UTILTIY FUEL INVENTORY MODEL

TIPS AND TRAPS

Prepared by
Applied Decision Analysis Inc.

Updated by
Stephen W Chapel
S.Chapel Associates

July 2007

UFIM Tips and Traps

CONTENTS

UFIM Tips and Traps

TIPS 'N TRAPS #1: THE CRITICAL RATIO TEST

PROBLEM

I can't figure out why UFIM decides to protect against some disruptions and ignore others. Shouldn't it be protecting against all disruptions?

EXAMPLE

UFIM is giving me an inventory level that seems dangerously low. If a disruption ever happened, we would run out of fuel. I can't go to management with this; my job is to ensure that we never run out of fuel!

Specifically, my inputs are:

Average Plant Heat Rate:	10,000 Btu/kWh
Heat Content of Fuel:	10 MBtu/Ton
Real Annual Cost of Capital (without inflation):	10%
Tax Portion of Revenue Requirement:	5%
Physical Holding Cost per month: $0.125/month:	$1.50/year
Disruption Arrival Rate:	1 time every 10 years (in any month)
Shortage Cost: $20/MWh Fuel Purchase Price:	$35/ton.

SOLUTION

For each disruption that you model, do this quick test:

Calculate your holding cost (HC) per year per ton:

> HC = physical holding cost + (income tax rate x fuel purchase price) + (capital cost x fuel purchase price).
>
> HC = $1.50 + (0.05 x $35) + (0.10 x $35) = $6.75/Ton

Calculate your shortage cost per ton:

> SC = Shortage Cost x (1 Average Plant Heat Rate) x Heat Content of Fuel x 1,000.
>
> SC = $20/MWh x (1 10,000 Btu/kWh) x 25MBtu/Ton x 1,000 = $50/Ton

Calculate the probability of the disruption:

> P = number of occurrences / number of years. P= 1 / 10= 0.1

If HC > SC x P, then UFIM will probably not recommend insuring fully for that disruption, giving you lower inventory levels. If HO < SC x P, then UFIM probably will recommend insuring for the disruption, resulting in higher inventory levels. In this example, HC = $6.75 and SC x P = $5, so UFIM does not recommend insuring fully for the disruption.

UFIM Tips and Traps

EXPLANATION

UFIM's purpose is to minimize costs, not to guarantee that you will never run out of fuel. UFIM balances holding costs and shortage costs to arrive at its inventory policy recommendations. The reasoning is as follows:

If your utility stocks up to have enough fuel to outlast the disruption, the cost is known; it is the holding cost per ton multiplied by the number of tons needed. In the example, assuming that you need 50,000 tons to survive the disruption, and that you stock up for the entire year, the annual cost is $6.75/ton x 50,000 tons = $337,500 (of course, if we only needed to stock up for part of the year, this cost would be less).

As the example shows, the holding cost can be very large, so UFIM determines whether the risk of the disruption really justifies spending the money.

If the disruption were certain to occur, then UFIM would simply compare the cost of running out of fuel with the cost of holding it. The cost of the disruption would be the shortage cost per ton multiplied by the number of tons short, or $50/ton x 50,000 tons = $2,500,000. If, as in the example, this shortage cost is greater than the holding cost, it would make sense to hold the extra inventory and UFIM would recommend stocking up to avoid running short. Of course, if this cost is not greater, then it would make no sense to stock up; no one wants to spend a lot of money to save a little.

Unfortunately, disruptions are by definition uncertain events. When uncertainty is present, it is necessary to decide whether the cost of being 100% protected is justified. An extreme example of this kind of decision would be trying to avoid any possibility of being involved in an automobile accident. To accomplish this, you would have to give up driving, riding in cars, taxis, buses and limousines, riding bicycles, and walking on sidewalks. If you knew with certainty that you would be involved in an auto accident next week, you might be willing to do without these conveniences for one week. However, since you know that an accident is unlikely, this cost is simply too high, and you will probably decide to live with something less than 100% protection.

In UFIM, the uncertainty of a cost is handled by using "expected" costs. The expected cost of the example disruption is the annual cost of the disruption, SC x Tons, multiplied by the annual probability that the disruption occurs, P, or

SC x Tons x P = $50/ton x 50,000 tons x 0.1 = $250,000.

The intuitive explanation of expected cost is that if you look at a fairly long period of time, say 10 years, and the disruption happens once during that time (P = 0.1) then the average cost per year is

(SC x Tons) 10 years, which is the same as SC x Tons x P.

Once the uncertainty of the disruption has been accounted for in the expected cost, UFIM can directly compare the known cost of holding inventory with the expected cost of running short. Obviously, if the expected cost of running short is greater than the cost of holding inventory, then the proper policy is to hold inventory. But if the expected shortage cost is less than the holding cost, then it is too expensive to protect against the disruption and UFIM will not recommend holding that inventory. A good approximation of UFIM's decision rule is, thus:

If HC < SC x P then hold inventory.

If HC > SC x P then do not hold inventory.

The rule can be rewritten to say that UFIM recommends holding inventory when HC / SC < P. This is called the "critical ratio test" and HC / SC is called the "critical ratio."

Applying this test to the example disruption, you can see that, if

HC = $6.75, and SC = $50.00, then HC / SC = 6.75 50 = 0.135.

Since you know that P = 0.1, and that .135 > 0.1, you can predict that UFIM will decide that the cost of holding sufficient inventory to protect against this disruption is too high to be justified and will recommend a lower inventory policy.

REFERENCES

UFIM 2.0 Manual pp. 1-4 through 1-16.

Morris, P., M. James Sandling, Richard B. Fancher, Michael A. Kohn, Hung Po Chao, and Stephen W. Chapel, "A Utility Fuel Inventory Model," Operations Research Vol 35, No. 2, March April 1987, pp. 169 184.

UFIM Tips and Traps

TIPS 'N TRAPS #2: CHOOSING DISRUPTIONS TO MODEL

PROBLEM

I have nine disruptions to model, and UFIM will only let me model 25 "disruption starting months"! How do I choose disruptions without losing important detail in my case?

EXAMPLE

Here are my potential disruptions:

Extra Hot Summer	Barges Fail
Extra Cold Winter	Coal Pile Freezes
Coal Unloader Outage	Frozen River
Transport Equipment Breaks	Rail Delay
Coal Miner Strike	

SOLUTION

Try following these four steps to generate a list of disruptions to model:

Step 1: List Possible Disruptive Events

Step 2: Check Disruption List for "Normal Times" Situations

Step 3: Check Disruption List for Disruptions That Can Be Combined

Step 4: Narrow Down the Final List.

UFIM Tips and Traps

EXPLANATION

Disruptions are modeled in UFIM as events that affect supply, demand, and other system characteristics differently from normal operations. Disruptions and their associated warnings are specified by the following characteristics.

- Annual frequency
- Seasonal incidence
- Duration
- Warning
- Fuel prices
- Demand distributions, by month
- Delivery constraints, by month
- Replacement power costs, by month.

Throughout the process of specifying and narrowing down disruptions, it is important to keep in mind the above list of ways in which disruptions may differ from normal times. As you will see, many potential disruptions can actually be modeled using the normal times inputs, while others cannot.

STEP 1 LIST POSSIBLE DISRUPTIVE EVENTS

The first step in defining disruptions is to make an exhaustive list of events (disruptions) that may affect your normal plant operations. Make sure to include every event that might be important, if it is not already characterized in normal times data. For example, if coal supply is always reduced in June, due to miners' vacations, you may have already characterized it as a normal times constraint, not as a disruption that occurs every June.

For example, the list of possible disruptions might be the same as the one on Page 1. In Step 2, we will re visit this disruption list.

STEP 2 CHECK DISRUPTION LIST FOR 'NORMAL TIMES SITUATIONS

Now that you have listed all the potentially important disruptive events you can, check the list one by one for events that can be incorporated into your normal times data.

For UFIM purposes, the following conditions must be represented as a disruption:

- Events that can last for more than one order period
- Events leading to disruption management ("burn reduction")
- Events causing fuel prices to differ from normal times
- Events causing replacement power costs to differ from normal times
- Events causing different order constraints.

UFIM Tips and Traps

For each of the disruptions on your list, check to see whether it meets any of these conditions. If it does, then (for now) leave it on the disruption list. If it does not, then it can be folded into the description of normal times.

A few tips: If you are in doubt, it maybe easier and clearer to treat an event as a disruption, even if it is not strictly necessary. Also, some of the events on your list may be combined, and treated as a single disruption, (see Step 3).

To continue the example using the list of nine disruptions from Step 1, you might know that "Rail Delay" causes a very short delay in the coal supply (just 2 to 3 days). In addition, "Extra Hot Summer' happens on average once every four years, and causes a demand increase that usually lasts no more than two to three weeks. This can be represented by adding a "high demand" outcome to the normal times demand distribution for June, July, and August, re scaling the normal times distribution accordingly. There is no need to model the event as a disruption. This leaves seven disruptions on the list.

To conclude Step 2, construct a revised list of potential disruptions, and revise normal times data accordingly.

STEP 3 CHECK DISRUPTION LIST FOR DISRUPTIONS THAT CAN BE COMBINED

Using the potential disruption list from Step 2, look for any disruptions that can be grouped together and modeled, for example, as one disruption that can happen in several different months. Disruptions that are grouped together **should be caused by similar events** -- otherwise modeling them can become confusing. In addition, they should:

- Have the same distribution on duration.
- Have the same effect on order constraints.
- Have the same effect on prices, demand, and replacement power costs, for any particular month. (These inputs are specified monthly.)

Disruption	Duration	Manage Burn Reduction	Arrives In	Fuel Demand	Max Order	Fuel Price	Replacement Power Costs
Unloader Outage	1 Week	No	Oct-Dec	Normal	0	$30/ton	Normal
Frozen River	3-4 weeks	No	January	Normal	0	$ 30/ton	Normal
Barges Fail	3-4 weeks	no	February	Normal	0	$30/ton	Normal

Table 1

A table like Table 1 may help you determine how disruptions can be grouped. Table 1 shows a summary of the characteristics of three possible disruptions, all of which would interrupt supply. Disruptions #1 and #2 have different durations, and cannot be grouped together. However, Disruptions #2 and #3 have the same distribution on duration and can be modeled as one disruption with chances of occurring in January and February.

UFIM Tips and Traps

Continuing the example, you might find that of the seven remaining disruptions, there are three pairs that each can be modeled as one disruption. "Frozen River" and "Barges Fail" can be modeled together as shown in Table 1. In addition, "Extra Cold Winter" and "Coal Pile Freezes" both occur in January and February and can last 1 2 weeks; they have the same effect, so are modeled together. Also, "Coal Unloader Outage" and "Transport Equipment Breaks" have the same effect; when the unloader isn't working, the coal can't be transported, and vice versa. Thus, these two disruptions are modeled as one.

- The list of potential disruptions is now:
- Cold Winter/Frozen Coal Pile
- Unloader or Transport System Outage
- Barge Fails/River Freezes
- Coal Miner Strike.

STEP 4 NARROW DOWN THE FINAL LIST

There are two methods for narrowing down the final disruption list. Each of these is a little tricky:

- The Critical Ratio Test, or holding cost/expected shortage cost tradeoff; this is discussed in UFIM Tips & Traps #1.
- Running UFIM with the most severe disruption(s) first, to see whether other disruptions are "covered" automatically.

In the Critical Ratio Test, the holding cost is compared to the expected shortage cost to determine whether to model a particular disruption. This can be a good way to identify those disruptions that have the most severe "critical ratio", or potentially high shortage cost compared to a low holding cost. Disruptions that are not potentially severe with respect to shortage cost may be deleted from the list. For example, the Unloader/Transport System Outage might have a critical ratio that is significantly greater than the probability of the event. If this were the case, you could drop the disruption from the list.

The second method for narrowing down disruptions is to run UFIM for the most severe disruption first. If the resulting inventory targets are so high that they would insure, not only the severe disruptions, but the less severe disruptions as well, then you may drop the less severe disruptions from the list. However, this method requires careful analysis and may be difficult to use when disruptions occur in several different months.

REFERENCES

UFIM Version 2.0 Manual pp. 1 3 to 1 13. UFIM Tips & Traps #1: The Critical Ratio Test. ADA/EPRI 1989 Fuel Inventory Planning Workshop Workbook.

UFIM Tips and Traps

TIPS 'N TRAPS #3: DETERMINING REPLACEMENT POWER COSTS

PROBLEM

I'm having difficulty specifying my "burn reduction" or "replacement power" cost curves for UFIM. How are these related to "shortage costs?" What are the different ways to represent them?

EXAMPLE

I specified my normal-times replacement power cost curve as:

Fractional Reduction	*$/MWh*
0.60	$15
0.60	$20
0.80	$20
0.80	$25

My delivered fuel price is $48/ton; heat content is 25 MBtu/ton; and heat rate is 10,000 Btu/ kWh. UFIM is giving me a fatal error message when I try to run my case because these replacement power costs are too low.

SOLUTION

UFIM uses the cost of replacement power and the price of fuel to compute shortage costs per ton for any fractional burn reduction. These shortage costs are used to compute overall costs for each policy UFIM evaluates. The average shortage cost perton of burn reduction is the difference between the average replacement power cost and the fuel price. Thus, the replacement power cost should be greater than the fuel price. Otherwise, shortages create savings rather than costs.

UFIM checks to make sure that your replacement power costs are greater than your fuel price. If they are not, UFIM gives you a fatal error message, because this implies that the average shortage cost can be negative. In the example above, UFIM would compute the average replacement power cost as

(0.6 x $1 5/MWh) + (0.2 x $20/MWh) + (0.2 x $25/MWh)

= $18/MWh, equivalent to $45/ton

and the fuel price is $48/ton, which is greater than the replacement power cost. This implies that the shortage cost is negative (-$3/ton), and UFIM gives a fatal error.

To avoid this fatal error, replacement power curves should reflect increased costs of shortages at the unit. Replacement power cost curves can be input as "step functions" that explicitly reflect the sharp increases in system costs as additional cycling plants are loaded and as expensive off-system power is purchased. They can also be input as smooth curves with gradually increasing replacement power costs, similar to a system marginal cost curve.

UFIM Tips and Traps

EXPLANATION

Two topics on replacement power costs are discussed in the sections below:

1) The two common types of curves, and how they are used to compute average shortage costs and converted into dollars per ton;
2) How the two types of curves are represented in UFIM.

TWO TYPES OF REPLACEMENT POWER COST CURVES

Figure 1 depicts the type of replacement power cost curve that UFIM users most commonly use. The curve shows the cost of replacing power that would otherwise be generated by the unit in question. The x-axis (Burn Reduction %) represents the fraction of time during which replacement power is available at the price indicated on the y-axis. UFIM assumes that fuel burn at the unit in question is reduced when the replacement power cost is lowest. Thus, if the unit is turned off 10% of the time (i.e. burn is reduced by 10%) then replacement power is available at $15 per megawatt hour during this time, as shown in the figure.

The form of the curve in Figure 1 is a "step function." The cost per megawatt hour makes discrete jumps at certain percentages of burn reduction along the curve. This is a convenient way to represent the rapid startup of each increasingly expensive source of replacement power as the need for the power increases. In the figure, the first source of replacement power for the coal unit would be an oil unit, available 60% of the time at a cost of $1 5/MWh; next, a combustion turbine would be used an additional 20% of the time at a cost of $20/MWh, and finally off-system power would be purchased for $25/MWh.

Figure 1 also shows the cost of purchasing fuel for the unit in question. The average shortage cost (discussed below) is the difference between the average replacement power cost and the fuel cost. Notice that the fuel cost should always be less than the cost of purchasing replacement power (although in the original example, it was not).

The second type of replacement power cost curve is depicted in Figure 2. This curve is "smooth", without discrete jumps in the replacement power cost. The curve is derived from a system marginal cost curve, which can be used to determine replacement power costs for a baseload unit. In a large system, the system marginal costs maybe fairly continuous. In this case, the replacement power costs are modeled best as depicted in the figure.

Shortage Cost. UFIM computes the average shortage cost for a fractional burn reduction as the difference between the average replacement power cost up to that fractional reduction, and the fuel cost. For example, the average replacement power cost forthe entire curve in Figure 1(100% reduction) is:

(0.6 x $1 5/MWh) + (0.2 x $20/MWh) + (0.2 x $25/MWh)

= $18/MWh.

UFIM Tips and Traps

FIGURE 1

SAMPLE REPLACEMENT POWER COST CURVE

"STEP FUNCTION"

SOURCE OF POWER	COST	%OF TIME
OIL	$15	60%
CT	$20	20%
PP	$25	20%

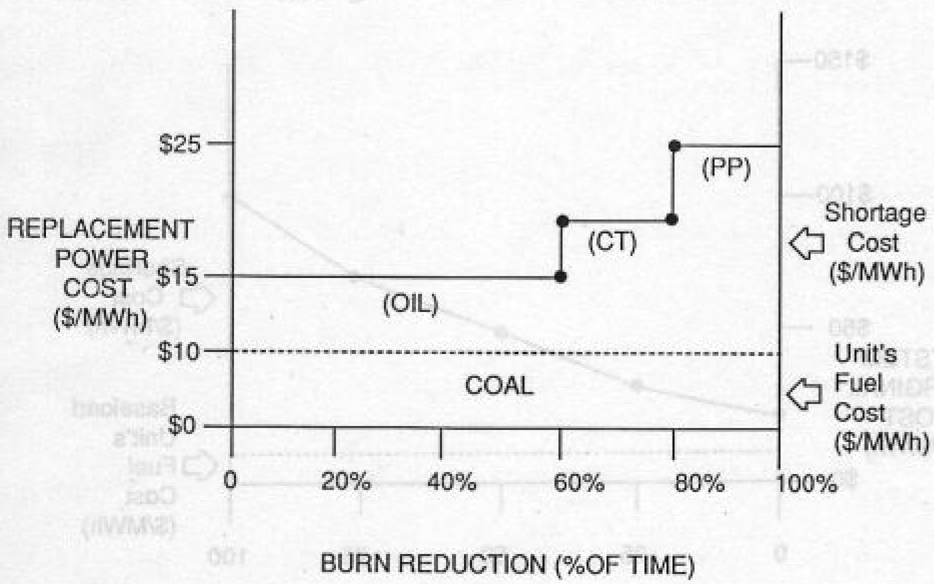

UFIM INPUT:

FRACTIONAL REDUCTION	$/MWh
0.60	$15
0.60	$20
0.80	$20
0.80	$25

UFIM Tips and Traps

FIGURE 2

REPLACEMENT POWER COST CURVE

"SMOOTH"

FOR BASELOAD PLANT, REPLACEMENT POWER COST ~ SYSTEM MARGINAL COST

COST	%OF TIME
$20	@0%
$30	@25%
$45	@50%
$75	@75%
$100	@100%

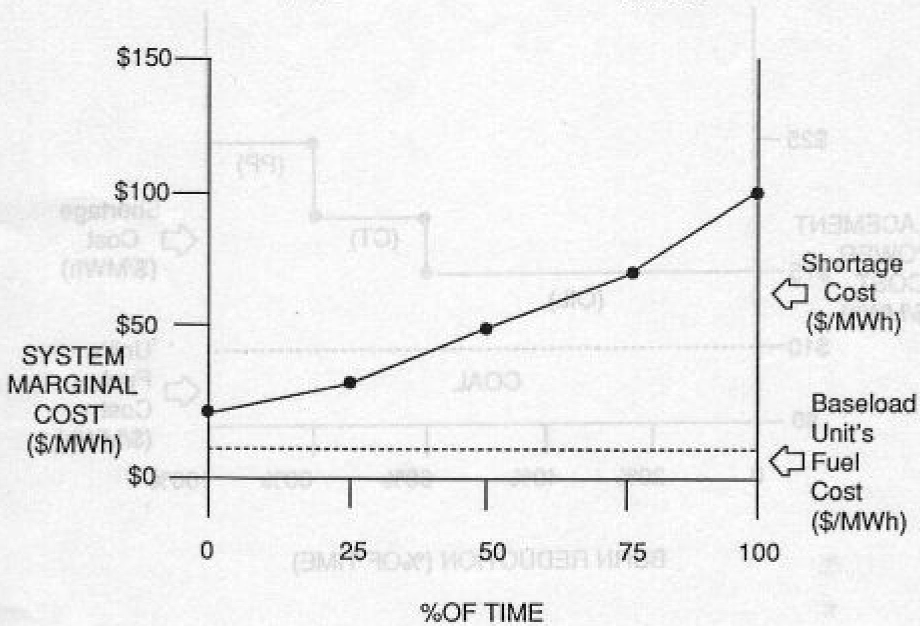

UFIM INPUT:

FRACTIONAL REDUCTION	$/MWh
0.0	$20
0.25	$30
0.50	$45
0.75	$70
1.00	$100

UFIM Tips and Traps

The fuel cost in this example is $1 0/MWh, so the average shortage cost per MWh is:

$18/MWh - $10/MWh $8/MWh.

UFIM would convert this cost into dollars per ton; the conversion process is given below.

Conversion of Costs to Dollars per Ton. Replacement power costs are specified in UFIM in units of dollars per MWh, while fuel shortages are measured in tons of fuel. UFIM converts replacement power and shortage costs into dollars per ton using a simple formula:

Cost ($/Ton) = Cost($/MWh) x 1 Average Plant Heat Rate
x Heat Content of Fuel x 1,000.

For example, using the inputs above, a shortage cost of $8/MWh would be converted to:

$8/MWh x (1 10,000 Btu/kWh) x 25 MBtu/Ton x 1,000

= $20/Ton.

A tip is: Check your UFIM SETUP output for the average fuel replacement cost. This cost is the area under the entire replacement power cost curve, converted into dollars per ton (without subtracting off the fuel cost). Make sure that you can reproduce this UFIM output using your own derivation.

REPRESENTATION OF REPLACEMENT POWER COSTS IN UFIM

For UFIM purposes, replacement power cost curves are specified as a set of (x,y) coordinates corresponding to the breakpoints on the curve. You only need to specify the breakpoints, and you may specify upto five of them. Figures 1 and 2 give examples of UFIM specifications of replacement power cost curves.

Two tips are: If the breakpoints specified on the curve do not include a point for 0% reduction and/or a point for 100% reduction, UFIM "draws" a straight horizontal line from 0% to the smallest reduction fraction specified, and another from the largest reduction fraction specified to 100%. Thus, any fractional reduction below (above) the smallest amount on the curve is assumed to cost the same per MWh as the smallest (largest) amount specified.

The easiest way to determine the replacement power inputs for UFIM is to carefully draw the curve and circle the breakpoints before entering them into UFIM.

REFERENCES

UFIM Users Manual (Version 2.0), pp. 2-4 to 2-9
ADA/EPRI 1989 Fuel Inventory Planning Workshop Workbook.

UFIM Tips and Traps

TIPS 'N TRAPS #4 THE MAXIMUM INVENTORY CONSTRAINT

PROBLEM

When I simulate target policies that are close to my maximum inventory level, I get very high price opportunity costs. What does this mean?

EXAMPLE

Some of my case data looks like this:

Fuel Price:	$40/Unit		
Demand Distribution – Amount	50kUnit	200 kUnit	350 kUnit
Prob.	0.25	0.50	0.25
Maximum Inventory	900 kUnit		

When I simulate these targets, I get the following cost table:

Target Policy	Total Cost	Annual Total Cost	Annual Burn Cost	Annual Holding Cost	Annual Oppty. Cost	Annual Shortage Cost
180	1368.83	95.82	88.95	0.73	0.00	**6.14**
270	1335.58	93.49	88.95	1.07	0.00	**3.47**
360	1314.55	92.02	88.96	1.42	-0.02	**1.66**
450	1307.96	91.56	88.96	1.78	-0.02	**0.84**
540	1308.72	91.61	88.95	2.15	0.00	**0.51**
630	1310.84	91.76	88.95	2.56	0.00	**0.29**
720	1314.00	91.98	88.96	2.88	-0.01	**0.16**
810	1318.42	92.29	88.95	325	-0.01	**0.10**
900	**1408.47**	**98.59**	**88.95**	**3.55**	**6.02**	**0.07**

At a target of 900, my maximum inventory level, the price opportunity cost is $6.02 million, which is very high. The price opportunity costs for the other target policies are zero, or very close to zero ("just noise").

SOLUTION:

To reduce the price opportunity cost, you should either increase your maximum inventory level, or decrease your targets. If you must keep the numbers as they are, keep in mind that most of the reported price opportunity cost is really the cost of "thrown away inventory."

EXPLANATION:

When the UFIM simulation calculates the costs incurred in a given period as a result of following a certain target policy, it first computes the fuel order and adds the cost of the fuel received to the total cost. The fuel order is based on the current inventory level, the expected demand, and the order constraints. Next it simulates the actual demand outcome and allocates the costof this amount offuel to the burn cost. If the actual demand is less than the expected demand, on which the order is based, the program stores the "extra" fuel in inventory. If the inventory level is well below the maximum, this causes no problems.

However, a problem does arise when a target policy is close or equal to the maximum inventory. The model assumes that the maximum inventory cannot be exceeded under any circumstance. If the ending inventory (i.e., starting inventory plus the order received minus the fuel burned to meet demand) is greater than the maximum inventory, the model effectively "throws away" the fuel in excess of the maximum inventory. The cost of this "thrown away" inventory is placed in the price opportunity cost category.

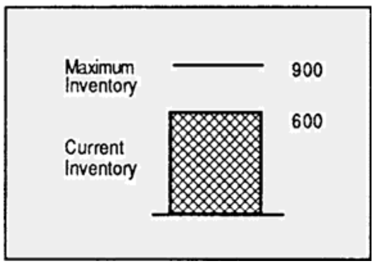

To demonstrate how this can happen, let's look at the example. We can see that the highest policy simulated has targets equal to the maximum inventory. Suppose that expected demand in the next period is 200 kilo units, but could actually be as low as 50 kilo units. The current inventory is 600 kilo units.

When it simulates this period, the model will order 500 kilo units, which is the target minus starting inventory plus expected demand, or900-600+200. The order is received, and the total cost is increased by, say, $40/unit multiplied by the order amount, or $20 million.

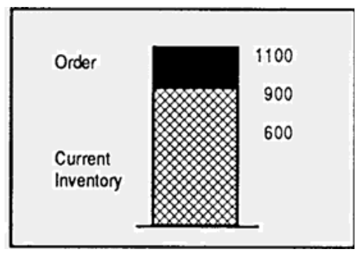

Next the demand outcome is chosen, and it is equal to a demand for 50 kilo units, 150 less than the expected demand. Thus, $2 million, or $40/unit multiplied by the demand, is added to the burn cost.

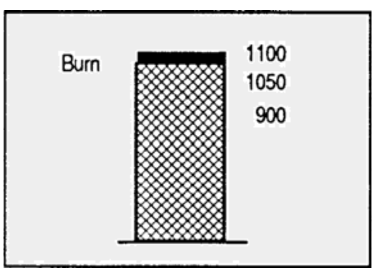

This leaves 150 kilo units of inventory in excess of the maximum, or600 + 450, which is 1050. The maximum inventory is 900, so the model sets inventory equal to 900, and effectively throws away the extra 150 kilo units of fuel. The cost of this thrown away fuel is allocated to the price opportunity cost, which goes up by $40/unit multiplied by 150,000 units or $6 million. If inventory must be "thrown away" very often, the total cost allocated to this category could be very high.

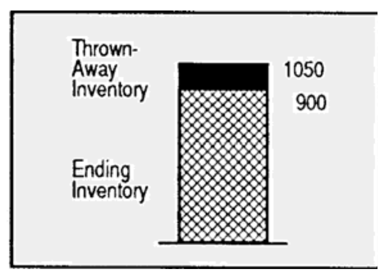

If you get a high opportunity cost because of this model assumption, think about your maximum inventory constraint input. If the maximum inventory level you are using is not a strict constraint, you may want to increase it when simulating high target policies.

Of course, sometimes a system does indeed have a strict inventory constraint. If such a system attempts to use targets close to this constraint, ft may at times incur costs for either denying deliveries or finding an alternative storage location. If this is the case for your utility, you may wish to use your UFIM output to roughly estimate these costs.

UFIM Tips and Traps

Tips:

1) Price opportunity costs will not usually be close to zero if prices vary over the course of the year, and/or are different during disruptions than during normal times. If your output shows high price opportunity costs and your case has varying prices, you should consider this when trying to determine the cost of "thrown away inventory that UFIM is reporting.

2) Remember that the UFIM simulation draws a distinction between expected demand and actual demand. UFIM assumes that fuel orders are placed to meet expected, or average demand. However, in many cases, the expected demand will never occur; the only possible demand outcomes are above or below the expected demand! The less uncertainty there is around demand, the less the "thrown away" fuel cost will be. Greater demand uncertainty can lead to high "thrown away" fuel costs, becasue it is more likely that the demand outcome will be less than the expected demand. Thus it is more likiely that the inventory will exceed the maximum at the end of an order period.

UFIM Tips and Traps

TIPS 'N TRAPS #5: MODELING JOINT DISRUPTIONS

PROBLEM

The UFIM case I am developing includes both a nuclear outage disruption and a disruption from a coal equipment failure. These disruptions have sometimes occurred simultaneously in the past. I want to model this in UFIM, but UFIM assumes that disruptions do not occur simultaneously. How can I include the joint occurrence of these disruptions in the model?

EXAMPLE

The nuclear outage happens in April, on the average of once every two years. During the outage, the levels in the demand distribution increase five steps. The equipment failure cuts my coal supply entirely for a month and can happen in any month. This event occurs once a year on average.

SOLUTION

The process for solving the problem of joint disruptions has two main steps:

1) Define the separate disruptions and the joint disruption as three distinct disruptions that cannot happen simultaneously, and

2) Determine the necessary model input data to enter for these distinct disruptions.

EXPLANATION

STEP 1: DEFINING DISTINCT DISRUPTIONS

The problem of joint disruptions is, of course, that two events are occurring simultaneously. The objective of the first step is to map out a list of all potential distinct combinations of single and multiple disruptions that can occur; "distinct" indicates that no two of the disruptions can occur at any onetime.

An easy way to create this list is to construct a tree diagram of the possible disruptions. Each node on the tree represents an uncertainty that has two or more possible outcomes; outcomes are represented by branches. The first node on the tree has two branches, representing whether or not the first event (here it is the "equipment failure") occurs. Each of these tower branches is then split to represent the possibility of the second event (in our

example "nuclear outage") occurring. At the bottom of the diagram, we have one branch for each of our distinct disruptions; these new disruptions are defined by the occurrence of a unique group of the original events. For the example given, the tree would look like the following:

UFIM Tips and Traps

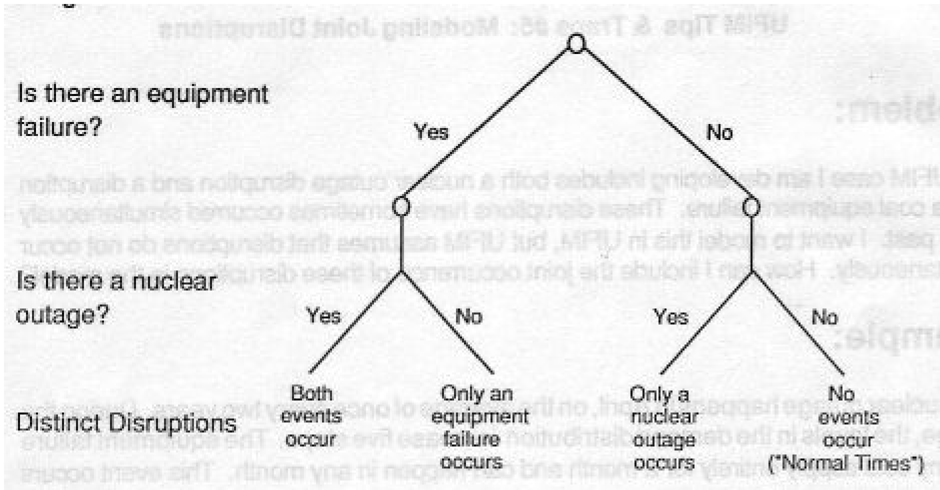

In this example, the disruption "Both events occur" is defined as the disruption where we have the simultaneous occurrence of both an equipment failure and a nuclear outage.

If we construct our disruption list to include all of these combinations of events - equipment failure, nuclear outage, and both events - we will cover all possible situations while not allowing any two disruptions to happen concurrently. The advantage in characterizing disruptions this way is that exactly one of them happens at any time; by construction there is no possibility that two can happen simultaneously.

This list, however, is only a preliminary list. In constructing a list of potential disruptions, you should test to see that all disruptions are actually significant enough to include in UFIM explicitly; you should also pay attention to the possibility of combining disruptions. Further details on how to refine the list can be found in Tips & Traps #1 (The Critical Ratio Test) and #2 (Choosing Disruptions to Model). In many cases the list cannot be finalized until after the second step (Determine model input data, discussed below) has been completed, because the critical ratio test requires data resulting from that step.

For our example we will assume that the critical ratio test indicates that all three of these disruptions are significant enough to model. (It should be evident that it is not appropriate to combine any of them, as they have very different effects.)

We have now tackled the first step, defining distinct disruptions that cannot happen simultaneously. The next step is to calculate the necessary data to enter for these disruptions.

STEP 2: DETERMINE THE NECESSARY MODEL INPUT DATA FOR THE DISTINCT DISRUPTIONS

Each of the three distinct disruptions defined so far is different in someway from each of the two original events. To model these disruptions successfully, it is crucial to think of them as different; no data should be "carried over" from the original events unless it is appropriate to do so. "Equipment failure only," for example, is different from the original event "Equipment failure" because the latter does not exclude the possibility that a nuclear outage occurs; data associated with "Equipment failure" should not be blindly copied and used for the disruption" Equipmentfailure only." The correct wayto construct these data inputs is to return to the sources of the original event data, perhaps historical data or expert consultation, and use that data to calculate the necessary inputs for the distinct disruptions.

Below we give examples of how the frequencies and monthly arrival probabilities might be calculated from historical datafor each of the three distinct disruptions.

Historical data for coal equipment failures and nuclear outages over the past ten years show the following monthly distributions:

Number of Years Disruption Occurred in Month:							
	Jan	Feb	Mar	Apr	May	Jun-Dec	**Total**
Equipment failure	3	2	2	2	1	0	**10**
Nuclear outage only	0	0	0	5	0	0	**5**

In April, one of the equipment failures occurred at the same time as a nuclear outage. Therefore the historical data for the three distinct disruptions are

Number of Years Disruption Occurred in Month:							
	Jan	Feb	Mar	Apr	May	Jun-Dec	**Total**
Both failure and outage	0	0	0	1	0	0	**1**
Equipment failure only	3	2	2	1	1	0	**9**
Nuclear outage only	0	0	0	4	0	0	**4**

From this updated historical data we can now calculate the frequencies and relative monthly probabilities for the distinct disruptions.

There has been one occurrence of the joint disruption ("Both failure and outage") in the past ten years, so the frequency is estimated to be ten years between disruptions. The disruption always happens in April, so the relative probability of occurrence is 1 in April, and 0 for all other months. (We assume that, because five nuclear outages have occurred in April and none have occurred in other months, there is no chance of a nuclear outage in months other than April. Expert judgment would be necessary to validate this assumption.)

For the disruption "Equipment failure only," the frequency is given by the following:

Frequency of the disruption

"Equipment failure only" = # years data / # of "Equipment failures only"

= 10/9

= 1.11 years between disruptions

UFIM Tips and Traps

The relative monthly probabilities are calculated as follows:

Relative monthly probability = #times in month /total occurrences

= 3/9= .333forJanuary

= 2/9 = .222 for February and March

= 1/9= . 111 for April and May

The historical data suggest that the "Equipment failure only" is three times as likely to occur in January, and twice as likely to occur in February or March, as it is to occur in April or May. Again, it would be prudent to verify (and possibly revise) these estimates using other expert data sources, as it would be dangerous to rely solely upon historical data from such a short period of time.

We now move on to similar calculations forthe disruption "Nuclear outage only."

In the past ten years, there have been five nuclear outages, one of which occurred at the same time as an equipment failure. This leaves four nuclear outages that have occurred alone in the past ten years.

Frequency of the disruption

"Nuclear outage only" = # years data! # of "Nuclear outage only"

= 10/4

= 2.5 years per disruption

As this disruption always happens in April, the relative probability of occurrence is 1 for April and 0 for all other months.

A tip is: Whenever possible, check that the new model input data is consistent with your intuition about the distinct disruptions. Consider the example above. Intuition would tell us that "Nuclearoutage" occurs more often than "Nuclear outage only"; the former includes the possibility of a joint event, whereas the latter does not. The disruption frequency calculated agrees with our intuition here. The frequency for "Nuclear outage only" is 2.5 years, while "Nuclear outage" occurs more often, once every 2 years on average (5 occurrences in the past 10 years = 2 years per disruption). Intuition can also help us understand the effect that a potential joint event has on monthly probabilities. The relative monthly probabilities for "Equipment failure only" are greater than those for "Equipment failure" in all months but April. This is because the joint event, which occurs only in April, is not included in "Equipment failure only." Excluding the possibility of some of these April events makes the disruption relatively more likely to occur in one of the other months.

UFIM Tips and Traps

After the frequencies and monthly probabilities have been determined for the distinct disruptions, it remains to characterize their demand distributions, burn reduction costs, and ordering data. Below is a chart of the probable effects that each of the three distinct disruptions will have on these data.

	Demand	Burn Reduction Costs	Maximum Order Allowed
Nuclear outage only	Higher	Higher	Normal
Equipment failure only	Normal	Normal	Less
Both outage failure	Higher	Higher	Less

The "Nuclear outage only" disruption increases demand and causes higher burn reduction costs due to heavier system-wide loads. The coal plant' s supply system is unaffected by this disruption, so it is reasonable to suppose that the maximum allowed order would remain normal.

The "Equipment failure only" disruption affects only the supply system. We would expect our ability to order coal to be restricted, but would expect no effect on the demand or burn reduction costs.

The joint disruption, "Both nuclear outage and equipment failure," gives us the undesirable effects of both previous disruptions. The equipment failure prevents normal order amounts, and the extra system-wide load from the nuclear outage translates into both higher demand and higher burn reduction costs.

A Tip is: In many cases, the combined effects of the individual events in a joint disruption might be much more severe than the sum of the effects of each considered alone. For example, think of the joint disruption of a coal strike and a transportation failure for alternative sources of fuel. These events may have similar effects individually, but in tandem may present drastically higher risks. The key to modeling such joint disruptions is to think of the joint occurrence as a unique disruption for which new data must be acquired.

REFERENCES:

UFIM User's Notebook, Tips & Traps #1, "The Critical Ratio Test."

UFIM User's Notebook, Tips & Traps #2, "Choosing Disruptions to Model."

UFIM Tips and Traps

TIPS 'N TRAPS #6: DISTRIBUTIONS ON INVENTORY LEVELS

PROBLEM:

I don't understand the different inventory distributions given in the UFIM outputs. What is the difference between the distributions given for disruption/normal periods and normal periods only? And why are some of my stockout probabilities so high?

SOLUTION:

To understand and use the probability distribution outputs from IJFIM, you must understand the following three concepts:

1) cumulative distributions,
2) the difference between distributions for both normal and disrupted periods and those for normal periods only, and
3) the difference between stockout and shortage probabilities.

EXPLANATION:

CONCEPT 1: CUMULATIVE DISTRIBUTIONS

When we speak of a "probability distribution," we refer to the set of potential outcomes of an uncertain event where each outcome is assigned a probability of occurring. The only requirements are that the outcomes do not overlap, and that their probabilities sum to one. For the example of a demand (probability) distribution, the outcomes are the different potential fuel demand levels. A typical demand distribution might look like the following:

Probability that demand is 1 Step = 0.2

Probability that demand is 2 Steps = 0.5

Probability that demand is 3 Steps = 0.3

This demand distribution tells us the probability that demand in the period will be any one of these three levels. The probability that demand is any other level is zero. Note that the demand levels do not overlap (a demand cannot be both 2 Steps and 3 Steps at the same time, for example), and their probabilities sum to one.

Suppose we knew that we had enough fuel to cover up to 2 Steps of demand. What would be the probability that we would have enough fuel to cover demand this period? It is simply the probability that the demand is 2 Steps plus the probability that the demand is 1 Step, which gives a total probability of 0.7.

A distribution that helps us answerthis kind of question is called a "cumulative distribution.' Instead of telling us "What is the probability that my uncertain demand is equal to a given level?" the cumulative distribution answers the question "What is the probability that my uncertain demand is less than or equal to a given level?"

Constructing a cumulative distribution from a regular probability distribution is not difficult. In fact we began constructing a cumulative distribution in the above example when we found the probability that the demand in the example was less than or equal to 2 Steps. To calculate a cumulative distribution, we simply sum up the probabilities that correspond to levels less than or equal to the one desired.

Constructing the cumulative distribution for the distribution given earlier:

Probability that demand ≤ 1 Step =	Prob(demand is 1 Step) =	 0.2
Probability that demand ≤ 2 Steps	= = =	Prob(demand is 1 Step) + Prob(demand is 2 Steps) 0.2 + 0.5 = 0.7
Probability that demand ≤ 3 Steps	= = =	Prob(demand is 1 Step) + Prob(demand is 2 Steps) + Prob(demand is 3 Steps) 0.2 +0.5 + 0.3 = 1.0

We summarize the results-

level	**Prob(demand ≤ level)**
1	0.2
2	0.7
3	1.0

(Note that the probabilities for a cumulative distribution do not sum to one.)

Plotting these probabilities gives a graph that looks like this:

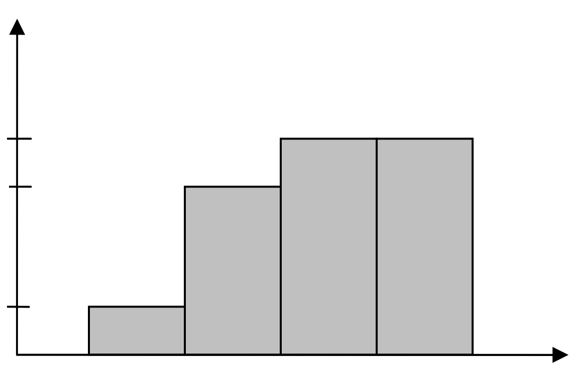

Note that the probability that demand is less than or equal to 4 steps is 1.0; this follows from the fact that demand cannot be more than 3 steps. Likewise, since demand must be at least 1 step, the probability that demand is less than or equal to say, one-half step is 0.

While constructing cumulative distributions is not difficult, it is tedious to make such calculations by hand for distributions that have many levels. Yet most UFIM users find cumulative distributions on inventory easier to interpret and use than regular probability distributions. For these reasons, UFIM calculates cumulative distributions on end-of month inventory and includes them among its outputs. These distributions are found in the files *Normal- Times* results.

CONCEPT 2: THE DIFFERENCE BETWEEN DISTRIBUTIONS FOR BOTH NORMAL AND DISRUPTED PERIODS & NORMAL PERIODS ONLY

UFIM includes two sets of cumulative distributions among its many outputs: the "Cumulative Distribution on Ending Inventory - Normal and Disrupted Periods," and the "Cumulative Distribution on Ending Inventory - Normal Periods Only." Both of these tables list, month by month, cumulative distributions on the amount of inventory (in Steps) left at the end of the month. The difference between the two tables, as their names imply, concerns the type of periods (Normal Only vs. Normal and Disrupted) which are included in the calculation.

A period is called a "disrupted" period if a disruption begins or is in process at the start of that period (month). If no disruption occurs oris in process at the start of a month, the month is called a "normal" period. For the purpose of aggregating data, "Normal and Disrupted periods" refers to all periods (months) that are either normal or disrupted, which is, of course, all periods. "Normal periods only" refers solely to those periods that are "normal", i.e., those periods that proceed without a disruption.

To investigate a given month's ending inventory, without regard to whether a disruption has occurred in that month or not, use the "Normal and Disrupted Periods" table. Suppose, for example, that you want to knowthe likelihood of ending Month 2 (or beginning Month 3) with inventory at 150 or fewer kilounits. Figure 1 shows a portion of the table "Cumulative Distribution on Ending Inventory - Normal and Disrupted Periods" with the desired answer circled; it is found by moving down the left column until we reach the row labeled "150" kilo units and reading the entry under the second month. The chance of having an inventory of 150 kilou flits or fewer at the end of month two is 0.30 (or 30%).

Suppose you wanted to see what effect the disruptions in Month 2 had on this probability distribution. Looking in the second column in the "Cumulative Distribution on Ending Inventory - Normal Periods Only" table, you could read the distribution on ending inventory in the second month assuming that no disruption was in process in the second month. Reading the "Normal Periods Only" table in Figure 2, we see that the probability of ending the second month with inventory less than or equal to 150 kilo units is 0.38 (or 38%), assuming that no disruption has occurred in the second month.

	Month		
Inventory (kUnits)	**1**	**2**	**3 • • •**
0.0	0.00	0.00	0.18
30.0	0.00	0.00	0.54
60.0	0.00	0.02	0.83
90.0	0.00	0.07	0.95
120.0	0.00	0.17	0.99
150.0	0.00		1.00
180.0	0.00	0.42	1.00
210.0	0.00	0.51	1.00
240 • • •	0.00	0.65	1.00

Figure 1
Cumulative Distribution on Ending Inventory - Normal and Disrupted Periods

	Month		
Inventory (kUnits)	**1**	**2**	**3 • • •**
0.0	0.00	0.00	0.18
30.0	0.00	0.00	0.54
60.0	0.00	0.02	0.83
90.0	0.00	0.09	0.95
120.0	0.00	0.22	0.99
150.0	0.00	0.38	1.00
180.0	0.00	0.51	1.00
210.0	0.00	0.61	1.00
240 • • •	0.00	0.73	1.00

Figure 2
Cumulative Distribution on Ending Inventory - Normal Periods Only

A tip is: If you want to compare a cumulative distribution in normal and disruption times with the distribution that would result if there were no disruptions at all, then you should remove the disruptions and run the case again to get the distribution without disruptions. The "Cumulative Distribution on Ending Inventory - Normal Periods Only" table under each month gives cumulative distributions on ending inventory assuming that only that particular month is always normal.

CONCEPT 3: THE DIFFERENCE BETWEEN STOCKOUT AND SHORTAGE PROBABILITIES

Both "stockout" and "shortage" are situations where there is no inventory on hand. The term "stockout" describes any situation where inventory is zero. "Shortage," however, describes only those situations where a demand has been made on inventory when there is none on hand. Thus a shortage is one type of stockout. If a stockout occurs, and no demand is made on the inventory, then no shortage occurs.

UFIM Tips and Traps

The difference between the two can also be expressed in terms of the cumulative distributions discussed above. If we consider a shortage as "negative" inventory (a situation in which replacement power is purchased), then a shortage is a situation where the inventory level is strictly less than zero. A stockout is a situation where the inventory level is less than or equal to zero. So the shortage probability is always less than or equal to the stockout probability (the actual number reported in UFIM) because a shortage probability doesn't include the chance that inventory is exactly equal to zero.

In UFIM, only strict shortages cause replacement costs to be incurred. There actually may be times when it is advantageous to run inventory down to zero, but not to run short of inventory. This would be reported in UFIM as a situation in which the ending month inventory is zero (a stockout) but no shortage occurs.

To illustrate why a high stockout probability might be advantageous, consider the extreme case where demand is certain and where there is no chance of a disruption. In order to minimize holding costs (and hence total costs, because no shortage need occur if demand is known), it would be best to order just enough fuel at the beginning of the month to last the entire month. The ending monthly inventory would be zero for each month, so the overall stockout probability would be 1.0 although no shortage ever occurs!

Reminder: UFIM considers all inventory to be "useable" inventory. If you have inventory that cannot be burned (say, for example, unuseable coal atthe bottom of apile), you should not include this inventory in your starting inventory level. All UFIM outputs will then refer to inventory levels ignoring you run useable inventory, with zero inventory representing zero useable inventory.

A tip is: If you are developing least-cost inventories with UFIM and the UFIM output indicates that stockout probabilities are high, check the UFIM-generated targets in *Summary* results. If the targets are close to zero, then your case is most likely experiencing the situation described in the above paragraph; that is, it may be economically advantageous, given the assumptions of your case, to have high stockout probabilities. If, on the other hand, your targets are high, then the situation is more complex. In this situation UFIM is saying that you would like to have high targets, but that for some reason these targets are not being attained. Most likely, there are constraints to achieving those high targets. There are many reasons why your case maybe "over-constrained" in this way: the normal-times supply might not be enough to satisfy demand, or perhaps the supply system has not "recovered" from severe disruptions of previous months. The key to unlocking the cause behind high stockout probabilities when targets are high is answering the question "Why am I not able to satisfy my target in this period?"

REFERENCES

Ross, Sheldon: A First Course in Probability Theory Macmillan Publishing Company, New York, 1984.

UFIM 2.0 User's Guide

TIPS 'N TRAPS #7: LEAST-COST POLICIES VS. LEAST-RISK POLICIES

PROBLEM

I've done my UFIM analysis, but my targets are far too low to use. My manager, the plant managers, and the PUG will never stand for this!

SOLUTION

Ask yourself why you think the targets are too low. Almost all objections to low targets fall into two categories:

- First, are the objections due to modeling assumptions? Do you overestimate supply capabilities, or under-estimate the shortage costs, frequency, or effects of possible disruptions?
- Second, are the objections due to a reluctance to assume the risk of a shortage? Have you overlooked the fact that UFIM recommends the targets that result in the least total expected cost -- a substantial portion of which may be expected shortage cost?

EXPLANATION

OBJECTIONS DUE TO MODELING ASSUMPTIONS

Frequently in UFIM applications, analysts present UFIM targets and the affected parties declare them to be infeasible. Plant managers say "We can't build up that fast" or the coal buyers say "We can't buy that much. Experience has shown that when estimating maximum supply, people will often give estimates for normal times unloading capacity, coal buying, or other factors that represent the maximum achievable during an extraordinary month -- which is quite different from what can be repeated month after month. When you consult others for UFIM inputs, be sure to ask them to make estimates for any normal month.

A signal that a capacity estimate may be over-optimistic is when it is calculated rather than looked up from historical data or managerial experience. Calculated capacity estimates frequently sound like this: "The equipment can handle one truck every fifteen minutes, which makes four trucks an hour, thirty-two cars a day, and 704 cars a month, bringing us to a grand total of 2112 tons per month." This type of estimation often disguises a strong assumption of a perfect no surprises" world. If someone gives you estimates like this, ask to see historical data on how many trucks are usually unloaded in a month.

Estimates for disruption months will probably be different, and may well include an expectation of some "extraordinary" effort. Try to make sure that the estimates you use are reasonable for the kind of disruption you are modeling.

In addition to over-estimating supply capabilities, many analysts under-estimate shortage costs. Shortage costs that are too low can lead to targets that are also too low. Refer to the new UFIM 3.0 User's Manual and other Tips & Traps for more information on modeling shortage costs.

UFIM Tips and Traps

Another common mistake is modeling disruptions too conservatively. If someone says, "This target is too low -- what if we lose all supply for three months while demand is extra high?" don't dismiss it by saying it's not worth modeling because the event is very unlikely. In many cases, it is perfectly appropriate to include a very severe disruption with a very small probability of occurring.

OBJECTIONS DUE TO RISK OF A SHORTAGE

Once you are confident that you have captured the situation your utility really faces, don't change your case just because someone says the targets are still "too low." Instead, ask if the targets are uncomfortable because they involve a risk of running out of fuel. Remember, UFIM is calculating the least-expected-cost policy, which is the policy that best balances the costs of fully protecting against a shortage against the expected costs of incurring that shortage. This means that UFIM will choose a slightly risky policy over a less risky policy if the former has a lower expected total cost. UFIM disregards internal politics, regulatory concerns, and fuel manager's ulcers. This is not to say that these other concerns aren't valid ones. What UFIM does is give you a way to put a price tag on these other concerns.

To see how, consider the following sample outputs.

Expected Cost Per Target Policy ($ Million)
One burn day = 22 KUnits

Target Policy (kUnits)	Total Cost	Annual Total Cost	Annual Burn Cost	Annual Holding Cost	Annual Oppty. Cost	Annual Shortage Cost
192	1237.06	86.59	83.02	1.09	-0.02	2.50
252	1226.94	85.89	83.03	1.43	-0.02	1.45
312	1221.43	85.50	83.03	1.77	-0.01	0.72
372	1220.32	85.42	83.03	2.11	0.00	0.28
432	1222.51	85.58	83.03	2.46	0.01	0.08
492	1226.54	85.86	83.03	2.80	0.01	0.02
552	1231.24	86.19	83.02	3.15	0.01	0.01
612	1236.18	86.53	83.02	3.49	0.02	0.00

We can see that the least cost policy is to hold seventeen (372/22) days of inventory with an annual expected cost of $85,420,000. If you follow this policy, however, you can expect to pay $280,000 in shortage costs on average each year, which means that you can expect to run out of inventory sometimes. Suppose, however, that your utility is concerned because the PUC frowns on running out of fuel, not to mention how bad the newspapers can make your utility look if it happens. You can see that the policy of holding twenty-eight days of inventory has zero shortage costs. This implies that twenty-eight days is the lowest target that involves virtually no fuel shortage risk. By comparing total costs, we can see that the average annualized cost of totally eliminating risk of a shortage is $1,110,000 ($86,530,000 minus $85,420,000) per year.

UFIM doesn't have an opinion on which policy you should use -- that's up to you, your management, and your PUC. UFIM does, however, show how much it costs to be ultraconservative; it explicitly estimates what you must pay to avoid any risk of running out of fuel. This can be valuable information -- use it to your advantage!

REFERENCES

Utility Fuel Inventory Model: Basic Concepts (New User's Manual), Chapters 1 and 2.

UFIM User's Notebook, Tips & Traps #1: The Critical Ratio Test.

UFIM User's Notebook, Tips & Traps #3: Determining Replacement Power Costs.

UFIM Tips and Traps

TIPS 'N TRAPS #8: SUPPLY CURVES

PROBLEM:

We buy a fixed amount of fuel on a long-term contract and a variable amount on the spot market. How can I model this fuel supply situation in UFIM?

EXAMPLE:

I am running a non-seasonal case. Our long-term contract requires us to buy 60 KUnits of fuel each month at a yearly cost of $21.6 million ($30/unit/mo.) with severe penalties for not taking delivery. The spot market conditions are such that we can reasonably expect to buy up to 180 KUnits if needed at $35/unit. A third source of fuel is is available at $40/unit.

The demand distribution is as follows:

Demand	Probability
3 Steps	0.2
5 Steps	0.5
8 Steps	0.3

Maximum fuel order is 300 KUnits. Step size is 30 KUnits.

SOLUTION:

In order to specify unit fuel prices that vary with the order amount, use UFIM's supply curve option. There are three important steps to using supply curves:

- Entering supply curves,
- Checking price assumptions,
- Interpreting output.

EXPLANATION:

STEP 1: ENTERING SUPPLY CURVES

Entering supply curves into UFIM is straightforward. The data UFIM requires for each supply curve are total purchase costs for fuel at the minimum and maximum order amounts, and any other "break" points where fuel prices change. A good way to obtain this information is to create a table of order amount versus total cost up to your maximum order. We specify such a table for the example above. Due to a severe penalty for not taking delivery on the contract, the contract fuel will be the first 60 KUnits purchased each month. The next 180 KUnits of fuel, if needed, can be bought on the spot market at $35/unit. The last 60 KUnits (up to the 300 KU nit maximum) will be bought on the spot market at the higher $40/unit price. The "break" points where the fuel prices change in this example occur at 60 KU flits, 240 (=60+180) KUnits and 300 (=60+180+60) KUnits. Calculating the total fuel costs at each break point yields:

Order Amount	Cost (in 1000's)
60 KUnits	$1,800 = $21,600 12
240 KUnits	$8,100 = $1,800 + $35*(24060)
300 KUnits	$10,500 = $8,100 + $40*(300240)

Obtaining these purchase costs allows usto draw a picture of the entire supply curve. Because fuel costs are constant between the order amounts, we draw the supply curve by simply connecting these points on a graph of order amount versus total cost:

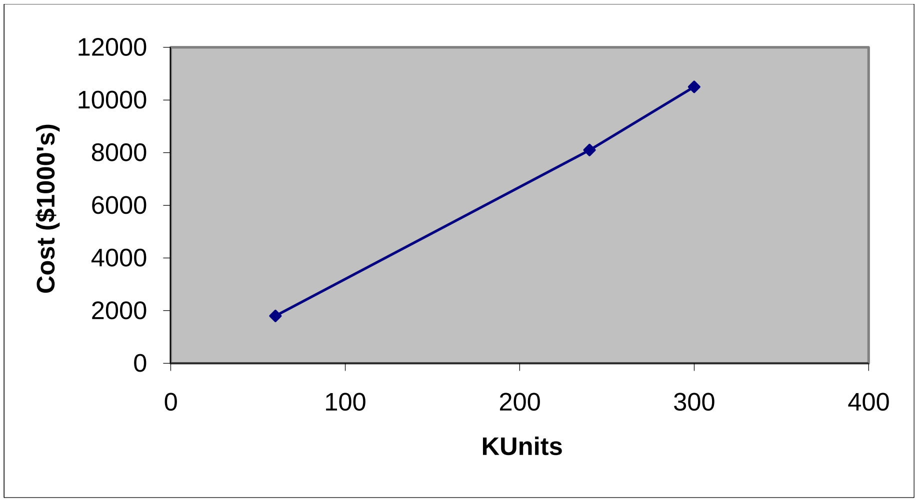

This graph (of a type called "piecewise linear" because it looks like a line broken into pieces) is the graph that UFIM uses to calculate fuel prices when supply curves are used. This curve shows the total price for any purchase amount from the minimum up to the maximum order. All you need to do is enter th epairs of order amount (in Steps, rather than KUnits) and costs, like those listed above, and UFIM will automatically "draw" its own lines. For the example above, the three pairs of order amounts and costs would be:

Order Amount (Steps / Month)	Cost ($1000)
2 Steps	1800
8 Steps	8100
10 Steps	10500

UFIM allows you to enter up to 5 "break" points for your supply curve (including the minimum and maximum orders), allowing you to model up to five separate marginal fuel prices. The points must be entered in the same order in which they are graphed, left to right.

A tip is: In this example, the "break" point representing the minimum order amount (2 Steps) coincides with the "break" point where the first price change takes place (also 2 Steps). If these do not coincide, you must enter them separately. For example, if the contract above required a purchase of at least 2 Steps, but allowed an optional additional purchase of 3 Steps at $30! unit, you would need to enter one pair of points forthe minimum 2 Step purchase, and another pair of points for a 5-Step purchase (where the price changes for the next step of fuel). Similarly, if the contract had no minimum required purchase, then you must include a point in the supply curve that corresponds to a minimum order of 0 Steps. The cost of such an order would be $0.

Supply curves area very flexible way to enter contract provisions and spot fuel prices. You can use supply curves to describe minimum and maximum orders, underlift penalties, fixed costs, quantity discounts, and other fuel supply characteristics. Note also that supply curves can be used to describe situations with increasing marginal fuel costs (such as our example), decreasing marginal fuel costs, or some combination of the two.

A tip is: If there is a fixed cost element to your fuel purchases, then you can indicate this by adding the fixed cost to each of the fuel order points listed. Note, however, that the fixed cost will only be incurred if an order is placed. (Ordering 0 KUnits is the same as not ordering any fuel.)

STEP 2: CHECKING PRICE ASSUMPTIONS

UFIM makes certain assumptions about fuel prices that the user may not violate, regardless of whether fixed fuel prices or supply curves are used. They are:

- .All fuel prices must be positive.
- Fuel price must always be less than average replacement power cost, whether in normal or disrupted times. (Otherwise the shortage cost, defined as replacement power cost minus fuel price, could be negative.)
- Fuel price during a warning or a disruption must be greater than or equal to the normal times fuel price. (Otherwise there would be incentive to enter a disruption, which is always considered a "negative" event.)

These conditions are easy to check with fixed fuel prices. Supply curves, however, imply that there are no fixed fuel prices. In order to check the latter two conditions in the presence of supply curves, UFIM calculates a "representative" fuel price for each supply curve to make the necessary comparisons. The single figure that UFIM uses to represent a supply curve is the average price of fuel required to meet the expected demand. This single figure takes into account, through the expected demand, where on the supply curve you order from on average, and hence the per-unit price of fuel on average.

UFIM Tips and Traps

We can restate the three assumptions UFIM makes, as they apply to supply curves during normal times:

- The marginal price of fuel must always be positive, i.e., it must always cost more money to buy more fuel. (Note, however, that it is not necessary for marginal fuel prices to be increasing.)
- In normal times, the average price of fuel at the expected demand must always be less than average replacement power cost. During disruptions, the fuel price must always be less than the average replacement power cost.
- The fuel price during warning or disruption must be greater than or equal to the normal times average price of fuel required to meet the expected demand.

If any of these assumptions is violated in a case, UFIM will terminate the run with a fatal error.

The representative fuel price for the supply curve in the example can be calculated as follows. The expected demand is 3(.2) + 5(.5) + 8(.3) = 5.5 Steps * 30 KUnits/Step = 165 KUnits.

According to the supply curve, the price of 165 KUnits of fuel is $1800 + (1 6560)*$35 = $5.475

million. The average fuel price to purchase the expected demand is then $33.1 9/unit (rounding up to the next cent).

Any period described by this supply curve must have an average replacement power cost in normal times of at least $33.19/unit, and a fuel price during warning and disruptions of at least $33.1 9/unit.

STEP 3: INTERPRETING OUTPUT

The key point to understand about supply curves is that they often produce order policies that are not equivalent to strict target policies. The justification of this statement is mathematical, but can be loosely understood as follows.

UFIM holds inventory to protect against future disruptions, supply uncertainties and demand uncertainties. During a particular period, UFIM knows what to expect "on average" in terms of these future uncertainties. A target policy addresses this average future by giving a single figure of inventory to "shoot" for each month. Given a monthly fuel price, UFIM purchases fuel until it reaches an inventory level where the next unit of fuel costs more to buy and hold than that unit saves in expected avoided shortage costs in the present or a future period. This level is usually called the period's target, and it depends upon the future uncertainties and also the fuel price in the period.

The key to remember is that if the fuel price in the period is constant, then this "target" inventory level will be the same no matter what inventory level you have at the beginning of the period. If, on the other hand, the fuel price varies during the period, then the inventory level that UFIM shoots for does depend upon the inventory level at the beginning of the period. This is because

the purchase price of an additional unit of fuel, and hence whether or not it is advantageous to purchase another unit, depends upon how much fuel you have already purchased in the period. Thus if you have increasing marginal fuel costs, as we do in our example, it may be cost-effective to bring your inventory up to 360 KUnits from 120 KUnits, but not from 90 KUnits. A strict target policy, barring ordering constraints, would have a 360 KUnit target independent of the initial period inventory.

Typically, an order policy resulting from supply curves will look like a target policy in some places, but not in others. Consider the output from a case that was run using the supply curve in the example, with the exception that no minimum purchase was required (i.e., a point was added with 0 KUnits costing $0), and also run with a constant fuel price of $30/unit. (See Table 1. on next page).

The constant fuel price yields a true target policy, in that ignoring the order constraint of 300 KUnits, the sum of inventory and orders is a constant amount (360 KUnits).

The optimal order policy forthe case with supply curves exhibits a constant amount of inventory plus orders, but only for beginning inventory levels between 120 and 200 KUnits. This is because the supply curve exhibits increasing marginal costs; the smaller the order placed, the cheaper the fuel price per unit purchased. The price structure provides incentive to purchase fuel in as small amounts as possible. This explains why the order policy suggests purchasing 60 KUnits of fuel even when inventory is at a high level. It is better to purchase fuel at the low per-unit price and hold it for a period or two, than to have to purchase the extra two units in a later month when your order size may be larger and hence your per-unit fuel price higher.

Using similar reasoning, we can explain why the orders flatten at 240 KUnits at the lower inventory levels. This is because an extra unit of fuel purchased would cost the maximum price ($40/unit) and it is apparently worth taking the increased risk of running short of rather than paying an extra $5/unit of fuel. Thus, in the example, it is always worth placing another 60-KUnit order for fuel, but it is never worth placing an order for more than 240 KUnits.

Notice in both cases that the orders "flatten" precisely at the break points in the fuel prices.

A tip is: Because the order policy resulting from the use of supply curves is usually not a strict target policy, UFIM will give an error message warning the user that the target policy given in the *Summary* results is a target policy approximating, but not necessarily equivalent to, the optimal order policy given in *Normal – Times* results. In order to evaluate how well the target policy approximates the optimal order policy in costs, you can compare the costs of the order policy from Run Level 2 with the costs of the optimal target policy from the Simulation (Run Level 4). Additionally, using the UFIM Enhancements, you can evaluate other user-specified order policies in which you might be interested.

REFERENCES

Utility Fuel Inventory Model: Basic Concepts (New User's Manual), Chapters 1 and 2.

UFIM User's Notebook, Tips & Traps #3: Determining Replacement Power Costs.

UFIM Tips and Traps

INVENTORY KUNITS	Fixed Price	Supply Curve
0	300.00	240.00
30	300.00	240.00
60	300.00	240.00
90	270.00	240.00
120	240.00	240.00
150	210.00	210.00
180	180.00	180.00
210	150.00	150.00
240	120.00	120.00
270	90.00	90.00
300	60.00	60.00
330	30.00	60.00
360	0.00	60.00
390	0.00	60.00
420	0.00	60.00
450	0.00	60.00
480	0.00	60.00
510	0.00	60.00
540	0.00	60.00
570	0.00	60.00
600	**0.00**	60.00

Table 1
LEAST-COST ORDER AMOUNTS (IN KILOUNITS)

UFIM Tips and Traps

TIPS 'N TRAPS #9: REPLACEMENT POWER COSTS REVISITED

PROBLEM:

I enter a single curve for replacement power costs during a disruption. How does UFIM calculate burn reduction costs and shortage costs from this one curve? And how does this differ from the way UFIM calculates shortage costs during normal times?

EXAMPLE:

During the disruption and during normal times, the following replacement fuel alternatives are

Source	% of time used for replacement power	Cost/MWh
Oil	First 60%	$15
Combustion Turbine	Next 20%	$20
Purchased Power	Last 20%	$25

My delivered price of fuel is $35/ton; heat content is 25 MBtu/ton; heat rate is 10,000 Btu/ KWh; and step size is 20 Ktons.

SOLUTION:

Whether or not burn reduction is allowed in a period determines how UFIM uses the replacement power curve to calculate burn reduction and shortage costs.

Burn reduction is the planned purchase of replacement power. Through burn reduction, a utility strategically uses its limited fuel stocks to decrease replacement power costs. In UFIM, burn reduction is allowed during disruptions whenever you enter a (sub)period with some inventory.

When burn reduction is not allowed, shortage costs are calculated differently. Times when burn reduction is unavailable include:

- Normal times
- Disruptions for which the user specifies that the burn reduction option is not allowed
- Whenever the inventory level hits zero during disruptions that allow burn reduction. (In this case there is no inventory that can be used strategically.)

EXPLANATION:

A replacement power curve is entered into UFIM as a series of points. These points define the price of replacement power for various percentage reductions in plant burn. As discussed in Tips and Traps #3, "Determining Replacement Power Costs," this curve can be entered as a step function, where jumps represent the addition of higher cost sources of replacement power, or as a "smooth" curve representing increasing system marginal costs. The replacement power costs given above, for example, describe the following replacement power curve:

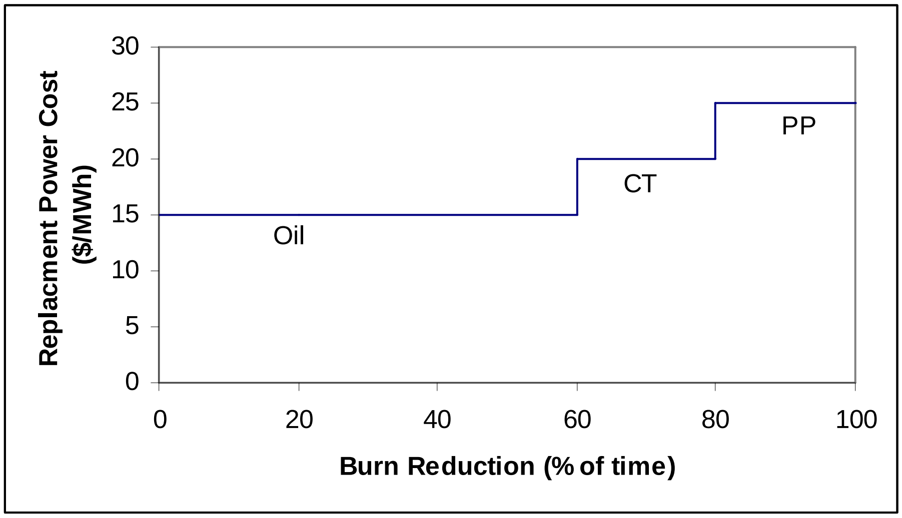

Tips & Traps #3 explains in more detail how to specify replacement power curves. This Tip & Trap explains how UFIM uses these curves to calculate the various components of replacement power costs.

REPLACEMENT POWER COSTS WHEN BURN REDUCTION IS ALLOWED

Burn reduction is a planned purchase of replacement power, and is available only if a plant has some fuel currently in inventory. When burn reduction is used, UFIM assumes that a utility will try to save as much money as possible by strategically burning scarce fuel at a plant during times when replacement power is most expensive, and by purchasing replacement power during times when that power is cheapest. If, for example, UFIM decides that the best disruption management policy for a given inventory level is to reduce burn at the plant by 75%, then the plant will burn fuel during the 25% of the time when replacement power is most expensive, and purchase power during the 75% of the time when the replacement power is cheapest. The average cost of the replacement power is determined by the cheapest 75% of the replacement power curve.

The percentage of burn reduction is calculated as the amount of burn reduction (in steps) divided by the monthly demand outcome. For example, suppose that during a supply disruption there are 2 steps of fuel in inventory, demand is 8 steps, and there are no fuel deliveries. If UFIM decides to reduce burn 6 steps, using only 2 steps of current fuel inventory, then the percentage reduction is

$$Percentage\ reduction = \frac{reduction\ amount}{demand\ \ outcome} = \frac{6\ steps}{8\ steps} = 0.75 = 75\%$$

The cost for this burn reduction, per megawatt-hour, would be calculated as the average cost of the cheapest 75% of replacement power. This cost is derived from the shaded area on the graph:

UFIM Tips and Traps

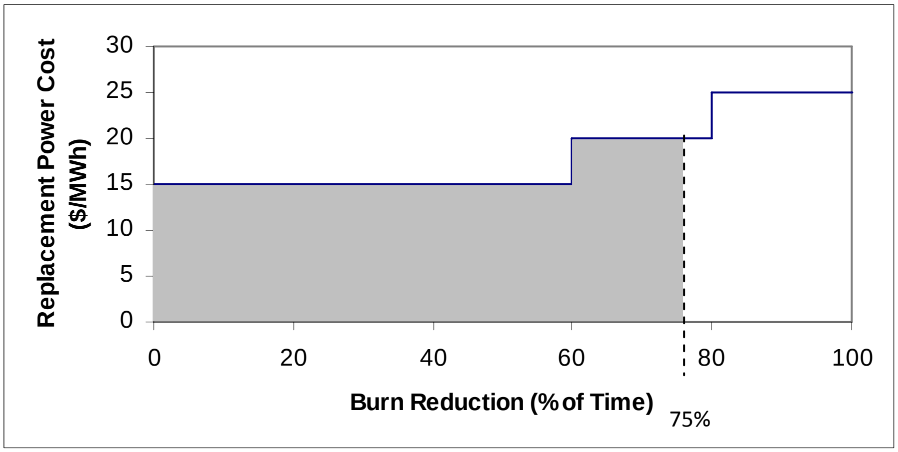

As the figure indicates, oil constitutes 60% of the fuel burn that month, the combustion turbine 15%, and current inventory stocks the other 25%. The average cost of the replacement power (oil and combustion turbine) is the weighted sum of the cost of these two replacement sources, divided by the fraction reduction that they represent:

$$average\ \cos t\ of\ replacement\ power = \frac{\$15(0.6)+\$20(0.15)}{0.75} = \$16/MWh$$

This replacement power cost is an actual cost; the utility pays, on average, $16 for each megawatt of replacement power used. But how much has it really cost the utility to run out of fuel? If it had not run out of fuel, it would still have had to pay $35/ton for fuel. How much worse is purchasing replacement power than using regular fuel stocks?

In order to make the extra cost of reducing burn or running short of fuel more explicit, UFIM divides this replacement power cost into two categories, fuel burn costs and shortage costs. Fuel burn costs are the costs the utility would have to pay to provide the same power using its regular fuel source. The shortage cost is the replacement power cost minus the fuel burn cost; it is a measure of how much extra replacement power really costs, as an alternative to usual stocks.

Note: Burn and shortage costs, as components of replacement power costs, do not exist in reality. That is, the utility does not separately pay burn costs and shortage costs when it purchases replacement power. These categories are only created in order to let you see what part of replacement power costs is due to plant demand and is unavoidable (fuel burn costs) and what part is due to plant management and is therefore potentially avoidable (shortage costs).

Note: Burn costs also exist when replacement power is not being purchased. UFIM aggregates burn costs overall periods and reports them in *Summary* and *Normal – Times* results.

In the above example, the order price of fuel during the disruption is $35/ton. The total cost of the six steps of burn reduction is

$$\textit{Total} \text{ cost } \textit{of replacement power} =$$
$$6\ \textit{steps}(20\textit{Ktons}/\textit{Step})(25\textit{MBtu}/\textit{ton})(1\textit{KWh}/10{,}000\textit{Btu})\$16\textit{MWh} = \$4.8\ \textit{million}$$

This is separated into burn costs and shortage costs as follows:

Burn Costs = 6 Steps * (20 Ktons/Step) * ($35/ton) = $4.2 million

Shortage Costs = $4.8 million - $4.2 million = $0.6 million = $600,000

Note that 8 steps of demand are satisfied this month from both coal and replacement power. Hence burn costs for the entire month are:

Burn Costs for entire month = 8 Steps *20 Ktons/Step * $35/ton = $5.6 million

REPLACEMENT POWER COST WHEN BURN REDUCTION NOT ALLOWED

UFIM assumes that when burn reduction is unavailable, the utility has no option to strategically purchase replacement power; when it runs out of fuel, the utility must purchase power immediately and has no chance to wait for times when replacement power is cheapest. Accordingly, the replacement power cost when burn reduction is unavailable is determined by taking into account the entire replacement power curve.

The average replacement power cost when burn reduction is unavailable is the weighted cost of all replacement power options:

Average replacement power cost = ($1 5*0.6) + ($20*0.2) + ($25*0.2) = $1 8/MWh

Replacement power purchased in the month will cost $18 for each megawatt hour obtained.

Again, as in the example above, this replacement cost will be split up into two parts for the accounting, fuel purchase costs and shortage costs.

REFERENCES

Utility Fuel Inventory Model: Basic Concepts (New User's Manual), Chapters 1 and 2.

UFIM User's Notebook, Tips & Traps #3: Determining Replacement Power Costs.

UFIM Tips and Traps

TIPS 'N TRAPS #10: DISRUPTION/WARNING ARRIVAL RATES

PROBLEM

The case I am running has only two disruptions, but when I run it, UFIM stops with a fatal error and tells me that the sum of my disruption/warning arrival rates is greater than 1. What is wrong with my case (or UFIM)? How does UFIM calculate disruption and warning arrival rates?

EXAMPLE

The first disruption is a frozen river disruption. It happens about once every three years and only in February. We can tell it is coming by unusually cold weather starting in January, but half the time this advanced warning is wrong. The data for my disruption are:

Years between disruptions:	3
Warning Length:	1 Month
Warning Accuracy:	0.5 (only 50% of the warnings are correct)
Disruption Length:	1 Month

My other disruption is an unloader breakdown, which tends to happen in December or January with equal chance and lasts about a month. This disruption happens once a year. It comes without a warning.

Years between disruptions:	1
Disruption Length:	1 Month

I have chosen not to model the possibility that these disruptions happen simultaneously.

SOLUTION

UFIM calculates the arrival rate of each disruption by dividing the relative monthly probability of occurrence by the length between disruptions. However, this calculation is adjusted to take into account both the length and accuracy of possible warnings.

EXPLANATION

UFIM thinks of disruptions in terms of "events." Each combination of a disruption and a period in which that disruption can start forms a different event. The unloader break-down described above can happen in either December or January. Accordingly, UFIM treats the unloader breakdown in December as one event, and the unloader breakdown in January as a separate event.

UFIM calculates an arrival rate for each event to ensure that you do not specify the possibility of more than one disruption, on average occurring in a single month. Arrival rates are associated with the beginning of the disruption events. When calculating arrival rates, UFIM first determines in what period each of your disruption events begins. For disruptions that have no warning, the beginning of the event is the start of the actual disruption. If the disruption has a warning, UFIM counts backwards from the start of the disruption to include the warning.

UFIM Tips and Traps

A tip is: Although we usually think of a warning as a hint during normal times that a disruption is about to occur, UFIM treats the warning as part of the disruption event. During a real life warning (for example, a one month warning that a strike might be imminent) you may wish to take advantage of certain fuel procurement options that you would not ordi¬narily consider during normal times. UFIM gives you the option of specifying supply data for the warning that differ from the normal times supply data. Also, the policies your utility follows during warnings can radically affect the severity of disruptions. Consequently, warnings should be thought of as events closely linked to disruptions, but separate from normal times. UFIM distinguishes warnings from normal times by including them in disruption events.

The disruptions listed in the above example describe three separate events. The unloader breakdown is two separate events because the breakdown can occur in December or January. These disruption events begin in December and January respectively because they have no warning.

The frozen river accounts for a single disruption event that occurs in January. This event occurs in January because the actual disruption (which starts in February) is always preceded by a one month warning. Thus the warning (and the event) always begins in January.

UFIM then calculates the rate at which each of these events occurs. For an event that has no warning, the rate is calculated by dividing the relatively monthly probability by the years between disruptions. In the example of the unloader outage above the disruption arrival rate in the month of January is:

$$\text{Arrival rate} = \frac{\text{Relative monthly probability}}{\text{Years between disruptions}} = \frac{0.5}{1} = 0.5$$

The January arrival rate for this disruption is 0.5, meaning that if we looked at all of the months of January over a period of time, the disruption would happen in about half of those Januaries. Another way to think of the arrival rate is as a chance that the disruption will happen in a given January. For this example, there is a 50% chance that an unloader outage happens in any particular January.

For disruptions that have warnings, UFIM adjusts the arrival rate to take into account the accuracy of the warning. Consider the frozen river disruption. This disruption happens in February, and only February, on average once every three years. The arrival rate of the actual disruption is:

$$\text{Arrival rate} = \frac{\text{Relative monthly probability}}{\text{Years between disruptions}} = \frac{1}{3} = 0.33$$

This says that in 33% of our Februaries, we will experience a frozen river disrup¬tion. But each of these is preceded by a warning of length one month. But because our warnings are only 50% accurate, the warning event will occur twice as often as the actual disruption. Half of the January warning months will be followed by a frozen river disruption and half will be false alarms. While there is a 33% chance that February will have a disruption, there is a 33% 0.5 = 66% chance that January will begin a warning event. In general,

UFIM Tips and Traps

$$\text{Arrival rate with warning} = \frac{\text{Relative monthly probability}}{\text{Years between disruptions x Warning accuracy}} = \frac{1}{3} = 0.33$$

Notice that the previous formula given, for disruptions without a warning, is the same as this formula, except that the warning accuracy is assumed to be 1.

Making these calculations reveals the problem with entering both the frozen river and unloader outage disruptions. Calculations reveal that 50% of our Januaries begin with unloader disruptions, and 66% of our Januaries begin with warnings of river freezings that may (or may not) turn into actual river freezings. If we add that up we get over 100%! This does not make sense because it says that on average more than one event happens every January. UFIM only allows one disruption event to happen in a month, and so to UFIM, this situation cannot occur.

A tip is: Having the sum of your arrival rates greater than one is equivalent to saying that "normal times" never exists for that month, because you will always be in a warning or disruption during that month. This situation generally implies that you should rethink your concept of "normal times." If a given warning arrives often during one of these months, you can enter warning data as normal times data (in the appropriate month) and remove the warning for the disruption. This will decrease the sum of your arrival rates. Alternatively, you can explicitly model the possibility that multiple disruptions occur simultaneously. For more information about modeling joint disruptions, refer to Tips & Traps #5, "Modeling Joint Disruptions"

A tip is: As this example shows, the use of warnings (and in particular inaccurate warnings) can drive your arrival rates very high. In fact, if your disruption is frequent enough, and your warning inaccurate enough, you can get the fatal error message with a single disruption. This is because the more inaccurate the warning, the more warning events will occur for every disruption that occurs. In general, if warnings are very inaccurate, or happen very often, you may wish to specify the warning data as normal times data and model the disruption without a warning.

REFERENCES

Utility Fuel Inventory Model: Basic Concepts.

UFIM User's Notebook, Tips & Traps #5: Modeling Joint Disruptions.

UFIM Tips and Traps

TIPS 'N TRAPS #11: OUTAGE COSTS

PROBLEM

I am currently running a case that includes the possibility of a supply shortfall so severe that we might have to cut electric service to some customers. How do we determine the "cost" of such an action? Are there any references on estimating outage costs?

EXPLANATION AND SOLUTION

UFIM finds least cost inventory strategies by balancing the costs of holding inventory with the risks of running short. In order for UFIM to make decisions consistent with your utility's operating environment, it must know your utility's "true" costs of replacement power. Determining the cost for typical replacement power options such as combustion turbine or purchased power can be (though not always) a straightforward calculation or a lookup on a system marginal cost curve. In the extreme case where all backup options have been exhausted, and your utility cannot supply the necessary power, the costs of the replacement power option "curtail delivery of electricity" are not necessarily evident.

Previous Tips and Traps (numbers 3 and 9) have explained how replacement power curves are defined and used by UFIM. The subject of this Tips and Traps is not how to use the replacement power curves, but rather what replacement power cost to use to represent the possibility of a customer outage.

BACKGROUND

A significant amount of research has been devoted to the subject of customers' outage costs, and several state PUC's have mandated large scale outage cost surveys to support resource planning at utilities. Unfortunately, the cost estimates from surveys and other research efforts have varied widely and there are currently no accepted "correct" outage costs or procedures for estimating outage costs.

A major cause of the difficulty in estimating outage costs is the lack of a substantial market in which outages are traded   for example, an outage insurance market in which electric consumers could buy insurance against outages. In such a market, the level of insurance purchased would directly indicate the costs of outages. The lack of market data has led to the use of varied approaches for estimating outage costs and varied estimates. Below we will briefly describe some of these approaches to estimation and the data they have produced.

We should note that interruptible / curtailable service and real time pricing are markets for partial outages. Some information from these markets has been analyzed and will be described below. Much information from these markets is unfortunately of limited value because the markets have few participants and usually only indicate the costs of partial outages. However, if such programs grow, the determination of outage costs may be greatly simplified.

APPROACHES TO OUTAGE COST ESTIMATION

UFIM Tips and Traps

Perhaps the single best source for information on outage costs is the 1989 EPRI report by Laurits R. Christensen Associates, "Customer Demand for Service Reliability" (EPRI, 1989). This report identifies four methods of estimating outage costs:

- Proxy methods. These methods examine prices in markets that are assumed to represent or reflect the cost of outages. Examples of cost proxies include: the cost of back up generators, the ratio of output measured in dollars to the electricity used to produce the output measured in kilowatt hours, the value of production in the home, the wage rate, and the price of electricity. The first two are usually used to measure the cost of outages in commercial and industrial uses, the next two are used to measure the cost of residential outages, and price of electricity is used to measure the cost of all outages. Outage cost estimates based on these methods have ranged from $1.27/kWh to $6.14/kWh (EPRI, 1989, p.2 4).
- Reliability demand models. In this approach, the demand for electricity is compared across areas that differ significantly in their electric service reliability. It is expected that, given similar prices, less electricity will be purchased in areas of poor service quality than in areas of high service quality. The differences in the demand for electricity across these regions will indicate the cost of outages. However, due to uniform levels of high service quality, this method is impractical in the U.S.; it has not proved satisfactory even in countries with greater service quality variations.
- Consumer surplus measures. In this method, the value of outages is determined based on those few markets, such as real time pricing and interruptible/curtailable services, in which service reliability is bought and sold. This method holds promise but has been limited by the fact that few customers choose such service options, the data on these programs is of poor quality, and that customers experience partial rather than full outages. An analysis of Orange and Rockland's Peak Activated Rate suggested an outage cost of $0.1 05/kWh for participants (EPRI, 1989, p. 3 22). A recent preliminary analysis of interruptible/curtailable programs suggest outage costs on the order of $3.57/kWh (Caves, et al, 1991, p. 27).
- Survey methods. Survey methods have been the most popular approach to outage cost determination, and therefore we will describe these methods in additional detail.

EPRI (1989) classifies survey methods into three groups:

- Direct Cost. An outage or curtailment is described and the customer is asked to state the cost of the outage or curtailment.
- Contingent Valuation. An outage or curtailment is described and the customer is asked to state how much he/she would pay to avoid the event (Willingness to Pay or WTP) or how much he/she would have to be paid to accept the event (Willingness to Accept or WTA).
- Contingent Ranking Method. Respondents are asked to rank or choose from a list of described outages, curtailments, or interruptible/curtailable rates.

Examples of the application of these survey methods can be found in the Energy Journal (1988). As with the other methods, survey methods have produced a wide range of outage cost estimates. For residential customers, outage costs range from $0.09/kWh to $14.61 /kWh (EPRI, 1989, p. 2 15); for industrial customers, outage costs range from $7.76/kWh to $22.46/kWh (EPRI, 1989, p.2 16); and for commercial customers, outage costs range from $1.54/kWh to $956.60/kWh (EPRI, 1989, p.2 17) (all in 1986 dollars).

MAJOR SOURCES OF INFORMATION

Four publications provide an excellent review of outage costs estimation methods and data.

1) The first (EPRI, 1986) is the proceedings from a number of seminars on the value of service reliability to customers. Papers include reviews of approaches to measuring the value of electric reliability and reports on individual studies of the value of reliability.
2) The second (Energy Journal, 1988) is a special edition of The Energy Journal that focuses on electricity reliability. Papers discuss the use of outage cost information in planning for electric systems and, also, present specific estimates of the value of reliability.
3) The third publication, which has already been noted, (EPRI, 1989) is a synthesis of the outage cost literature. This publication provides an annotated bibliography and assesses the data and methods presently available for estimating the market for reliability differentiated service. Three tables at the end of this report (EPRI, 1989, p. 4 2 to 4 4) provide and excellent summary of outage cost estimates.
4) The fourth publication is similar to the third in that it provides and updated bibliography and assessment of the data and methods. However in this case it is for estimating the value of service reliability to customers. This publication also provides a simple model for estimating the value of reliability to customers.

REFERENCES

Caves, D.W.; Herriges, J.A.; and Windle, R.J. (1991). "The Cost of Power Interruptions in the Industrial Sector: Estimates Derived from Interruptible Service Programs," Working Paper. To appear in Land Economics in 1992.

Energy Journal (1988). Vol. 9, Special Electricity Reliability Issue.

EPRI – Electric Power Research Institute (2000)., Customer Needs for Electric Power Reliability and Power Quality: White Paper. Feinstein, C.D., Morris, P.A., Downs, C., Report 1000428, October 2000..

EPRI – Electric Power Research Institute (1986). "The Value of Service Reliability to Consumers," prepared by Criterion Inc., EPRI EA 4494, Project 1104 6, May 1986.

EPRI – Electric Power Research Institute (1989). "Customer Demand for Service Reliability, A Synthesis of the Outage Costs Literature," prepared by Laurits R. Christensen Associates, EPRI P 6510, Project 2801 1, September 1989.

UFIM User's Notebook, Tips & Traps #3, "Determining Replacement Power Costs."

UFIM User's Notebook, Tips & Traps #9, "Replacement Power Costs Revisited."

UFIM Tips and Traps

TIPS 'N TRAPS #12: TARGET VS. ORDER POLICIES

PROBLEM

UFIM gives me two normal-times least-cost policies, a least cost target policy (in *Summary* results) and a least-cost order policy (in *Normal – Times* results). What is the difference between them? Which one should I follow?

SOLUTION

The "real" policy that UFIM calculates is the least-cost order policy. The least-cost order policy tells you how much to order each month once you know the starting inventory level for that month. Because order policies are somewhat lengthy (an order policy for a single month would consist of one number for each possible starting inventory level), UFIM converts the order policy into a simpler type of policy, a target policy. A target policy tells you to order up to a certain inventory level, no matter what inventory level you begin the month with.

In many cases, the target policy is exactly equivalent to the least cost order policy; both policies tell you to do the same thing. Other times, however, the target policy is an approximation of the order policy. In this case, the target policy is useful to get an idea of how much fuel UFIM suggests you have at the end of the period, though the order policy is a more accurate description of what you should do. It is the order policy that you usually use to make ordering decisions.

EXPLANATION

UFIM FINDS THE LEAST-COST ORDER POLICY

During normal times, a utility can control inventory policy only through the amount of fuel ordered. Consequently, UFIM was designed to answer the question "how much fuel should we order"? When UFIM is trying to figure out the least-cost inventory strategy, it takes each possible period and each possible inventory level at the start of that period, and searches among all possible order amounts to find the order amount with the least expected cost.

The result is a least-cost order policy, or, in other words, a list of suggested order amounts based on what inventory you have at the start of the period.

An order policy is explicit, for any starting inventory level, it tells you what to order. However, the fact that it is explicit also makes it cumbersome. If we assume that inventory level at the beginning of the month can be any one of 30 possible levels, then communicating the order policy for the month to someone else means giving them a list of 30 possible order levels, one for each beginning inventory level. If each month has a different policy, then that list would have 1 2 times 30, or 360 numbers!

UFIM Tips and Traps

THE TARGET POLICY AN ORDER POLICY WITH A SPECIAL FORM

In many circumstances, fortunately, the order policy can be simplified. If you run a case without supply curves and without supply variation, the order policy usually has a special form. The first two columns of this table are a sample order policy from a non-seasonal case.

Beginning Inventory	Order Amount (in Ktons)	Inventory + Order
0	200	200
10	200	210
20	190	210
30	180	210
40	170	210
50	160	210
60	150	210
70	140	210
90	130	210
90	120	210
100	110	210
110	100	210
120	90	210
130	80	210
140	70	210
150	60	210

The third column shows what is special about this order policy. For most of the inventory levels, the sum of the beginning inventory level and the order is a fixed number. In other words, the policy above can be summarized as follows: "Make an order such that the order added to the current inventory equals 210 Kilotons, if possible." (In this example there is a maximum order constraint of 200 Kilotons. This makes it impossible at very low inventory levels to reach the 210 Kiloton mark.)

While this is much simpler, it still does not give you an idea of how much fuel you should have at the end of the period to prepare for the next month. If we subtract the expected burn from this number, it tells us at what inventory level we can expect to end the month, on average. This number is called the target. A target policy provides a target number for each period of the year. Suppose the distribution on burn is as follows:

Burn	Probability
120 kilotons	0.25
150 kilotons	0.50
180 kilotons	0.25

UFIM Tips and Traps

Expected burn is:

(0.25)x120 + (0.50)x150 + (0.25)x180 = 150 kilotons.

The monthly target is therefore (210 - 150) = 60 kilotons.

In this case, the order policy listed above and the target policy here are the same. Notice, however, that the target policy has two advantages:

- It is more concise. You don't need an exhaustive list of how much to order at each inventory level (one number to remember each period versus 30 numbers with the order policy).
- It gives you an idea of how much inventory you will have at the end of the period on average.

A trap is: It is important to notice that the target is calculated from your expected burn level, so it is a goal that you shoot for on average. Because burn is uncertain, you cannot guarantee that you will reach a particular inventory level at the end of the month. The best you can do is make an estimate of your expected burn, and make an order taking into account this potential burn (as well as potential future disruptions). Some months (with low burn) you will end up above the target, and some months (with high burn) you will end up below the target. If you do not end a month with exactly the target level of inventory, it does not necessarily mean that you have failed to follow the target policy. For example, according to the order policy above and your burn distribution, if you order so that your inventory plus order amount equals 210 Kilotons, then the distribution on your end-of-period inventory will be

210 -180 = 30 kilotons	0.25
210 – 150 = 60 kilotons	0.50
210 -120 = 90 kilotons	0.25

Notice that with .50 (.25+.25) probability, you will end up at either 30 or 90 Kilotons of fuel at the end of the month--each of these is 30 Kilotons away from the target. But note that the average inventory level is indeed 60 kilotons, exactly the target. When you end up with 30 kilotons you have still followed the target policy; the burn outcome was higher than average, so the ending inventory was lower than average (i.e., lower than the target).

WHEN THE ORDER POLICY DOES NOT HAVE THE SPECIAL FORM

What happens when the order policy is such that the sum of your initial inventory and monthly order is not a fixed number? (This usually occurs when you use supply curves; UFIM adapts its order policies to take advantage of special price structures such as quantity discounts.) The order policy will be an accurate description of the least-cost strategy, but it will not have the special property that makes it the same as a target policy. As we mentioned above, a target policy can be very useful for communicating and understanding the inventory situation. When the least-cost order policy is not a target policy, UFIM will find a target policy that approximates the order policy.

This picture represents what is happening:

UFIM Tips and Traps

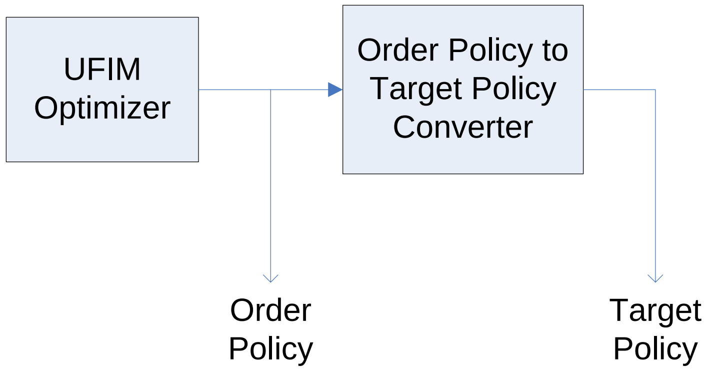

In many cases (when there are no supply curves), the order and target policies describe the same actions. When the order policy does not translate exactly into a target policy, the target policy is an approximation of the order policy.

A tip is: When the order policy is not exactly a target policy, UFIM will give you a warning in *Run Log* (accessed via the *Results* menu). When you see this warning, you should check the order policy to see what it looks like. The order policy may be significantly different from the target policy. If you have entered supply curves, the order policy may give you information about conditions under which you should place especially small or large orders.

A trap is: When UFIM creates an approximate target policy, the reported costs from Run Level 2 (Disruption Management and Policy Development Submodels) are always the costs associated with the original order policy. If you want to know the costs associated with the target policy, run the target policy through the "User-specified policy" enhancement.

UFIM Tips and Traps

TIPS 'N TRAPS #13: TARGETS AND EXPECTED END-OF-PERIOD INVENTORIES

PROBLEM

There are two tables in the *Summary* results that list inventory levels. One is called "Least-Cost Normal Times Target Policy" and the other is "Expected Stocks at the End of Each Period". Sometimes these tables contain similar inventory levels, sometimes not. The expected stocks supposedly result from following the target policy, so shouldn't they match the end-of-period target levels?

EXAMPLE

During most months I have no minimum order level, but in January my fuel supply contract requires a minimum order of 150 Kilotons of coal. I also sometimes have demand disruptions in July; every two years we have a particularly hot month and the demand for electricity doubles. Hot months come with no warning. I can order during a disruption but cannot utilize burn reduction.

UFIM gives me the following tables:

Month:	"End-of-Period" Targets:	Expected Stocks at the End of Each Period:
Jan.	60	75
Feb.	60	60
Mar.	60	60
Apr.	60	60
May	60	60
Jun.	60	60
Jul.	60	90
Aug.	60	64
Sep.	60	60
Oct.	60	60
Nov.	60	60
Dec.	60	60

SOLUTION

Expected ending inventories can differ from end-of-period targets for any combination of the following reasons:

- Disruptions
- Supply Constraints (i.e. a minimum or maximum order)
- Supply Curves.

UFIM Tips and Traps

EXPLANATION

DISRUPTIONS

Remember that end-of-period targets are goals you try to achieve on average during normal times, by following the normal times order policy. Expected ending inventories are stock levels averaged over all periods, including disruptions. Inventory levels following disruptions are often much different than normal times inventory levels. Therefore expected inventory levels over all periods (normal and disrupted) differ from the normal times targets, which are goals for normal periods only.

In this example, if there is no heat wave in July then we follow the normal times order policy. The end-of-period target level is the goal we expect to reach on average when using this policy. The expected ending inventory for a normal July would therefore be 60 Kilotons, the same as the target level. During a disrupted July both the demand distribution and the demand variance increase. We order according to the disruption management policy, and the expected ending inventory is different: 120 Kilotons. (We leave out the calculations as they are not particularly relevant.)

The actual expected ending inventory for July lies between these two figures. Especially hot months occur on average once in two years, so in a given July there is a 50% chance of a disruption occurring. Therefore the expected ending inventory for July is:

(1/2)x120 + (112)x60 = 90 Kilotons,

somewhat higher than the normal times end-of-period target.

Why do the end-of-period targets only reflect the results of following the normal

times order policy? Disrupted inventory scenarios are often so critical and so different than normal periods that you need information dealing specifically with them. UFIM gives detailed disruption management information in *Disruption* results.

FUEL ORDERING CONSTRAINTS

As explained in Tips 'n Traps #12, UFIM's least-cost order policy is often of a form that is equivalent to a target policy. That is, for each month the sum of the month's order and the initial inventory is fixed. (In this example assume that the order policy fixes the sum at 210 Kilotons.) UFIM uses this fixed sum to calculate its target policy.

Unusually tight supply constraints can sometimes prevent you from ordering "up to" this fixed level. In this example a minimum purchase requirement in January forces us to order at 'east 150 Ki'otons. Suppose we start January with 90 Ki'otons. Then we have to order "up to" 150+90=240 Kilotons, 30 more than the fixed level UF1M assumes when calculating its target policy.

Whenever you can be forced to order an amount different than the one used by UFIM to calculate its target policy, your expected ending inventory will be different than the end-of-period target level. In January the minimum order contract sometimes forces us above the 210 Kilotons UFIM assumes for its target policy, so the expected ending inventory for January is higher than the target level. If we weren't forced to order this minimum amount, the expected ending inventory would be the same as the target.

Note that expected ending inventory for August is also higher than the end-of-period target. As in January, this happens because we are sometimes forced to order more than the level assumed by UFIM for its target calculation. If a disruption occurs in July, both our demand for fuel and the variance of our demand will increase. As explained above, this leaves us on average with a much larger ending inventory than during normal times. Occasionally we will even end a hot July with more than 210 Kilotons of inventory; in that case, even by ordering nothing we will begin August with more than the fixed sum assumed by the target policy. The result is a slight increase in the expected ending inventory level for August.

SUPPLY CURVES

Tips 'n Traps #12 explains that if you enter supply curves, UFIM's least-cost order policy won't be entirely convertible to a target policy. UFIM will come up with a target policy, but it will only be an approximation. The expected ending inventory levels are calculated from the order policy, not the target policy. Therefore when you follow the least-cost order policy, your expected ending inventory levels will be slightly different than the approximate target levels.

A Tip Is: Large differences between target inventory levels and expected ending levels may signal a way for you to save money. For example, if you find that the difference is due to a supply constraint, try running the same case with the constraint removed. The difference in cost between the two runs is an indicator of how much the constraint is "costing" you. Removing constraints (for example, removing minimum orders from contracts or obtaining more railroad capacity) may decrease your overall inventory costs.