IEEE Transactions on Power Systems, Vol. PWRS-1, No. 1, February 1986

FUEL CONTRACTING UNDER UNCERTAINTY

Richard B. Fancher
Decision Focus, Inc.
Los Altos, California

James F. Wilson
Decision Focus, Inc.
Los Altos, California

Howard A. Mueller
Electric Power Research Institute
Palo Alto, California

Abstract - Uncertainty about fuel requirements and fuel market conditions has made long-term utility fuel planning risky. A methodology is presented that explicitly incorporates uncertainty into the analysis of the fuel "contract mix" problem. The approach allows utility analysts to optimize contract mix in the face of uncertainty, to identify and quantify risks and to measure the value of fuel procurement flexibility.

INTRODUCTION

Over the past decade, utilities have found it increasingly difficult to make sound fuel planning and procurement decisions. The uncertainty in fuel burn requirements and fuel market prices causes much of this difficulty.

Until the late 1960s, fuel markets were relatively orderly: primary fuels were reliably available, and fuel price changes were gradual and smooth. Since the late 1960s, fuel markets have been marked by substantial cost increases and substantial volatility. For example, prior to the late 1960's, fuel generally comprised some 20% of utility costs. In the United States they now account for almost 50% of costs or more than $40 billion annually.

Recently traditional fuel procurement practices have often resulted in expensive outcomes for both utility ratepayers and shareholders. Inflexibility in the mix of contracts has led to surplus inventories, unfavorable contract termination costs, and contract or "captive mine" costs significantly above market prices.

This paper presents a quantitatively explicit methodology to help understand, improve and explain utility fuel procurement strategies. The methodology is an application of the decision analysis framework described in [1] and [4]. By explicitly incorporating uncertainty, the approach allows utility analysts to optimize fuel mix in the face of uncertainty, to identify risks and to measure the value of fuel procurement flexibility.

Following the Problem Statement, the methodology and solution approach are briefly described. The largest section of the paper consists of examples. These demonstrate how uncertainty is represented and the kinds of information and insights available under this approach.

PROBLEM STATEMENT

The utility fuel procurement decision involves selecting a set of quantities from the available supply alternatives--both long-term contract and spot. Fuel contracts represent a utility's commitment to purchase and a supplier's commitment to supply a stream of a specified quantity of fuel over time at a price which is unambiguously defined at the time of delivery. Contract procurement usually involves a two step process. The first step involves the decision to make a long-term contract, and the quantity to commit to; the second step involves deciding what quantity to have delivered during a specific time period. Delivering unneeded fuel or curtailing deliveries under contract may result in extra costs to the utility.

The fuel planner strives to minimize cost while taking into account reliability, fuel quality and shareholder welfare. Often these other concerns are not explicitly quantified but are included as constraints. The traditional fuel procurement analytical approach consists of the deterministic application of rules-of-thumb based on experience, such as "an 80% contract fraction." As in other areas of utility planning, uncertainty was dealt with through the use of scenarios, if at all. The result has been the familiar set of problems: high fuel inventories due to fuel contract commitments larger than required to meet burn, contract terminations, contract fuel prices well above market, and problems with captive mine costs.

Incorporating Uncertainty

Explicitly including uncertainty in a quantitatively structured manner helps improve fuel procurement analysis. Two areas of uncertainty complicate the problem--the uncertainty in fuel requirements (demand) and the uncertainty in future fuel prices and availability (supply).

The methodology uses probability distributions to characterize uncertainty. For example, Figure 1 illustrates burn uncertainty. The top half presents the probability density; the bottom half presents the equivalent information as a complementary cumulative probability distribution. The cumulative is analogous to an equivalent load duration curve and will facilitate our presentation of subsequent examples.

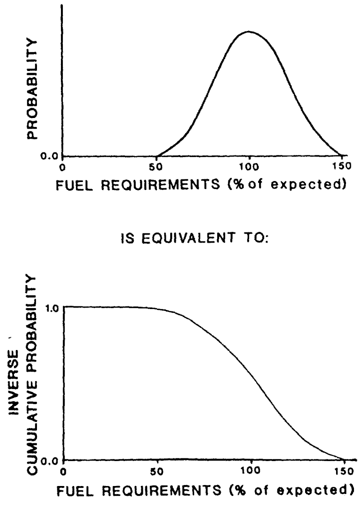

Figure 1. Equivalent Descriptions of Fuel Requirements Uncertainty

Requirements for a specific fuel can depend on the utility electric demand, supplies available from lower priced resources and the availability of the generating equipment being fueled. The fuel burn of a base-loaded generator depends only on its availability; however, all the enumerated factors

85 SM 466-8 A paper recommended and approved by the IEEE Power System Engineering Committee of the IEEE Power Engineering Society for presentation at the IEEE/PES 1985 Summer Meeting , Vancouver, B. C. Canada, July 14-19, 1985. Manuscript submitted September 18, 1984; made available for printing June 6, 1985.

0885-8950/86/0001-0026$01.00©1986 IEEE

affect the fuel demands for other generating resources. Thus burn uncertainty often results from the combination of several uncertainty sources.

As one approaches a specific period the uncertainty concerning that period's demand usually diminishes. For example, we may know that a planned lower-cost resource is commercial or that it has been delayed. This characteristic--that uncertainty partially resolves--is an essential component to the methodology's incorporation of uncertainty within fuel procurement analysis. Figure 2 illustrates the concept.

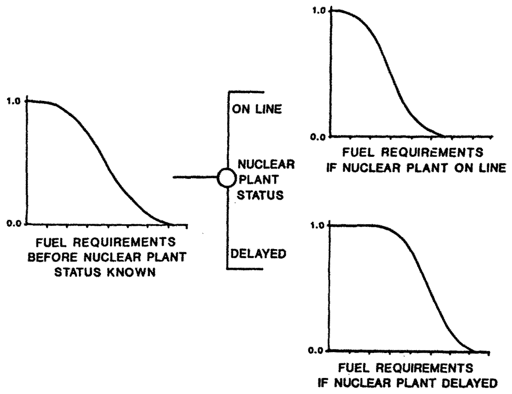

Figure 2. Illustration of Resolution of Fuel Requirements Uncertainty

Fuel Procurement Options

Utilities supply their fuel demands with short-term (spot), medium-term and long-term contracts. Figure 3 presents a spot market supply curve. The figure indicates that larger volumes of spot purchases will entail higher incremental costs, illustrating a common perception of the spot market's price characteristic. Contracts are characterized by a set of factors: nominal quantity flow rate, base price, escalation, duration, leadtime, and the cost incurred for taking other than the nominal quantity. Figure 4 illustrates a common structure of the last item. In this example the utility is allowed to adjust deliveries between 80 percent and 100 percent of nominal flow without cost. We term this the "flexibility range." The figure also illustrates that deliveries of less than 80 percent incur a charge for each unit of fuel refused. We term this charge "underlift;" in Figure 4 it is $10 per ton.

On the supply side of the fuel procurement, considerable uncertainty often exists in the price and quantity characteristics of the spot market and the escalation outcomes for contracts. Again, these uncertainties often partially resolve over time.

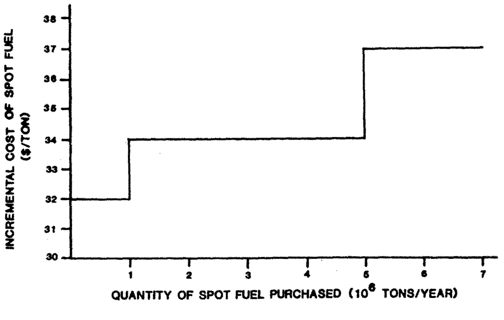

Figure 3. Example of Spot Market Price-Quantity Forecast

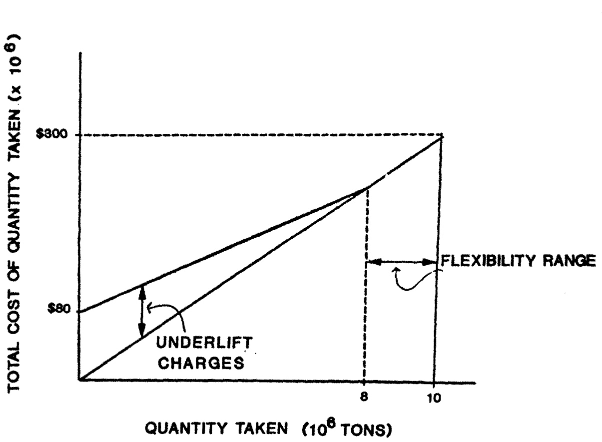

Figure 4. Example of Contract Structure (Underlift Charge = $10/ton)

When explicitly considering uncertainty, to select a set of quantities from available supply alternatives requires that the range of possible outcomes and corrective contingent decisions be analyzed in an integrated fashion. The resulting best set of decisions differs significantly from the process of optimizing for a single assumed set of future circumstances. Differences arise because a single assumed future does not give credit for flexibility nor consider contingent decision alternatives. The analysis containing uncertainty explicitly makes tradeoffs between price, contract flexibility and the availability of alternatives if expected forecasts are not realized.

By explicitly tracing out the range of possible outcomes, the fuel analyst can also consider the risks posed to the ratepayer or potentially the shareholder. This information will have significant value to utility decision makers if they perceive that their regulatory review process creates shareholder incentives which do not exactly correspond with minimizing expected ratepayer costs.

MODEL DESCRIPTION

The methodology has been incorporated into a computer model that represents a simulation of fuel planning and fuel procurement decision-making. The model may be used to evaluate the costs and risks associated with contracting strategies in specific settings. The structure of the model can be understood as including a short-term, tactical fuel management model within a longer term framework of uncertain events and strategic decisions. The general model structure is diagrammed in Figure 5. An analogy may be drawn to capacity planning models. Such models often include a short-term, tactical model of system operation within a strategic model representing the construction and retirement of generation resources.

The entire model evaluates a contract mix strategy. The probability distribution of total discounted fuel cost for a given set of strategic choices is determined. This is done by simulatng the possible ways that fuel demand and market conditions may evolve (hereafter the possible "states of the world"), simulating the fuel planner's strategic choices in each state, and executing the tactical fuel management model to determine the cost outcomes. Much of the model structure represents the timing of uncertainty resolution and strategic decisions, corresponding to a decision tree or decision diagram as described in [1].

Tactical Model

The tactical model is a representation of short-term fuel procurement management. The decisions involve scheduling

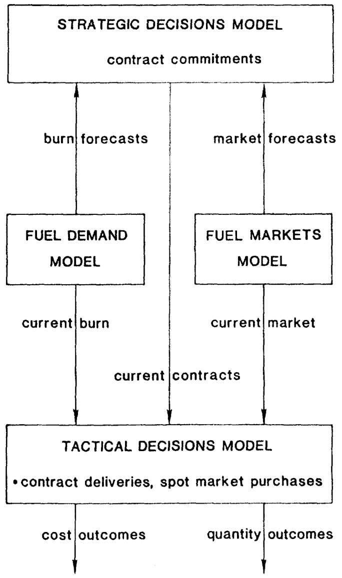

Figure 5. Contract Mix Model Structure

deliveries under contracts already in place and purchasing fuel on the spot market to meet current burn requirements. There are no uncertainties in the tactical model except the short-term burn variation. Market conditions and contract commitments are given.

The tactical model is highly simplified. The spot and contract procurement opportunities are ranked in order of increasing incremental cost, resulting in a short-term supply curve for fuel. For a level of fuel burn, the least-cost mix of available contract fuel and spot fuel is selected. The result of the tactical model is a probability distribution on fuel cost, where the range of outcomes is due to the short-term burn uncertainty. Thus the tactical model determines the fuel manager's cost outcomes based on a set of contracts already in place.

Figure 6 shows the possible contract and spot quantity outcomes of the tactical model for a simple example. The diagram indicates a situation where a contract for 85% of expected burn has been arranged. Assuming fuel under the contract is cheaper than spot, the mix of spot and contract fuel selected to meet each possible level of burn is seen by looking horizontally across the diagram. For instance, when the burn equals the expected burn, the entire contract quantity is delivered and spot fuel is purchased to meet the additional 15% of burn. At the minimum level of burn, the quantity underlifted equals 35% of expected burn; the maximum spot exposure occurs at the maximum level of burn, when spot purchases equal 65% of expected burn.

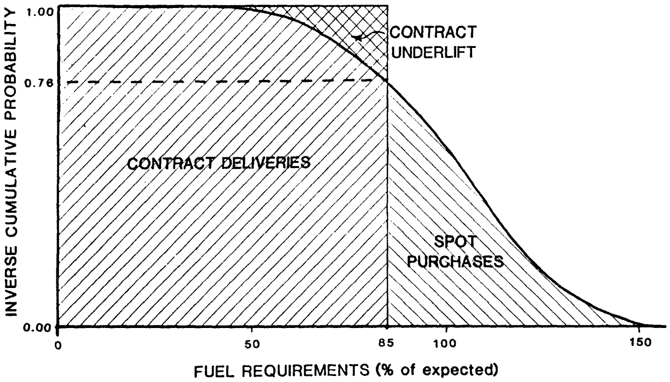

Figure 6. Illustration of Contract and Spot Quantity Outcomes with Uncertain Burn

The hatched regions of Figure 6 have important interpretations--they represent the expected quantities of contract deliveries, contract underlift, and spot market purchases. The point at which the burn distribution intersects the contract quantity line indicates the probability that spot purchases will be necessary, about 0.76 in the example.

The fundamental tradeoff between the risk of contract underlift or inventory holding costs and the risk of high-cost spot purchases is clearly seen through this diagram. If the contract quantity is increased (the line representing the contract quantity moves to the right), the expected quantity of spot purchases decreases but the expected amount of contract underlift increases.

Again there is a useful analogy between the burn distribution and a load duration curve. Drawing the contract quantity on the burn cumulative is analogous to the stacking of generation resources under the load duration curve. The contract may be thought of as the "base-loaded" option (low variable cost, inflexible), with the spot market representing a more expensive and flexible "peaking" option.

Strategic Decisions

The long-term, strategic decisions include the quantities of contracts or strategic investments to commit to. (This could include, for instance, purchasing a mine or building a refinery.) In addition to current alternatives, a model application may involve various types of contracts available at future times. The fuel planner is assumed to take advantage of any information that becomes available in time for a decision. Therefore future strategic decisions may be contingent on how uncertainty has resolved up to the time of the decision. A future option requires one decision for each possible state of the world at the time.

The strategic options may involve lead times or durations of commitment. As there is uncertainty about future burn and market conditions, there is uncertainty about how valuable an option will be in the future. The current attractiveness of a contract offer depends on the possible cost outcomes in the many possible future states of the world.

Optimization of Strategic Decisions

An additional feature of the model is an optimization procedure that determines the strategic choices that minimize expected fuel costs. The optimization procedure invokes the model iteratively, incrementally improving the strategic choices until little further improvement is possible. The minimum attainable expected discounted fuel cost, and the strategic choices that attain that cost level, are determined.

The optimization method used is nonlinear, based on the method of steepest descent [3]. The derivatives of fuel cost with respect to each strategic decision quantity are calculated within the tactical model and are used in the optimization.

The derivative of cost with respect to a contract quantity indicates the attractiveness of an additional incremental quantity of that contract, given all other strategic choices. The attractiveness of a contract will depend on the possible outcomes of fuel burn and market conditions for the duration of the contract. The derivatives therefore are sums of probability-weighted values computed (in the tactical model) under various states of the world.

METHODOLOGY EXAMPLES

When uncertainty is explicitly considered in a model, the range of the kinds of information and comparisons that result is greatly expanded. The following examples, based on very simple models, are intended to demonstrate the implications of incorporating uncertainty into the analysis of contract mix decisions.

The examples are based on a group of coal plants with an expected burn of 5 million tons per year (Mtpy). The burn is considered to be uncertain with a standard deviation of 1 Mtpy. For simplicity, only uncertainty about fuel burn is represented. The spot market options are as represented in Figure 3. The contract option used for the first examples has no flexibility range, a price of $30/ton, and an underlift charge of $6/ton. The strategic decision in the first examples will be the quantity of the contract option to commit to.

Uncertainty and Contract Mix Decision-Making

Suppose a fuel planner in our example could select a contract quantity after finding out his level of burn. He could then always contract for his exact burn level (contract is cheaper than spot), and pay $30/ton for all fuel purchased. His expected total cost would be $30 times the expected burn (5 Mtpy) or $150.0M.

Figure 7 shows how the expected value of fuel cost in our example depends on a contract decision made under burn uncertainty. As we would anticipate, expected costs are higher than the $150.0M achievable under certainty. There is a contract quantity (90% of expected burn in this example) that minimizes expected cost. Higher contract fractions involve excessive underlift charges, and lower contract fractions involve excessive spot market purchases. The 90% fraction trades off underlift and spot in the optimal proportions. This minimum point depends heavily on the assumptions made in the example: the degree of burn uncertainty, the structure of the contract, and the price-quantity relationship in the spot market.

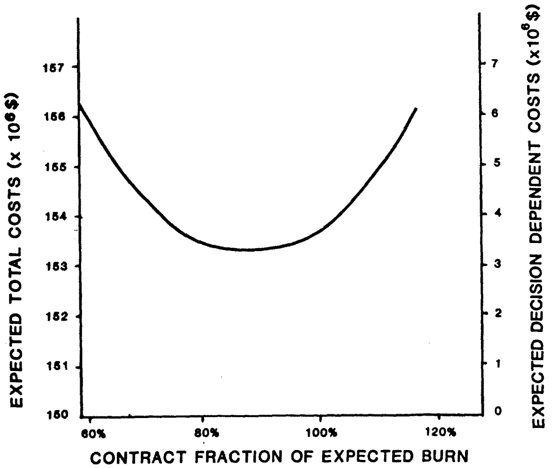

Figure 7. Expected Value of Cost

We see that even if a 90% contract fraction is selected, expected costs are $153.1M, $3.1M higher than if burn information is available in time for the decision. It could be said, then, that the cost of uncertainty in this example is $3.1M. The fuel planner should be willing to pay up to this amount to obtain perfect information about what his burn will be before committing to a contract quantity. In decision analysis this value is called the Value of Perfect Information on fuel burn [2].

The fuel planner cannot avoid the $150M cost even if perfect information is available. We will therefore call costs above this "perfect information" level the decision-dependent costs. The expected value of decision-dependent costs is displayed on the right-hand scale in Figure 7.

Figure 8 shows how burn uncertainty translates into cost uncertainty for three representative contract fractions. Cost uncertainty is lower the higher the contract fraction. The two tics on the horizontal axis show the range of expected values of the three distributions. Clearly the variation in cost outcomes is large compared to the differences in expected cost. It can be concluded, then, that risk may be an important consideration in situations such as the one in this example.

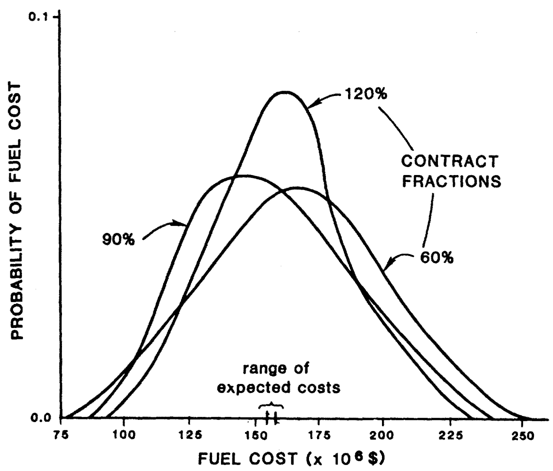

Figure 8. Probability of Cost Outcomes

Risk of High Cost Outcomes

When uncertainty is considered, the possibilities of extreme fuel burn outcomes or of extreme market conditions are represented. The expected value of fuel cost includes these extreme cost outcomes, weighted by their probabilities. If the fuel planner chooses to minimize the expected value of fuel cost (rather than the value of fuel cost he most expects), the risk of extreme cost outcomes is considered.

Risk information may be of interest for several reasons. Planning for contingencies requires an awareness of the possibilities. Additional options, flexibility or information may be sought to reduce exposure to risks. A utility may be called upon to defend a decision, which would involve demonstrating an awareness of the risks involved with the actions selected and other alternatives.

Risk information will also be needed when certain outcomes are to be especially avoided. In this case the objective is not simply to minimize expected cost. The fuel planner may wish to weight certain outcomes more heavily than others for the purpose of decision-making, or perhaps he wishes to insure that certain cost outcomes cannot occur. There could be many reasons for this, including regulatory treatment of fuel costs or corporate policies.

The following three figures represent types of information that may be generated for the purpose of considering risk in contract mix decision-making.

Figure 9 shows the probability of cost exceeding high levels under the 60%, 90% and 120% contract strategies. These high cost levels occur when a high burn level occurs. From this measure of risk it appears that a contract fraction of 120% of expected burn is least risky, and 60% involves the greatest risk.

Figure 10 shows the possible average cost outcomes (total cost divided by burn) under the three strategies. The risk of a high average cost might be of interest if regulatory treatment of fuel costs (either explicitly or otherwise) tended to penalize a utility for having fuel costs that were considerably above market prices. Unlike high total costs, high average costs occur when burn is very low (and are infinite when burn is zero!). Based on this measure, the 120% contract fraction appears most risky, and a 60% contract fraction least risky - the opposite conclusion to the previous figure.

The range of outcomes of decision-dependent costs is displayed in Figure 11. For each possible burn outcome, the decision-dependent cost is the total cost minus the cost under perfect information (= \$30 times burn). It is the extra cost over what would be incurred if the fuel planner had perfect foresight. Unlike the previous two measures, the high decision-dependent cost outcomes occur when burn is either very high or very low. Note that the 90% contract fraction appears less risky than either the 60% or 120% strategy under this measure.

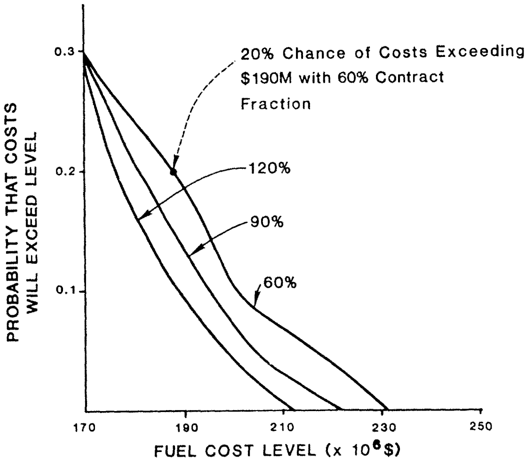

Figure 9. Probability of High Cost Outcomes

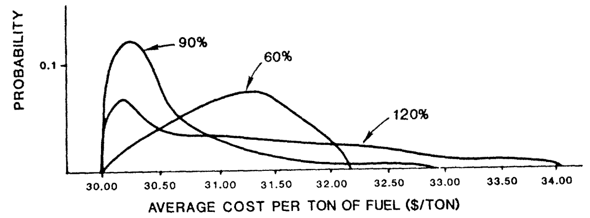

Figure 10. Probability of Average Cost Outcomes

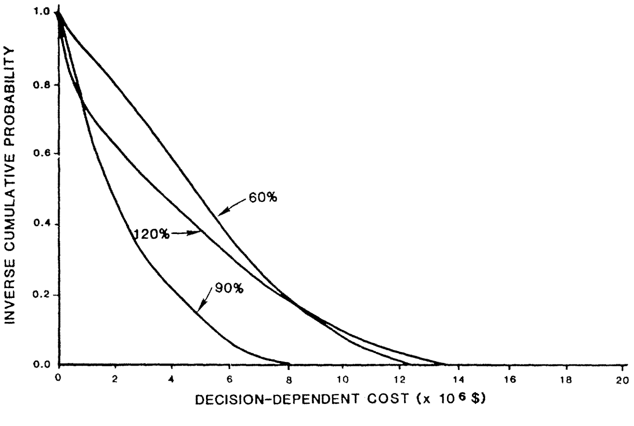

Figure 11. Probability of Decision-Dependent Cost Outcomes

Another common approach to determining the best decision under risk involves representing the risk aversion through a function that defines the value or "utility" of cost outcomes for the decision-maker [1]. The decision may then be found that maximizes the expected value of this function.

Fuel planners need more information about the possible outcomes of procurement strategies than expected costs. The reasons may involve risk considerations that are highly specific to each situation, and the appropriate risk measures will vary accordingly. A methodology that explicitly incorporates uncertainty is capable of generating the kinds of information needed to assess risks and make prudent decisions.

Uncertainty and Contract Flexibility

For the next examples, we introduce a new contract structure with a higher base price and more flexibility than the contract used in the previous example. The new contract's base price is \$0.50/ton more expensive, but there is a 20% flexibility range and only a \$3.00/ton underlift charge. Recall that the previous contract had no flexibility range and a \$6.00 underlift charge. Under the new contract, quantities as low as 80% of the contractual amount may be taken without penalty; below 80%, \$3.00 must be paid for each ton not taken. Henceforth, the previous contract will be called the inflexible contract, and the new structure will be called the flexible contract.

Figure 12 shows how the contract quantity that minimizes expected costs depends on the level of burn uncertainty and the degree of contract flexibility. For either contract, if there is no uncertainty a 100% contract fraction minimizes expected costs. For higher levels of burn uncertainty, the contract fraction increases or decreases depending on contract structure.

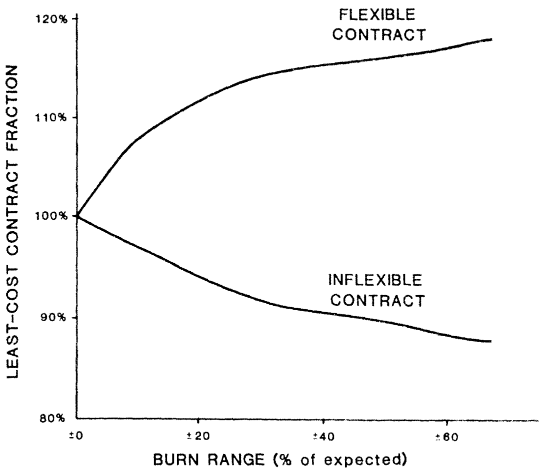

Figure 12. Expected Cost Minimizing Quantities of a Flexible or an Inflexible Contract

With an inflexible structure, the quantity tends to track the low end of the burn distribution; with a flexible structure, the quantity is determined by the high end of the burn distribution.

The optimal contract decision depends upon both the degree of burn uncertainty and the degree of contract flexibility. For this reason, rules of thumb can become outdated as conditions change. The optimal contract fraction from an expected value standpoint may be highly system- and plant-specific and may change over time.

In the next example, mixes of the inflexible and flexible contract are explored. Recall that the base price of the flexible contract is higher—otherwise it would always be preferred over the inflexible contract. Figure 13 shows how the mix that minimizes expected cost varies with the degree of burn uncertainty. Even when there is very little uncertainty, a mix of flexible and inflexible contract minimizes expected costs. With increasing uncertainty, the fraction of flexible contract increases.

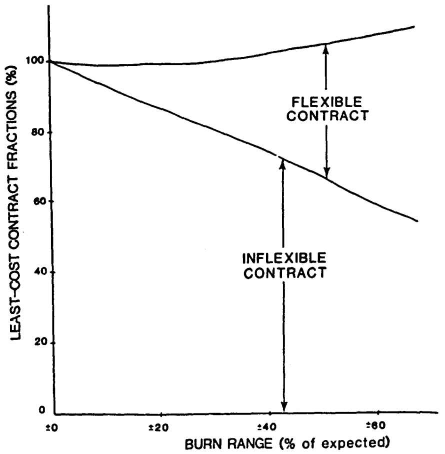

Figure 13. Expected Cost Minimizing Mix of Flexible and Inflexible Contract

Figure 14 shows how the expected costs vary under the three sets of options--flexible only, inflexible only, or a mix. In all cases, expected costs increase with greater uncertainty. With a flexible contract, costs increase with uncertainty at a slower rate. At any level of uncertainty, the mix provides the lowest expected costs.

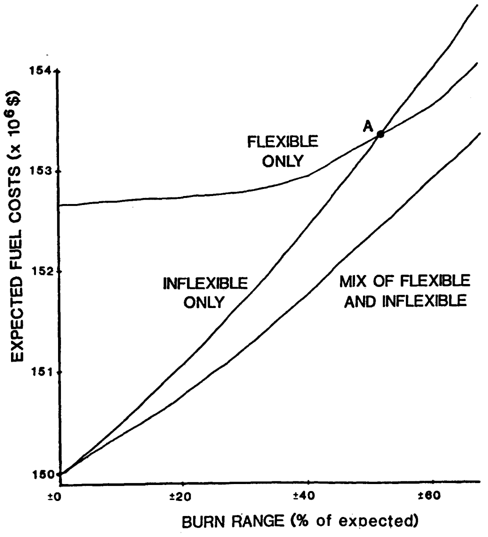

Figure 14. Minimum Attainable Expected Cost

The value of contract flexibility under the circumstances of the example may be deduced from this figure. The value starts at zero when there is no burn uncertainty and increases with increasing uncertainty. At the point where the costs are equal for the flexible and inflexible contract (marked A in the figure), the value of flexibility must be the difference in base contract price, or $0.50/ton. Incorporating uncertainty in the analysis therefore permits valuing contract flexibility.

The possible outcomes of decision-dependent costs are plotted in Figure 15. The flexible contract involves considerably less risk than the inflexible contract; but the price is a certain decision-dependent cost of at least $2.4M. The flexible contract is like an insurance policy relative to the inflexible contract. The mix of contracts reduces risk still further while also considerably reducing the fixed portion of the costs. Flexibility and the mix of contracts are even more attractive when risk is considered.

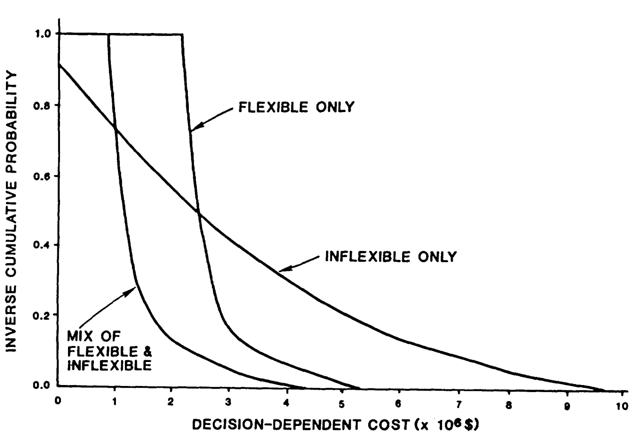

Figure 15. Probability of Decision-Dependent Cost Outcomes

Resolution of Uncertainty and Contingent Decisions

The previous examples have all been entirely static - time was not a factor. In this context we have demonstrated that consideration of uncertainty can be important in assessing both the costs and risks of contract decisions.

Uncertainty, however, changes over time. Information is gained about the likely outcomes of future uncertain events. Trends may develop or disappear, eliminating or reducing previous causes of uncertainty. Usually, our uncertainty declines over time. As we have seen, uncertainty has a cost, so how uncertainty resolves over time may be important when the timing of decisions is considered. Our final group of examples is intended to demonstrate that when uncertainty resolves, and what alternatives the fuel planner has available when it does, are also very important aspects of the contract mix problem and can influence current decisions.

We now add one more factor to our simple example model. Suppose that in addition to the previous uncertainty about burn, there is also a 50% chance that a nuclear plant will be coming on line. If the plant comes on line (is not delayed), the expected burn at our example plant will be reduced by 2 Mtpy, to 3 Mtpy. If the plant does not come on line, the expected burn will still be 5 Mtpy. For simplicity we assume that, in either case, there will still be the same variation around the expected burn. For this example, the only contract alternative is the inflexible contract described above.

If the nuclear plant is delayed, the fuel planner's choice to minimize expected costs will be as in the earlier examples, 4.5 Mtpy or 90% of expected burn (see Figure 7). If the plant will be online, the least-cost contract quantity turns out to be 2.5 Mtpy. The results given that the fuel planner knows the nuclear plant status at the time of his contract decision are summarized in Figure 16.

Now consider the situation where the fuel planner will not know whether the nuclear plant will be on line at the time he must commit to a contract quantity (perhaps his suppliers require a one-year lead time before deliveries commence). In this case his least-cost contract quantity is 3.4 Mtpy (Figure 17). This quantity is clearly a compromise--he is most likely over contract if the plant comes online, and under contract if it does not.

Without the benefit of the knowledge about the nuclear plant status, expected costs are $1.5M higher. The information about the nuclear plant status therefore is worth $1.5M to the fuel planner. He should be willing to pay up to $1.5M to find out the status of the nuclear plant in time for his decision. Alternately, he should be willing to pay up to $1.5M to be permitted by his supplier to delay his commitment. Clearly, when uncertainty resolves is an important factor in this example.

Suppose next that the fuel planner has a new contract option that has a base price 1% more expensive than the previous contract, and has no lead time. In this case, he minimizes costs by contracting for 2.6 Mtpy of the previous contract, and adding 1.9 Mtpy of the new, more expensive contract only if the nuclear plant turns out to be delayed (Figure 18). If the nuclear plant is on line no quantity of the new contract is selected, as its only advantage is its short lead time. Costs are now only $0.2M greater than if the information about the nuclear plant is available in time for all decisions.

With the addition of the new contract option, the cost-minimizing quantity of the previous contract to be selected is reduced from 3.4 Mtpy to 2.6 Mtpy. This example demonstrates that options available in the future may influence current decisions.

The final figure in this section (Figure 19) shows that not only does the information about the nuclear plant reduce expected costs, but also exposure to very high cost outcomes. The value of the information may be even higher when risk is considered.

We've seen that the timing of the information available to the fuel planner, and the future options available to him, may determine the best current strategy. Commonly-used analytical methods, which usually consider future scenarios individually, cannot determine the fuel planner's likely state of information at the time of future decisions. By representing how much the fuel planner will know and when he will know it, the methodology presented here permits simulation of a realistic reaction to events. The result is that the range of possible outcomes associated with any strategy is more accurately identified and the strategy may be more accurately evaluated.

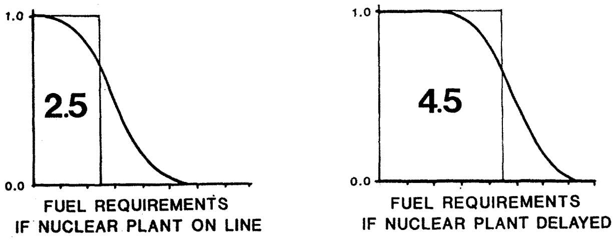

Figure 16. Expected Cost Minimizing Contract Decisions with Knowledge of Nuclear Plant Status. Expected Decision-Dependent Cost = $3.1M.

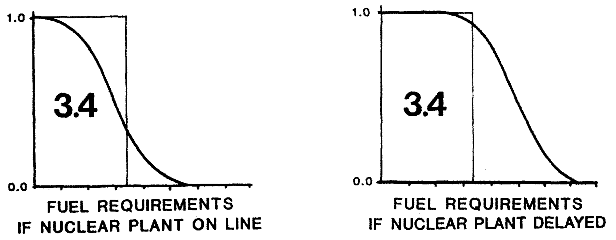

Figure 17. Expected Cost Minimizing Contract Decisions Without Knowledge of Nuclear Plant Status. Expected Decision-Dependent Cost = $4.6M.

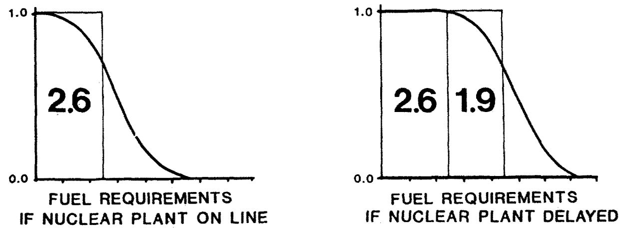

Figure 18. Expected Cost Minimizing Contract Decisions Without Knowledge of Nuclear Plant Status and With Short Lead Time Contract Option. Expected Decision-Dependent Cost = $3.3M.

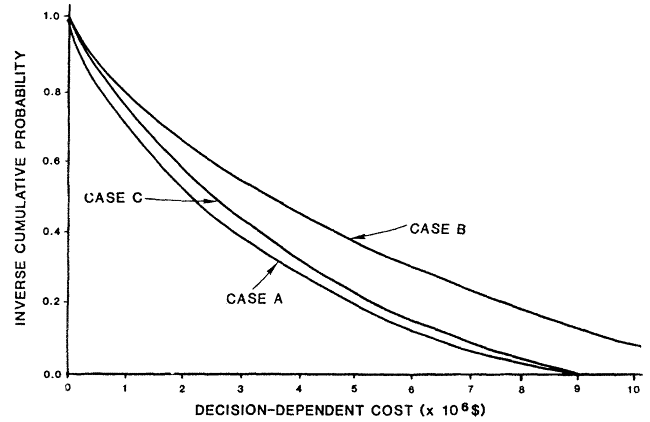

Figure 19. Probability of Decision-Dependent Cost Outcomes

SUMMARY

Through incorporating uncertainty explicitly in a comprehensive modeling framework (rather than piecewise through an approach such as scenario analysis) the costs and risks of contracting strategies may be identified and evaluated. The methodology permits generating types of information and comparisons that are necessary for sound fuel planning in the face of uncertainty and are not available using other methods.

Uncertainties result in higher expected fuel cost. The cost of uncertainty implies a value for flexibility that can reduce exposure to undesirable outcomes. The cost of uncertainty also implies a value for information that reduces the uncertainty. The value of specific types of information or flexibility may be quantified under a particular utility's unique set of circumstances.

Due to risk considerations, a fuel planner's objective may not be to minimize expected fuel costs. Incorporating uncertainty in the analysis allows generation of the kinds of information needed to evaluate exposure to various types of undesirable outcomes.

ACKNOWLEDGMENTS

The authors thank Charles G. Thompson, Ram C. Mukherji and George Clark of Southern California Edison Company, and Michael S. Hyrnick and Thomas Gaston of Ohio Edison Company, for furthering our understanding of the problems they face. The project was funded by the Electric Power Research Institute (RP2359-20).

REFERENCES

1. C. A. Holloway, Decision Making Under Uncertainty: Models and Choices. Englewood Cliffs, New Jersey: Prentice-Hall, Inc. 1979.

2. R. A. Howard, "Information Value Theory." IEEE Transactions on Systems Science and Cybernetics, vol. SSC-2, pp. 22-26, August 1966.

3. D. G. Luenberger, Introduction to Linear and Nonlinear Programming. Reading, Massachusetts: Addison-Wesley Publishing Co. 1973.

4. D. W. North, "A Tutorial Introduction to Decision Theory." IEEE Transactions on Systems Science and Cybernetics, vol. SSC-4, pp. 200-210, September 1968.