and the vendor works hard to establish rapport with regulators.

Likelihood of Occurrence
The DOE bears the risk: Low cost = 10 percent, medium cost = 40 percent, high cost = 50 percent.

Shared risk: Low cost = 30 percent, medium cost = 60 percent, high cost = 10 percent.

The vendor bears the risk: Low cost = 25 percent, medium cost = 50 percent, high cost = 25 percent.

Costs are lowest when the risk is shared and highest when the DOE bears the risk, because both the DOE and the vendor can influence outcomes and have the incentive to do so under a shared-risk scenario.

Modeling

We defined four main risk-allocation strategies for further analysis (Table 2). Each strategy consisted of specific allocations for each of the nine risks. All of these strategies were within the realm of possibility. Prior to engaging us to work on the problem, the DOE was heading toward something close to the vendor-bears strategy (Strategy 3). This strategy is quite similar to the 1995 draft RFP, and even after we started work, a senior DOE official declared that the DOE would "shift all risks to the vendor." Potential vendors in turn had indicated that just about every risk should be borne by the DOE (Strategy 1), although, predictably, the DOE was skeptical of such a position. As we analyzed individual risks, the analysis (and other factors) influenced the RFP. By the time we finished analyzing the various strategies, Strategy 2 had

Strategy	1	2	3
Name	DOE	Shared	Vendor
Risk name			
Interest rate	DOE	Vendor	Vendor
Inflation	DOE	DOE	Vendor
Changes in law	DOE	DOE	Vendor
Permitting	DOE	Shared	Vendor
Waste stream C	DOE	Vendor	Vendor
Appropriation	DOE	Vendor	Vendor
D and D/RCRA	DOE	Shared	Vendor
Lawsuit	DOE	Shared	Vendor
Sabotage	DOE	DOE	Vendor
Earthquake	DOE	DOE	Vendor
Tornado	DOE	DOE	Vendor
Not to exceed	DOE	Shared	Vendor

Table 2: For each risk-allocation strategy, we defined how each risk would be allocated. Sharing risks did not mean that each risk is shared: some risks are shared, but other risks could still be allocated entirely to one party or the other.

Strategy	1 ($)	2 ($)	3 ($)
Vendor mean cost increment	157	32,660	106,152
Vendor standard deviation	55,876	61,780	94,655
DOE mean cost increment	336,731	107,625	60,400
DOE standard deviation	133,409	82,517	70,675

Table 3: We generated key summary statistics from the risk simulations (figures in thousands). These results described the impact of risk sharing, but we still needed to incorporate the main cost, direct payments from the DOE to the vendor.

support at some levels but this support was neither complete nor final. At management's request we also included Strategy 4 (not shown), which differed only slightly from Strategy 2 and under which the DOE fared slightly worse.

We used a Monte Carlo simulation with 1,000 iterations to produce profiles for each risk. For each iteration, we summed the increments from the base case due to each risk. The result, after accounting for risk interactions, was a frequency distribution over the total increments (in which higher costs in the chart correspond to reduced profits) to the base case, including payments for both the vendor and the DOE, for each of the strategies considered (Table 3).

As a measure of the reliability of the simulation results under repeated simulations, we used the standard deviation of the average cost for a given strategy, which was approximately $2 million.

Vendor Response

We then had to predict the price the vendor would charge under each strategy, which we could combine with the DOE's risk profile to obtain a net-cost distribution for the DOE.

Based on extensive conversations with the privatization financial task leader for the TWRS and with external financial consultants, we developed equations to predict how the vendor would set a price in response to any given set of risks and risk-allocation decisions. We validated these equations using experts' predictions regarding specific changes to the risk profile for which subjective estimates were fairly easy to provide. We then had other experts confirm that the parameter values we used were realistic. Our purpose was to codify judgments rather than to describe exactly and in detail what any particular firm would do.

One important variable used in the financial-risk model was the probability of default, for which we used as a proxy the probability that the vendor would face an after-tax loss larger than its equity stake. Specifically, we calculated the probability that $h - i < -(j + k)$, where

h is the baseline NPV of vendor cash flows discounted at the risk-free rate,

i is the net impact on vendor cash flows from all risks, also in terms of dollars discounted at the risk-free rate (we approximated the probability distribution on i by fitting a normal distribution to the simulation results),

j is the amount of equity the vendor puts up, and

k is the relative tax benefit of reduced profits in unfavorable cases compared to the baseline.

We determined cash flows by entering the proposed financing terms and price into a financial model (Weimar and Paananen 1995) based on standard accounting practice and developed by the privatization finance task manager to predict the DOE's and the vendor's costs in a deterministic base case and to clarify the literal implications of contract terms (Figure 2).

We assumed that the vendor would raise prices in response to being forced to take on additional risk in the following ways: (1) Risks that are predominantly downside would decrease the expected return; the vendor would raise the price to bring the expected return back to its baseline value. (2) Risks that increase the probability of default make it harder to obtain debt financing, so we assume the interest rate for debt financing corresponds to the degree of risk and the percentage of equity required, both of which we assume would increase when downside risk increases. (3) Equity providers are concerned with predictability on both the upside and the downside, so they would raise the required return on equity when the overall risk level increases.

The risk profiles and the vendor-response models determine the vendor's required internal rate of return (IRR); we must link this IRR to the main financial model to determine a target price for the vendor and the corresponding distribution over cash flows. These cash flows in turn determine debt terms (fraction of debt and debt interest rate) and costs, from which we calculate the probability of default. The equations below implicitly define an upward sloping curve with capital costs on the y-axis and probability of default on the x-axis. For given financing terms, there is also an implicit curve for price versus probability of default, which is downward sloping because raising the price gives the vendor a cushion that lowers the probability of default. We then use a modified binary search algorithm to identify the minimum price at which the vendor could still finance the project, subject to the following constraints: that the IRR is no less than that given by Equation (1), that the debt fraction is no greater than that given by Equations (2) and (3), and that the debt interest rate is no less than that given by Equation (4). The solution lies where the curve representing price versus probability of default intersects the curve representing the financier's capital cost versus probability of default.

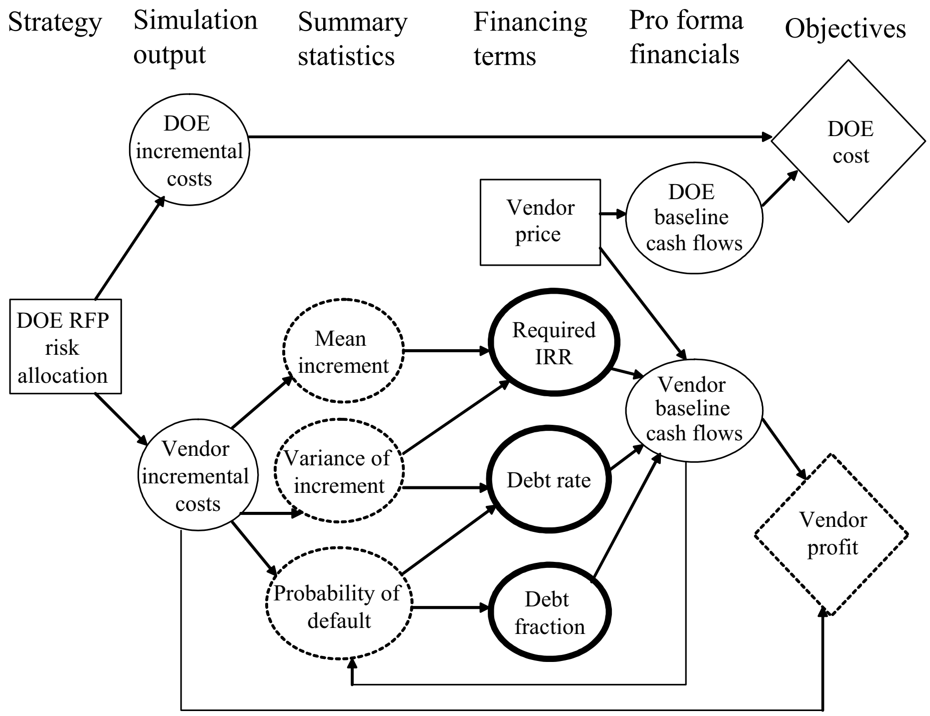

Figure 2: This stylized influence diagram represents the financial relationships in the model. The model predicts price from risk allocation and other assumptions. The influence diagram has two features that make it unconventional—the circular relationship involving price and probability of default and the fact that mean, variance, and probability of default are summary statistics derived from the probability distribution over the vendor cost increment (derived from end points of the decision tree in Figure 1).

In other words, the solution lies where the project's pro forma NPV including debt costs is exactly 0 when discounted at the predicted IRR requirement.

Equations Used to Represent Vendor Financial Response to Programmatic Risks

After obtaining advice concerning a few key considerations, we postulated the following relationships. The financial experts based their estimates of parameter values on their knowledge of other debt offerings, including situations in which debt ratings (which have a known relation to default rates) and corresponding terms were available.

The required after-tax internal rate of return (IRR) is given by

$$\mathrm{IRR} = k_1 + k_2 \times A + k_3 \times B, \quad \text{where} \tag{1}$$

$k_1 = 0.03$ (the risk-free rate),

$k_2 = 0.0275$ (the increase in IRR required to accommodate \$10 million of average downside), and

$k_3 = 0.025$ (the sensitivity of interest rate to variance in NPV of cash flows).

In this equation and in the ones that follow,

A denotes the variance in the vendor cost distribution (expressed in dollars$^2 \times 10^9$), and

B denotes the vendor mean cost increment (in dollars $\times 10^5$) from the simulation results.

If the probability of default (P) is less than five percent, the maximum debt fraction (DF) is given by

$$\mathrm{DF} = k_4 - k_5 \times P - 0.03 \times (A/k_6)^{k10}, \quad \text{where} \tag{2}$$

$k_4 = 0.95$ (the maximum debt fraction for a hypothetical case with zero percent probability of default),

$k_5 = 1.5$ (the sensitivity of the debt fraction to the probability of default),

$k_6 = 3.23$ (the variance corresponding to the minimum vendor risk case, including the assumed level of background risk), and

$k_{10} = 3.0$ (the exponential sensitivity of the debt fraction to variance). We used the last part of the equation in particular to calibrate predicted financing terms for experts' judgments. We were aware that it lacked a further theoretical basis, but given our time constraints, we used it, albeit cautiously. In later applications, we implemented this element of the model more elegantly.

If the probability of default is greater than or equal to five percent, the debt fraction is given by

$$\mathrm{DF} = k_4 - k_5 \times P - 0.03 \times (A/k_6)^{k10} - k_7 \times (P - T), \quad \text{where} \tag{3}$$

$k_7 = 16$ (marginal sensitivity of debt fraction to default rate for risky projects), and

$T = 0.05$ (the probability of default corresponding to financiers' threshold of great concern). This threshold is analogous to a downgrading of the debt on the project, which we assessed subjectively as the point at which the project would change from business as usual into what financiers would view as risky.

The minimum required debt interest rate (DR) is given by

$$\mathrm{DR} = M + k_8 \times P/Y + k_9 \times A/k_6, \quad \text{where} \qquad (4)$$

$M = 0.0737$ (the T-bond yield at the time of the study),

$Y = 0.07$, considered to be near the maximum allowable probability of default,

$P =$ the probability of default (computed by the model),

$k_8 = 0.138$ (the coefficient of the sensitivity of the debt interest rate to the probability of default), and

$k_9 = 0.015$ (the sensitivity of the debt interest rate to the variance in cash flows).

Where possible, the financial advisors identified analogous financing packages in the market; for example, bonds with a given rating have a historical annual default rate and a known risk premium over T-bonds. By finding a bond rating corresponding to $Y = 0.07$ and assuming a linear relationship between yield and default rate, one can estimate k_8. We derived some parameters directly from such market data, estimated some using subjective judgment informed by such data, and estimated others, such as k_{10}, by obtaining expert judgments about what financing terms ought to be for a small number of specified artificial scenarios and fitting the values to mimic those expert judgments. We compared the implications of the model, for a wide range of cases, with the intuitions and expectations of the financial experts and of the privatization team management. We concluded that the financing terms and the corresponding price for the baseline case were reasonable. The model's predicted impact of risk allocations on financing terms is of the same order as the financial advisors' estimates of those impacts for several particular reference risks that were fairly easy to judge. Our experts thought these parameters would provide reasonable results over the range of strategies we examined and, in particular, for considering the incremental impact of changes to the risk-allocation strategy.

Numerical Results

By optimizing our model, we obtained vendor financial terms (Table 4). By adding the DOE's payment to the vendor to the DOE's simulation summary statistics, we obtained the financial measures of direct interest to the DOE (Table 4).

The distributions over vendor cost (Figure 3) and the DOE's net costs for the entire contract period (Figure 4) for each strategy show some expected characteristics, such as the low risk to the vendor and the wider spread for the DOE when the DOE bears the risks. The vendor's risk distribution showed much greater potential for loss than for gain, which explained why, in the initial discussions of the draft RFP, vendors indicated they would require higher payments than the DOE had expected. Most risks were to be assigned to one party or the other; however, the middle case (RFP) actually reduced the DOE's total downside risk. Some of the distributions had spikes near \$0, corresponding to cases in which the vendor would be made whole (that is, the DOE would reimburse the vendor for all expenses including financing costs) after the project ended.

Strategy	1	2	3
Vendor financial measures			
Probability of default (P)	3.4%	5.17%	8.23%
Required internal rate of return (IRR)	11.6%	14.3%	30.3%
Debt fraction (%) (DF)	86.6%	82.7%	30.6%
Debt interest rate (DR)	9.3%	9.8%	11.6%
Price paid to the vendor by DOE			
DOE payment to vendor (\$K)	\$602,255	\$636,530	\$1,397,921
DOE financial measures			
DOE mean cost increment	\$336,731	\$107,625	\$60,400
DOE standard deviation	\$133,409	\$82,517	\$70,675
DOE total cost	\$938,986	\$744,155	\$1,458,321

Table 4: Key financial measures of interest to the vendor were tracked for each strategy. When the vendor bears all risks (Strategy 3), the difficulty of financing the project results in the DOE having to make much higher payments to the vendor. Although the DOE could minimize its direct payments to the vendor by accepting all risks, this would usually lead to higher total costs.

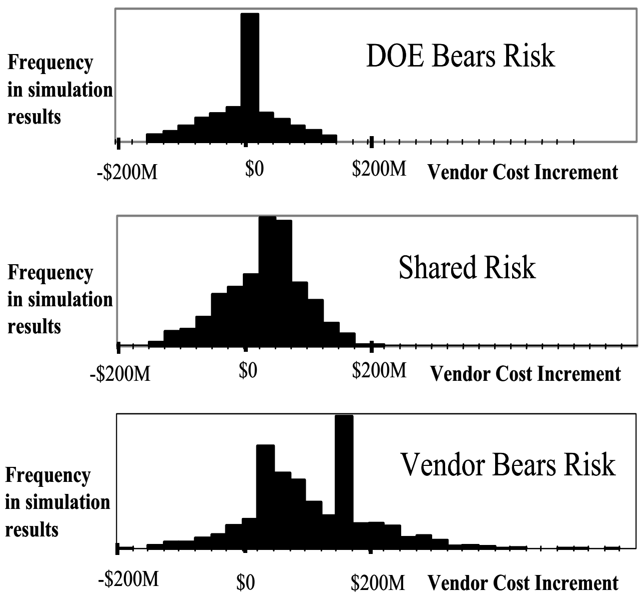

Figure 3: The vendor's distribution of cost increments (the deviation from contracted payments in the base case) depends on the risk allocation. Under the shared-risk strategy, the vendor has somewhat greater spread over its potential costs, but its mean costs do not increase much.

In analyzing the strategies and the results for each, we demonstrated that, indeed, sharing the risks would lead to savings. Allocating all the risks to the DOE would mean the vendor would not do enough to keep costs low, while allocating all the risks to the vendor would cause the vendor to demand too high a risk premium; thus, the DOE would obtain the lowest costs with Strategy 2. By bearing all risk, the DOE would obtain the lowest price, because the vendor would have a lower mean cost increment and variance and thus the best financing terms. By sharing risks, the DOE would make payments to the vendor that would be six percent ($34 million) higher than those under Strategy 1, but the savings to the DOE from impacts of residual risks would lead to a net cost savings for the DOE of over 20 percent ($194 million). Compared to the strategy in which the vendor bears all risks (Strategy 3), Strategy 2 requires the DOE to pay costs due to residual risks that are 78 percent higher ($47 million) but the total cost savings to the DOE are nearly 50 percent ($714 million) because it would pay a smaller risk premium to the vendor. The price for Strategy 3 would make the project infeasible. In spite of the cost difference, the standard deviation of the DOE costs under Strategy 2 was only slightly higher than under Strategy 3. Strategy 3, the costliest approach, was the DOE's preferred strategy going into this effort, but the figure given, $714 million, perhaps overstates the potential savings accruing under Strategy 2: if bids came in that high, the DOE could have issued, with difficulty, a revised RFP.

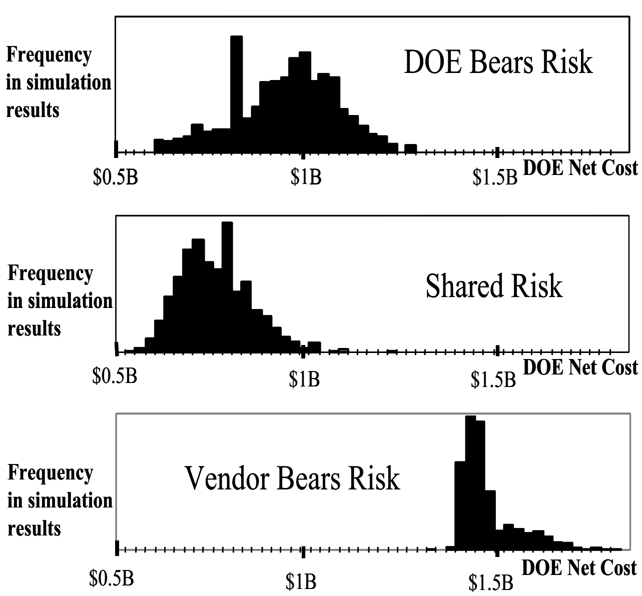

Figure 4: The histograms we generated for each risk-allocation strategy show that sharing risk stochastically dominates the two extreme risk allocations and so reduces both the DOE's expected cost and its risk.

Our presentation of these results triggered an animated discussion. At that time, the DOE managers essentially switched their support from an RFP that allocated most of the risks to the vendor to one that shared risks. We then refined the proposed RFP by performing a one-way sensitivity analysis for allocation of each risk, using Strategy 2 as a basis. We identified the allocation for each risk that minimized the DOE's cost when other risks were left unchanged. From these allocations we constructed Strategy 5, which resembled Strategy 2, except that the DOE and the vendor shared the interest rate, changes-in-law, waste-stream-C, and appropriation risks. In particular, by completely allocating to the vendors risks that they believed were under the DOE's control, for example, appropriation, D and D, and permitting, the DOE would increase its cost: vendors would assume the worst and incorporate such risks into their price.

For comparison, we added Strategy 6, in which all risks are shared, as well as many other variations (not shown) that we considered at the DOE's request. When the DOE and the vendor share all risks, the mean cost increment to the DOE is higher than in Strategy 2, as is the payment to the vendor, and so Strategy 6 is $24 million worse than Strategy 2. Strategy 5 had a mean cost increment to the DOE of $5 million more than Strategy 2, in exchange for a $15 million reduction in payments to the vendor, for an additional savings of $10 million along with a slight reduction in the standard deviation of $2 million.

By changing the allocation of the other risks, the DOE would obtain relatively minor potential cost increases or cost savings. This implied that if the vendors seemed very concerned about risks for which the model predicted little increase in costs (for example, if the vendor demanded an additional $50 million to accept waste-stream-C risk), then the DOE could make concessions on those risks (for example, accept the risk rather than pay the premium). Such risks might have symbolic value for the vendor or the vendor might have a more pessimistic view of the situation than it merited. Conversely, if the vendors were more willing to accept a risk than the model predicted, the DOE could try to obtain concessions from the vendor to reduce its total cost. The entire portfolio of risks matters and must be considered as a unit. For example, although the impact of shifting the NTE risk entirely to the vendor from Strategy 2 (in which it is shared) is $41 million, the cost of shifting the NTE risk to the vendor when the vendor already bears all other risks would be much larger ($198 million). In other words, the vendors can bear a little risk comfortably, but beyond a certain threshold, vendors will not assume additional risk without penalizing the DOE dramatically.

The final RFP was very close to our recommendation of Strategy 5. It had some refinements based in part on further application of our model. The DOE received several bids and selected two vendors for the technology-development stage. Although the cost included financing charges that the DOE perceived as high, the bids came in lower than the DOE's previous cost estimates and lower than bids that it had expected prior to the reallocation of risk in the RFP.

Decision-Support System

When the DOE issued the RFP, we anticipated it would need to evaluate bidders' suggestions and requests quickly during the bid-selection and negotiation phases. We improved some parts of the model to speed the DOE's response and to take advantage of insights gained along the way. We then created a decision-support system (DSS) rather than a one-off model. The resulting system answered negotiators' what-if questions in about half an hour, instead of the half a day or more the original model could take, depending on the complexity of the question. To make some of the improvements, we relied on standard methods to make models faster, more flexible, and more user-friendly. In particular, we restructured the model to use modular components and automated most tasks. We also made conceptual improvements that made the model more efficient than it was in our earlier efforts; we recommend such improvements as a starting point for similar efforts others might undertake.

We used neater risk templates to make it easier to structure and assess large models. We developed four basic templates to characterize risks that would be common to many procurement contracts, specifically permitting, processing, early termination, and ongoing financial risks. For each stage of the contracted project, users must complete one type of risk template, possibly for several risks of that type. The DOE would have to successfully resolve risks from each stage to move on to the next stage. This fact led to the next feature. We modeled risk interactions using a precedence matrix, in which a cell entry of 1 or 0 indicates whether the DOE must successfully resolve the risk in the corresponding column before going on to encounter the risk in the corresponding row. This matrix made it easy to add individual risks without affecting the rest of the simulation and allowed us to represent interactions in a more compact and comprehensive way.

We refined the vendor-response function and made it more transparent based on further meetings with the project's external advisors. The simulation (and templates) in the DSS specifically noted potential default events and, separately, more severe write-off events in which lenders would lose principal. In the original model, financier concerns about write-off events were reflected in the exponential penalty function we used for situations in which the probability of default was above five percent. Considering write-off risks explicitly is more precise, and by doing so we replaced the exponential penalty for high default probability with linear functions of the default and write-off probabilities. The DSS records future write-off events and default events as parts of the outcomes of individual risks as it simulates them. It does not directly calculate the debt fraction but instead derives it from a debt-service-coverage-ratio requirement found in practice. This structure facilitated assessments (with which everyone was more comfortable) of the judgmental model relating project summary statistics to financing terms. We ended by streamlining and automating the optimization subroutine for finding the vendor's minimum price.

Organizational Impact

The insights from our analysis produced consensus within the privatization team and within the DOE. The DOE managers presented the conceptual summary diagram (Department of Energy Privatization Working Group 1998, §5, Figure 4.6) on risk sharing to the then secretary of energy, Hazel O'Leary. One purpose of our effort was to provide the DOE with evidence about the potential consequences of insisting that the vendors alone should bear the risk, and we successfully showed that doing this would have a high cost. The DOE included our results in its report to the US Congress on the privatization project (Department of Energy Privatization Working Group 1998, §3.4), which described the model's impact:

> The DOE also concluded that the level of uncertainty with respect to design, financing, and regulation at the end of Part A was such that fixing prices would require an excessive price to compensate for the risk faced by the vendor. Thus, a design phase (referred to as Part B-1 in the contract) was defined to reduce this uncertainty and to provide the DOE with various reviews and decision points prior to proceeding with construction and operations. The design phase will allow time to verify technology performance on Hanford-specific wastes and to optimize debt and equity arrangements and technical requirements.

The flexibility that the design phase added proved important. Our new approach to risk allocation kept privatization viable up to that point. Our work convinced the DOE to take a more businesslike attitude toward risk. The DOE's report to Congress (Department of Energy Privatization Working Group 1998, §3.2) included the following lessons learned:

> *Need for an equitable risk allocation*. In the early stages of developing the Phase I contracting approach, the DOE recognized that privatization is effective in shifting significant performance risk to the vendor, but some risks would have to remain with the DOE. For example, the DOE recognized that the private sector would not accept the risk of potential fluctuations in yearly budget appropriations. In addition, the DOE recognized the need to absorb the risks associated with its own performance in areas such as waste characterization and preexisting conditions.

The DOE privatization working group (Department of Energy Privatization Working Group 1998, Case study 6) prepared a report on privatization for the secretary of energy. Its list of lessons learned from this project included the following:

> *Risk Allocation.* This is the single most important feature of the "deal" to be established. There must be an equitable allocation of risk between the DOE and the vendor. The vendor must be held responsible for the risks that it can control—for example, technical and performance risks. The DOE needs to be responsible for risks that it can control—for example, government-furnished items and minimum waste quantities.

The overall report, which helped lead to establishment of the DOE's Office of Privatization and Contract Reform, contained the following recommendation (Department of Energy Privatization Working Group 1998, §2, recommendation 5):

> *The DOE should better integrate the perspectives of the business community into its privatization initiatives.* One of the keys to successful privatization at the Department is the ability to structure business deals that will attract the business community by, for example, ensuring a proper balance between risk and reward.

This represented a complete turnaround from the view that the vendors should carry all risks, which had led to discouraging comments in response to the draft RFP.

Afterword

Through 1998, the DOE used the model often on an ad hoc basis in finalizating the RFP and selecting vendors and negotiations with the remaining vendor (of two initially selected) going into what was called Phase 1B-1 (the first stage was Phase 1A). In this case, the DOE used the model to evaluate whether the vendor's proposed risk allocation was reasonable and to negotiate an agreement (McLaughlin et al. 1998). Over the rest of the demonstration phase up to the start of full-scale production, the DSS forecast the optimal strategy to have expected costs of \$3 billion to \$5 billion less than strategies allocating almost all risk to either the DOE or the vendor.

The larger privatization effort had a disappointing end, however, for reasons that were beyond the scope of the RFP risk-allocation decisions. About a year after the last agreement, after developing much of the technology, the vendor suddenly raised its price estimates for the next stage by more than 100 percent. The reasons behind this move were controversial (Welch 2000). The risk of increased costs had been allocated so that when costs became higher than the DOE considered acceptable, the vendor gave up the technology it had developed and lost out on all incentive payments, while the DOE paid for the work completed and incurred holding costs while seeking a replacement. The DOE used its option to abandon its contract with the vendor (and with it, the privatization initiative) at the end of the design phase and then adopted a more conventional government-owned, contractor-operated approach.

Summary

Our new approach led to the DOE's successful issuance of an RFP that at one time had large expected savings in the hundreds of millions of dollars or more. In this case, we cannot trace good decisions about risk allocation to good ultimate outcomes, but this approach allowed the privatization program the good intermediate outcome of continuing; it attracted two vendors for Phase IA, when zero response would have been likely given the DOE's original attitudes toward assigning risk. In retrospect, the DOE may have made too much of a stretch in attempting to switch to a privatization regime under which it would have to behave as just another business.

That should not deter others—certainly businesses involved in procurement and probably governmental agencies as well—from considering this approach. Government agencies should consider the Hanford cleanup a rather extreme case. In many simpler situations, this approach could help a bureaucracy to overcome its reluctance to share risks. The DOE's documented lessons learned apply in such cases. Companies doing business with government agencies may also find our approach useful as a way to help the agency understand the company's positions in negotiations that include the allocation of risks.

Acknowledgments
We thank Mark Weimar for his advice, his enthusiastic cooperation, and his comments on this paper. We thank Bill Richmond and Megan Lerchen, subject matter experts who allowed us to conduct interviews and improve our understanding of complex programmatic risk topics. We thank Carol Sohn for her interest and support and other members of the privatization team, including Paul Kearns and Jerry Straalsund, for their assistance. We thank Raymond A. DiPrinzio and Brian Oakley of Scully Capital Services for their invaluable assistance in the development and refinement of the contractor response model. We also thank Mark J. Jusko of Argonne National Laboratory for his assistance with the refinement of the model as well as Chris Klaus for his assistance running scenarios. We thank two anonymous reviewers for their helpful suggestions. Finally, we thank Detlof von Winterfeldt for reviewing the initial proposal for this effort. The original work on which this paper is based was funded by the US Department of Energy, contract number DE-AC06-76RL O1830.

References

Department of Energy. *Report to Congress: Treatment and Immobilization of Hanford Radioactive Tank Waste*. Retrieved February 27, 2002 www.hanford.gov/docs/twrs-atp/toc.html (www.hanford.gov/docs/twrs-atp/section3.html).

Department of Energy Privatization Working Group. 1998. Report to the Secretary of Energy. *Harnessing the Market: The Opportunities and Challenges of Privatization*. Retrieved February 27, 2002 www.osti.gov/privatization/report/, DOE/S-0120. (www.osti.gov/privatization/report/chapt-2.htm).

Keisler, J. M., W. A. Buehring. 1996. Summary of the technical report on estimating the impact of key programmatic risk allocation decisions on Phase 1 bids and U.S. Department of Energy costs. ANL/DIS/TM-41, Argonne National Laboratory, Argonne, IL.

McLaughlin, P. D., M. A. Robershotte, W. A. Buehring, R. G. Whitfield. 1998. Analysis of the risk allocation for TWRS privatization phase IB. PNNL-12028 (limited distribution). Prepared for US Department of Energy under contract DE-AC06-76RLO 1830, November.

Raiffa, H. 1982. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, MA.

TWRS Privatization. 1995. *TWRS Privatization Request for Proposals* (draft). DE-RP06-96RL13308, November.

TWRS Privatization. 1996. *TWRS Privatization Request for Proposals.* DE-RP06-96RL13308, February.

Weimar, M., J. Paananen. 1995. Financial market analysis of Phase 1 for TWRS privatization. Report on file in TWRS Privatization Backup Documentation Library, June 9.

Welch, C. 2000. Skyrocketing costs may delay nuclear cleanup at Hanford. *Seattle Times*, May 1. Retrieved February 27, 2002 seattletimes.nwsource.com/news/local/html98/hanf01m_20000501.html.

George B. Mellinger, Senior Program Manager, Pacific Northwest National Laboratory, 902 Battelle Boulevard, PO Box 999, Richland, Washington 99352, writes: "I was The Deputy Manager of the Waste Disposal Integration Team (WIT) at the Hanford, Washington Department of Energy (DOE) site from 1995–1998. The WIT was the Pacific Northwest National Laboratory team (also referred to as the Privatization Team) responsible for assisting the DOE with developing, implementing, and managing the contracting strategy for the cleanup of massive amounts of nuclear waste stored in underground storage tanks. Under the risk and decision management sub-task for the WIT, extensive work was done in risk allocation modeling to support this multi-billion dollar DOE Tank Waste Retrieval System (TWRS) Privatization effort.

"I verify that the authors of paper 54-702-JK (Keisler, Buehring, McLaughlin, Robershotte, and Whitfield) did in fact do the work described for the TWRS Privatization team in the roles they detailed, and that their efforts did have the impact they describe. The concept of proper and appropriate risk allocation was key in garnering private industry support for this huge and complex cleanup effort. Specifically, the Request for Proposal (RFP) was modified based in part on their recommendations. The DOE received bids in response to the modified RFP, which enabled the initiative to proceed at that time."