AREA INVESTMENT STRATEGY MODEL

Version 1.6

USER’S MANUAL

User’s Manual

Area Investment Strategy Model

Version 1.6

Prepared by

Stephen W. Chapel, EPRI
Charles D. Feinstein, VMN Group LLC
Peter A. Morris, VMN Group LCC
Mukund N. Thapa, Stanford Business Software

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS PACKAGE WAS PREPARED BY THE ORGANIZATION(S) NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) NAMED BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS PACKAGE, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS PACKAGE IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS PACKAGE OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS PACKAGE.

ORGANIZATION(S) THAT PREPARED THIS PACKAGE

Stanford Business Software
PO Box 60398
Palo Alto, CA 94306-0398

ORDERING INFORMATION

Requests for copies of this package should be directed to 1-800-313-3774.

PREFACE

EPRI Area Investment Strategy Model: A Utility Perspective

Recognizing the need to establish a business case outlook toward distribution system investments, Wisconsin Electric has been actively involved in the EPRI project to produce a PC-based Area Investment Strategy Model. Version 1.6 is the third commercial version of that model. This perspective provides an overview of the model and suggests how the model might be used.

Model Overview

The Area Investment Strategy Model directly addresses the issue of minimizing the cost of providing reliable distribution system capacity in the face of uncertain load growth. The question of larger versus smaller capacity additions has always been considered by the distribution planner in seeking that plan which most closely tracks the need, and is thus most economical. Often, several load growth trends have been evaluated to test the sensitivity of a plan to different load conditions. That approach, however, examines only smooth trends and not the discontinuous changes in growth that will be potentially experienced in a particular area of the distribution system over time. This new model explicitly recognizes random yet related changes in load growth trend and provides a better means of comparing large versus small increments of capacity.

This new model does not change the basic distribution planning process, but rather augments that process by allowing further insights to be gained into the robustness of a plan. The planner determines area capacity and capacity problems, including thermal, voltage, and reliability, as usual, and then constructs capacity addition plans that meet the identified needs. If the planner wishes to use the model to compare several plans, certain inputs related to 1) the load growth characteristics of the area and 2) the characteristics of the specific capacity additions must be determined.

The load growth characteristics are placed in a matrix which defines several growth rates, usually the expected rate and a low and high limit, and the probabilities of staying in each of these rates or transitioning to one of the other rates. There is a procedure and spreadsheet available to help the planner assess these parameters.

The characteristics of specific capacity additions include the amount of capacity, capital cost, operating costs, and impact of the capacity additions on losses and customer value of unserved energy. These capacity additions can easily include distributed resources. A scenario with only the costs of losses and unserved energy may be constructed to capture the most likely impact of making no improvements in an area. Capacity additions can be easily constrained to follow a certain installation order or to be mutually exclusive, if the planner's judgment deems this to be more realistic.

The model can typically be run in less than one minute. It first quickly determines a tree structure that defines all possible load paths and then fits capacity additions to those paths in order to find the most economical strategy of meeting any given load trajectory. This optimum strategy is output along with supporting information. It is possible to view other paths to ascertain the robustness of the optimum strategy and it is possible to easily change parameters to test the sensitivity of the optimal strategy to changes in the parameters.

The Version 1.6 User's Manual is a complete description of the Area Investment Strategy Model.

Suggestions for Model Use

The model is best applied in those situations where the installed cost of the larger capacity items exceeds $50/kW or the initial utilization cost (installation cost divided by the load placed on the facility in the year of installation) exceeds $100/kW. This would tend to exclude a new package substation located under an existing transmission circuit and feeding into an existing distribution infrastructure. The need for a new substation with a lengthy transmission supply or one located in a highly urbanized setting would provide a likely candidate. Of course, the judgment of the planner or the existence of special circumstances, such as anticipated public opposition to a project or a system area that is well suited to a distributed resource solution, should govern application of the model. The dollar limits cited are intended only as a rough guide. Use of the model in examining plans will provide valuable insights into whether a large investment in distribution capacity is wise at a given point in time or, if instead, a small increment should be added now and further additions made contingent on actual growth.

Paul Freischmidt
Wisconsin Electric

Acknowledgments

We are pleased to acknowledge the valuable guidance we have received while working on projects with our colleagues at Wisconsin Electric, Green Mountain Power, Ontario Hydro, Commonwealth Edison, Public Service Electric and Gas, Nashville Electric Service, the members of the EPRI Area Investment Planning Project Advisory Group, and Jerry Bloom at EPRI.

We also acknowledge useful discussions on model design with Ren Orans, Brian Horii and Greg Ball of Energy and Environmental Economics. Greg and Brian collaborated extensively with us on the early model design and were key contributors to the Ontario Hydro and Wisconsin Electric case studies.

Bob Chow of Ontario Hydro motivated the original model development and hosted the first case study. He contributed time and technical expertise and his observations about the nature of load uncertainty were pivotal in the final design of the load model. Paul Freischmidt and John Nesbitt of Wisconsin Electric provided extensive assistance throughout the model development and testing process. Paul helped the project team understand key engineering-economic issues especially with respect to unserved energy and losses. John Nesbitt has been a champion from the beginning for developing an approach to distribution investment planning under uncertainty. John encouraged EPRI to develop the area investment methodology. He reviewed our work throughout the project and provided the guidance necessary to keep us on track.

Finally we would like to acknowledge Jonathan Lesser, formerly of Green Mountain Power. Jonathan's advice improved our understanding and representation of operating cost issues related to distributed generation and load control programs. He has been a model user and tester from the beginning. He was and continues to be our most effective beta tester. In the process of using the model to solve real-world problems, he has discovered and helped us fix literally dozens of modeling and programming problems. Jonathan has made important contributions to the model design and to the robustness of the software. We thank Jonathan for his persistence, patience and belief in quantitative economic analysis.

S.W.C.
C.D.F.
P.A.M.
M.N.T

TO REPORT PROBLEMS AND GET HELP RUNNING THE MODEL, CONTACT EPRI USER SUPPORT:

1. The EPRI project manager
   Steve Chapel
   (650) 855-2608
   schapel@epri.com

2. The Software developer
   Mukund Thapa
   Stanford Business Software
   (650) 856-1695
   mukund@sbsinc.net

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1 INTRODUCTION

1.1 Installing the Model

The accompanying CD disk contain the setup program that will install the model. The system requirements for the model are:

- Windows 95 or a more recent version of Windows,
- 5mb of RAM,
- 8mb of free space on a hard drive,
- Acrobat Reader 4.0 installed on the host machine.

Acrobat reader is used by the model to display the help file. If the user does not have the software they can download it for free from Adobe's website (www.adobe.com).

To install the model insert the disk in the CD-ROM drive. If the setup program does not start automatically, click **START**, select **Run**, enter **D:\SETUP** in the dialog box and click **OK** (if **D:** is not the drive identifier for the CD-ROM, enter the appropriate drive identifier). Then, follow the directions on the screen to complete the setup.

Unless the user specifies otherwise during installation, the program and case files are placed in the *c:\Program Files\EPRI\AreaInvest* directory. The DLL files are placed in the *c:\windows\system* directory.

Note: Four example data sets are provided with the software. The input files are ***case1.aip***, ***case2.aip***, and ***real_case.aip***. The outputs have the same names but with extension ***.Report*** (outputs are stored in a sub-directory labeled *Sample Outputs).* There is also an input data set for the Load Assessor tool. The example load assessor files are ***assessor.drs*** and a***ssessorJump.drs***. These data sets are the inputs used for the tutorial in CHAPTER 4 and the "Real Utility Example" in CHAPTER 5.

If the user wants to verify that the model is correctly installed and working, he or she can run each of the four cases (using the files with extension *.aip*) and compare the resulting output files with the outputs stored in the sub-directory *Sample Outputs*. To do this the user should:

1) Open the program by clicking *START*, selecting *Program*, and highlighting the *Area Investment Strategy* program, and double-clicking the light bulb icon.

2) Load a case. To do this, open the *File* menu and open one of the sample cases supplied, which have the extension *.aip*, by selecting *Open Case*.

3) Run a case. To do this open the *Analysis* menu and select *Run* and click *RUN* when the dialog box appears. Successful execution of the case occurs when a cost summary report appears.

4) Compare the resulting output files with the appropriate output files in the subdirectory *Sample Outputs*.

NOTE: The user should NEVER open a new input file and attempt to run the case using default values (by opening the *File* menu and selecting *New Case*). In addition, the first time user should NEVER attempt to build a case from scratch. Users should use the data sets distributed with the software as templates for constructing inputs. There are two reasons:

First, building an input data set from scratch requires a good understanding of both the analysis methodology and the functionality of the software. Developing this understanding requires, at a minimum, working through and studying the tutorials in CHAPTER 4 and CHAPTER 5.

The default values in a new data set are not valid data inputs. If the user attempts to run the model with these inputs, the model will return data error messages and the model will not run. *The decision to not provide a valid default data was part of the software design*. To correctly analyze infrastructure strategy options, utility analysts should think carefully about how to structure their problems and about the inputs that are appropriate given their specific problems. Providing a default data set would not encourage the necessary thinking.

1.2 Learning To Use The Model

The Area Investment Strategy Model is designed to help utilities solve distribution infrastructure planning problems. Successful application of the model requires learning not only how to enter data and run the model but also requires some understanding of the analysis methodology behind the software

Before attempting to use the model to solve real utility problems, the user should work through the step-by-step tutorial in CHAPTER 4. *To successfully complete the tutorial, the user must EXACTLY follow the steps described in CHAPTER 4.* The tutorial is designed to help the user begin to understand the functionality of the model and the analytical methods that underlie the model.

After completing the tutorial, CHAPTER 4, we recommend that the user work through the real case application described in CHAPTER 5. This chapter, like the CHAPTER 4 tutorial, provides a set of step-by-step instructions that illustrate an investment analysis. After completing CHAPTER 4 and CHAPTER 5 the user should be ready to begin to apply the model to their specific investment problems.

1.3 Error Traps

The software traps most data inconsistencies at the user interface level. However, there are certain situations where error conditions are recognized only after the model is processing the data. In order to ensure robustness of the software, the software traps run-time generated errors and reports it back to the user in a Window as a RUN-TIME ERROR. When this happens the user is asked click OK, the model returns to the main window.

In some (hopefully rare!) instances, the error message is one that should not have occurred and. occurs only because of an internal inconsistency in the software. In this case, the names and phone numbers of contacts at EPRI and Stanford Business Software, Inc. (SBSI) will be displayed. Please contact at least one of them with the exact message as it appears on your screen.

In other situations, the error message will point you in the right direction and you can attempt to modify the data for the case and re-run the model (after restarting the software). However, in this case too, should the message be unclear or if you have any questions, please do not hesitate to contact EPRI or SBSI.

If the model detects that there is not enough capacity to meet capacity needs under all potential load growth trajectories; it reports an error because then there is no feasible solution. You can get past this error in four ways: a.) adding additional alternatives, b.) decreasing load growth, c.) decreasing the planning horizon, or d.) modifying the constraints. Ideally, the model should automatically handle this situation and report a warning. However, such a scheme is not as simple as it sounds and we are attempting to find a good solution for a future release of the model.

1.4 What's New

Several new features and enhancements have been added to the model since the release of Version 1.0.

1) A *Do Nothing* investment option has been added. *Do Nothing* is used to represent the decision to delay any capacity additions until a given amount of load growth has occurred. See the "Capacity Alternatives" section in CHAPTER 2.

2) Band-Aids can be salvageable or not. This allows for better modeling of investments that can be removed and reused. Portable substations and distributed generators are examples of potentially salvageable investments.. See the "Capacity Alternatives" section in CHAPTER 2.

3) A load uncertainty assessment tool has been added to the model. This tool facilitates development of the required load-growth input parameters. For a description see section 2.4, Using The Load Assessor Tool.

4) The speed of the load forecasting algorithm has been increased dramatically.

5) There are additional output reports and an improved user interface.

1.5 Contacting Us

TO REPORT PROBLEMS AND GET HELP RUNNING THE MODEL CONTACT EPRI USER SUPPORT:

1. The EPRI project manager
   Steve Chapel
   (650) 855-2608
   schapel@epri.com

2. The Software developer
   Mukund Thapa
   Stanford Business Software
   (650) 967-6998
   mukund@sbsinc.net

CHAPTER 2 MODEL USER’S GUIDE

2.1 Introduction

This User's Guide for the Area Investment Strategy Model describes the user interface and the way inputs and outputs are handled by the interface. The user interface is a standard Windows implementation. It was written in Visual Basic. Because of the Windows implementation, users should be able to edit data sets, run the model and view results with little or no reference to this User Guide. However the modeling process requires considerable thought and it is highly recommended that model users carefully study the methodology in the rest of this User's Manual.

2.2 Running The Model

Open the model by clicking **START**, selecting **Program**, and highlighting the **Area Investment Strategy** program, and double-clicking the light bulb icon.

When the model is started, a start-up window is displayed for about ten seconds. This screen has the name of the program, the version number, and the list of individuals that developed the program. If you want to close the screen before it automatically closes, double click on the light bulb in the middle of the screen.

Once the opening screen disappears, the model displays a dialog box asking if you want to work on a new case or an existing case. If you want to work on an existing case, you must select the case. If you choose not to create a new case or work on an existing case, you can later make the selection through the File menu option.

Following this, the main window for the program is displayed through whose menu system you can edit files, save files, run the model and view results. This window is shown below in Figure 2.1.

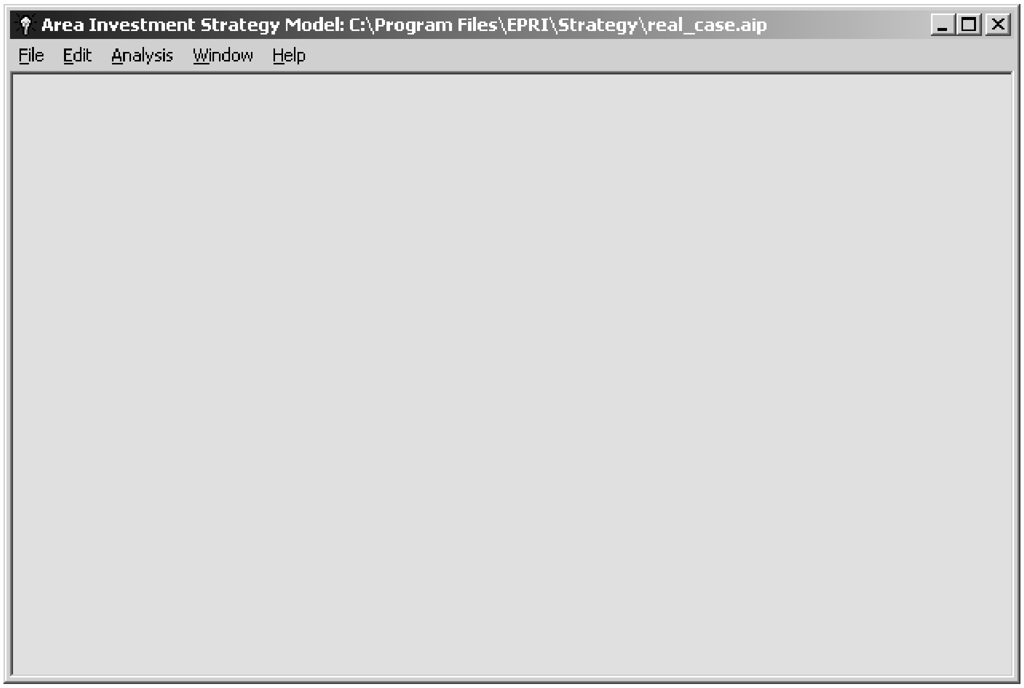

Figure 2.1 Main Screen

It would be a good idea to verify the correctness of your installation now. To do this open the File menu and open one of the sample cases supplied, which have the extension ".aip." Open the Analysis menu and run a case. Successful execution of the case occurs when a cost summary report appears; in this case, the model has been correctly installed.

2.3 Using the Strategy Model User Interface

2.3.1 The File Menu

The file menu is used to create, open, and save input data sets. This menu is shown in Figure 2.2. The file menu has five options, which are described next.

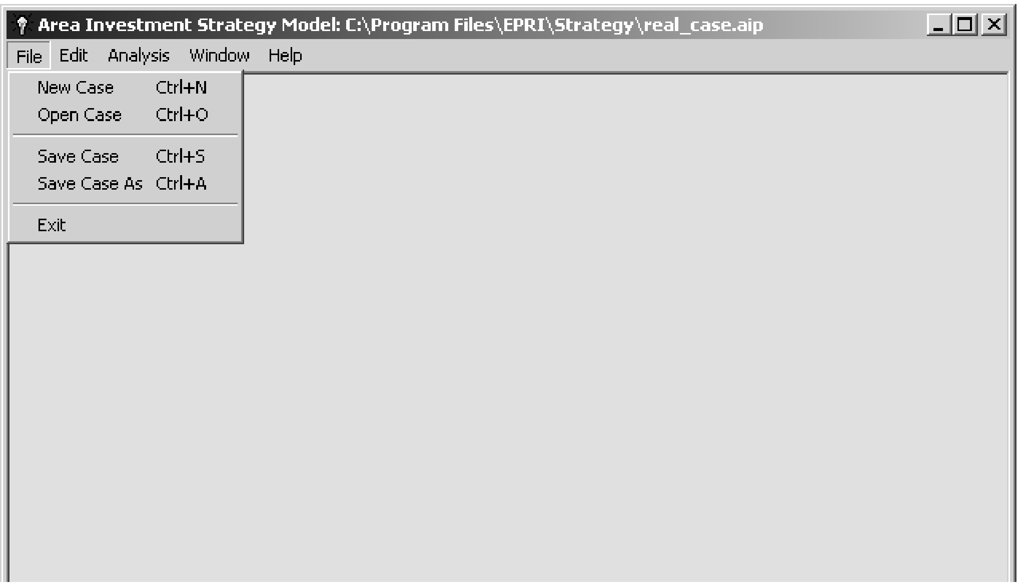

Figure 2.2 File Menu

New Case

Selecting this causes the program to clear all data and let the user enter a new set of data. If a case is being worked on, the user is given an option to save the current data if any of the values have changed.

NOTE: The user should NEVER open a new input file and attempt to run the case using default values (by opening the *File* menu and selecting *New Case*). In addition, the first time user should NEVER attempt to build a case from scratch. Users should use the data sets distributed with the software as templates for constructing inputs.

Open Case

Selecting this causes the program to prompt the user for an existing case through the standard Windows open dialog box. If a case is being worked on, the user is given an option to save the current data if any of the values have changed.

Save

Saves the case data to disk. If it is a new case the user is asked to select a name and location through a standard Windows file save dialog box.

Save As

Saves the case data to disk under a user-specified name and location. The user is asked to select a name and location through a standard Windows file save dialog box.

Exit

Exits the program. If the case data has changed and not been saved, the user is prompted for confirmation.

2.3.2 Editing Data Sets

The Edit menu is used to edit both new and existing data sets. Here we outline the different user data entry options.

Section 3.2 (Model Description), summarizes the input data requirements for the investment strategy model. In addition, CHAPTER 1 contains the technical documentation. The user should review both of these sections to gain familiarity with the methodology that underlies the model. To use the model effectively, the user will need to study CHAPTER 1 and 5 in detail. That study is probably best done in the context of doing a company-specific analysis.

Figure 2.3 shows the items in the edit menu.

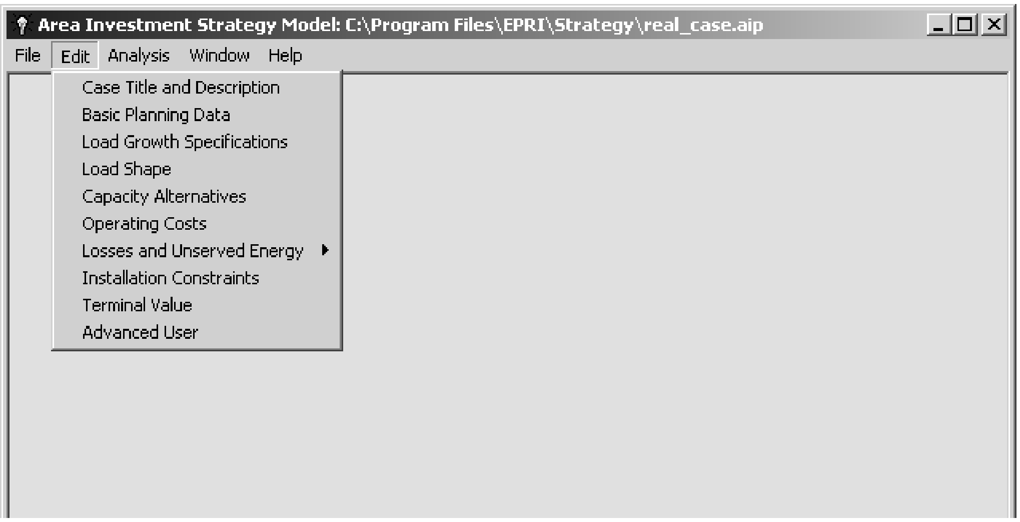

Figure 2.3 Edit Menu Items

Selecting an item in the Edit menu brings up a data dialog box. Selected dialog boxes are shown below along with brief explanations of input parameters.

Operational Features of the Editing Interface

Some of the input forms are interrelated. For example the number of capacity alternatives entered in the Capacity Alternatives form determines the number of rows in the Operating Cost form. Because of this, the user interface will not allow certain forms to be open at the same time.

The input forms are designed so that the user can tab from one data field to another. The one exception is grid inputs. Tabs can be used to move between cells within a grid but a tab cannot be used to exit the grid.

All grids can be edited by selecting a cell in the grid and hitting F2. This is a standard feature of the windows interface. Thus, if you want to edit exiting data in a grid, you can use the F2 function key.

Case Title and Description

Figure 2.4 shows the form used to input the case title and description. Note all input screens use data from the "real utility" example presented in CHAPTER 5.

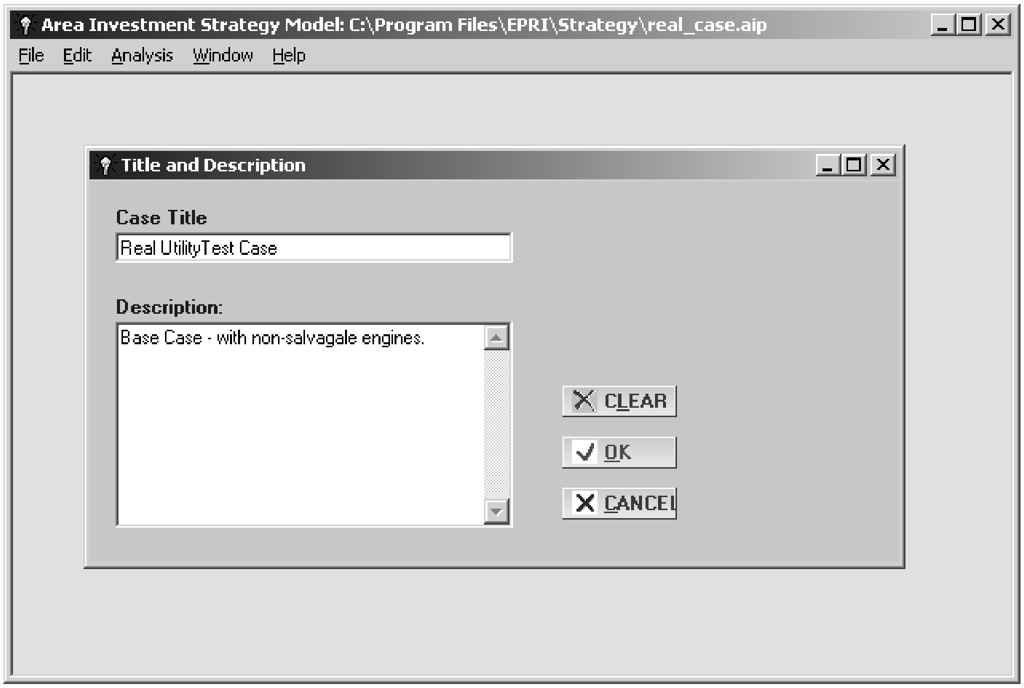

Case Title: Real UtilityTest Case

Description: Base Case - with non-salvagale engines.

Figure 2.4 Case Title and Description

<u>Input</u>	<u>Explanation</u>
Case Title	A name for the case (40 characters or less) – appears on all of the output forms.
Description	A description of the case.

Basic Planning Data

Figure 2.5 shows the basic data input form. Note that the discount rate is in "real" terms (the effects of inflation should not be included in the discount rate parameter). CHAPTER 6, Section 6.4, describes the accounting methods.

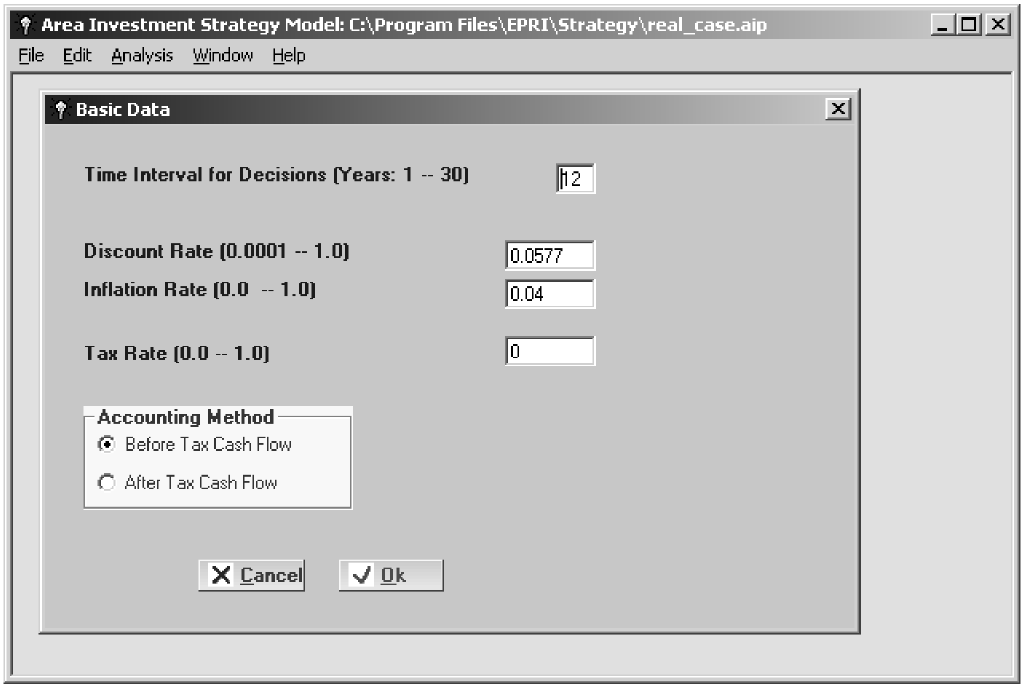

Time Interval for Decisions (Years: 1 -- 30) 12

Discount Rate (0.0001 -- 1.0) 0.0577
Inflation Rate (0.0 -- 1.0) 0.04

Tax Rate (0.0 -- 1.0) 0

Accounting Method
● Before Tax Cash Flow
○ After Tax Cash Flow

Figure 2.5 Basc Planning Data

Input	Explanation
Time Interval for Decisions	The time interval over which the decisions and uncertainties are modeled in detail.
Discount Rate	The discount used for calculating present values. This rate is "real" and thus should not include the effects of inflation.
Inflation Rate	The general inflation rate assumed for the study. This defines the basis for escalating costs.
Tax Rate	This is the overall corporate tax rate used when applying the after tax cash flow method. The user should read CHAPTER 6 Section 6.4 for technical details. It is recommended that a default of .34 be used - the federal tax rate for large corporations.

Accounting Method	The user must select either the Before Tax Method or the After Tax Method. This determines whether cash flows are calculated on a before tax or after tax basis. If the after tax method is chosen, the user should specify the depreciation schedule in the Capacity Alternatives screen.

Note: The Strategy Model uses real cash flows for economic calculations. The equation that underlies all Strategy Model present values calculations is:

$$PV = \sum CF_t (1 + ESC)^t / [(1 + RDR) (1+ INF)]^t$$

Where PV is present value, CF_t is the cash flow in year t measured in year t=0 dollars, ESC is the escalation on the cash flow component, RDF is the real discount rate, and INF is the inflation rate.

Load Growth Specifications

Figure 2.6 is the form used to record load growth specifications. The user should read the "Modeling Load Uncertainty" section in CHAPTER 6 before entering data here. The load modeling approach is based on the idea that it is useful to describe load growth in terms of multiple growth trends that persist for uncertain durations. The model requires a relatively small number of input variables. The user must specify 1) information on load growth trends (growth rates and trend-to-trend transition probabilities), 2) the initial load condition, and 3) load saturation parameters.

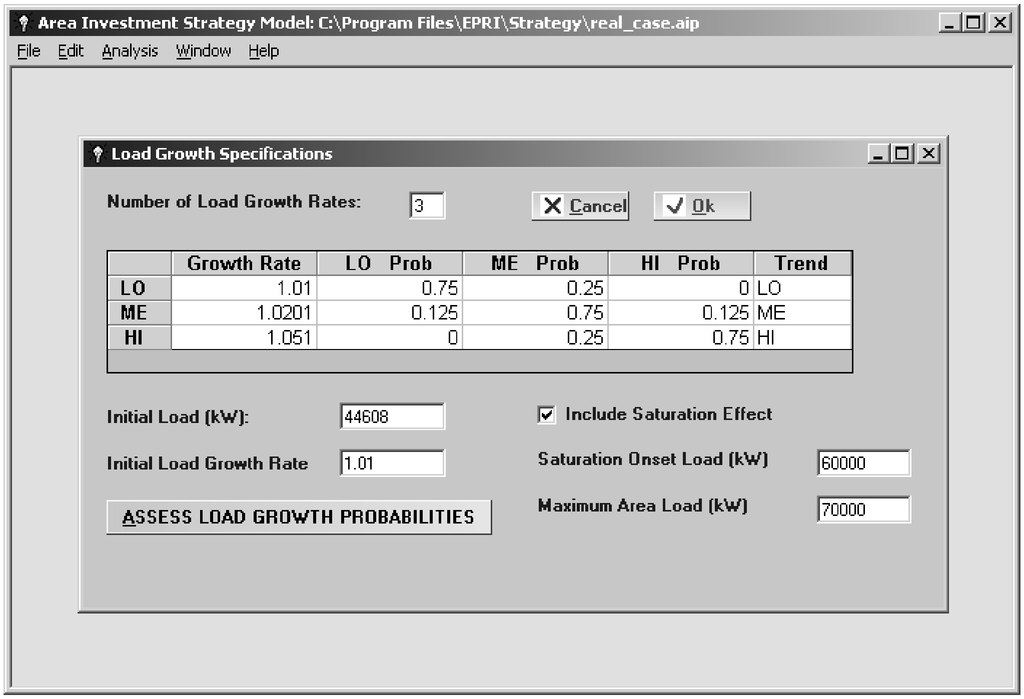

Area Investment Strategy Model: C:\Program Files\EPRI\Strategy\real_case.aip

Load Growth Specifications

Number of Load Growth Rates: 3

	Growth Rate	LO Prob	ME Prob	HI Prob	Trend
LO	1.01	0.75	0.25	0	LO
ME	1.0201	0.125	0.75	0.125	ME
HI	1.051	0	0.25	0.75	HI

Initial Load (kW): 44608

Initial Load Growth Rate 1.01

☑ Include Saturation Effect

Saturation Onset Load (kW) 60000

Maximum Area Load (kW) 70000

Figure 2.6 Load Growth Specifications

<u>Input</u>	<u>Explanation</u>
Number of Load Growth Rates	Load growth uncertainty is modeled using up to five load growth trends. In the example input form above there are three growth trends.
Growth Rate	For each load growth trend there is an annual growth factor, expressed in decimals as (1 + annual load growth rate).
Transition Probabilities	In the example input form above, the rows in the matrix are the probabilities of jumping to each of the growth trends (LO, ME and HI) for the next year, conditional on the current trend. Think of the rows as "from" and the columns as "to." An assessment tool designed to help the user develop these input probabilities is included with the software. The tool is activated by clicking the "ASSESS LOAD GROWTH PROBABILITIES BUTTON.

Trend	These are user inputs used to label each of the growth trends (LO, ME, and HI for the example shown in above.
Initial Load	The peak load in kilowatts at the start of the planning period.
Initial Load Growth Rate	The growth at the start of the planning period. This input cannot be lower than the lowest growth trend or higher than the highest growth trend.
Include Saturation Effect	If this box is checked, load in the area will be constrained by the model to not exceed the maximum area load (the saturation level). If the box is checked, the user must specify the two additional inputs listed below.
Saturation Onset Load (kW)	Peak loads above this level grow progressively more slowly and approach the maximum area load asymptotically
Maximum Area Load (kW)	The maximum level that load can grow to in the area -- the load saturation level.
Assess Load Growth Probabilities	If the user clicks this button an add-in assessment tool is activated (Load Assessor). This tool requires that the user answer a set of questions concerning current load level and growth, and future load growth scenarios. Based on the answers, the tool estimates the inputs for the growth trends and transition probabilities. See Section 2.4, for a description of the Load Assessor Model, including the input and output screens.

Load Shape

The Annual Load Duration Curve is entered using the Load Shape form - Figure 2.7. This input is explained in the "Model Description" section in CHAPTER 3

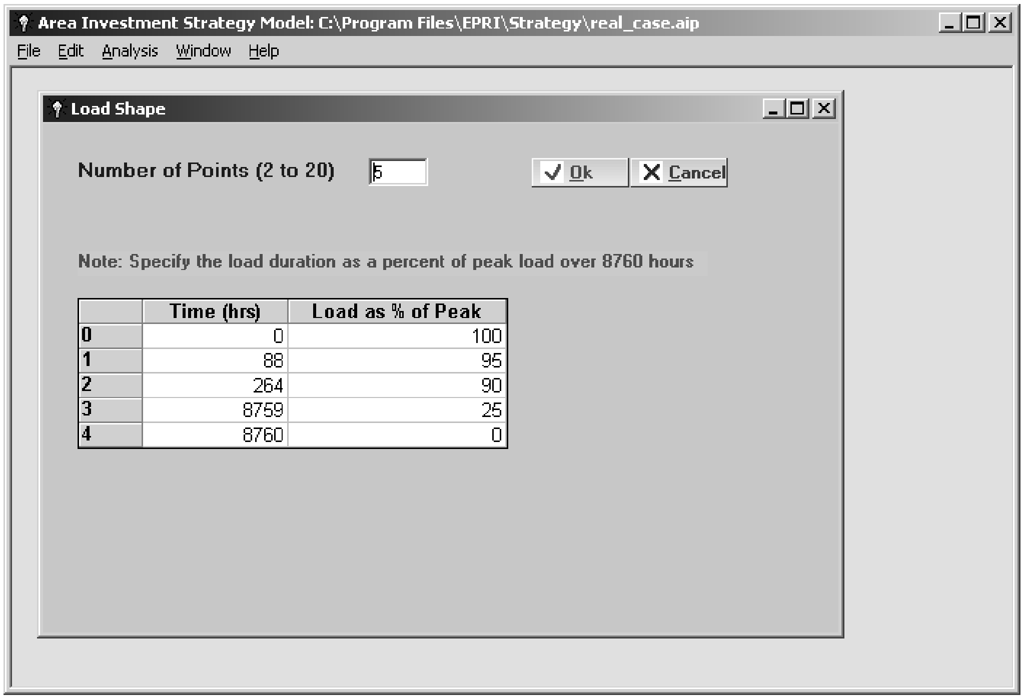

Number of Points (2 to 20) 5

Note: Specify the load duration as a percent of peak load over 8760 hours

	Time (hrs)	Load as % of Peak
0	0	100
1	88	95
2	264	90
3	8759	25
4	8760	0

Figure 2.7 Load Shape

Input	Explanation
Number of Points	This is the number of points used to specify the load duration curve. The curve is forced to pass through the two points (0, 100%) and (8760, 0%). The user can specify intermediate points. If none are specified, the load duration curve is assumed to be a triangle.
Time and Hours	These are pairs of time and greater-than-or-equal-to-probability points: for example in Figure 2.7, load is ≥ 95% of peak load for 88 hours.

Capacity Alternatives

Capacity alternatives, capital costs, escalation, lead-times and depreciation schedule are entered using the Capacity Alternatives form - Figure 2.8.

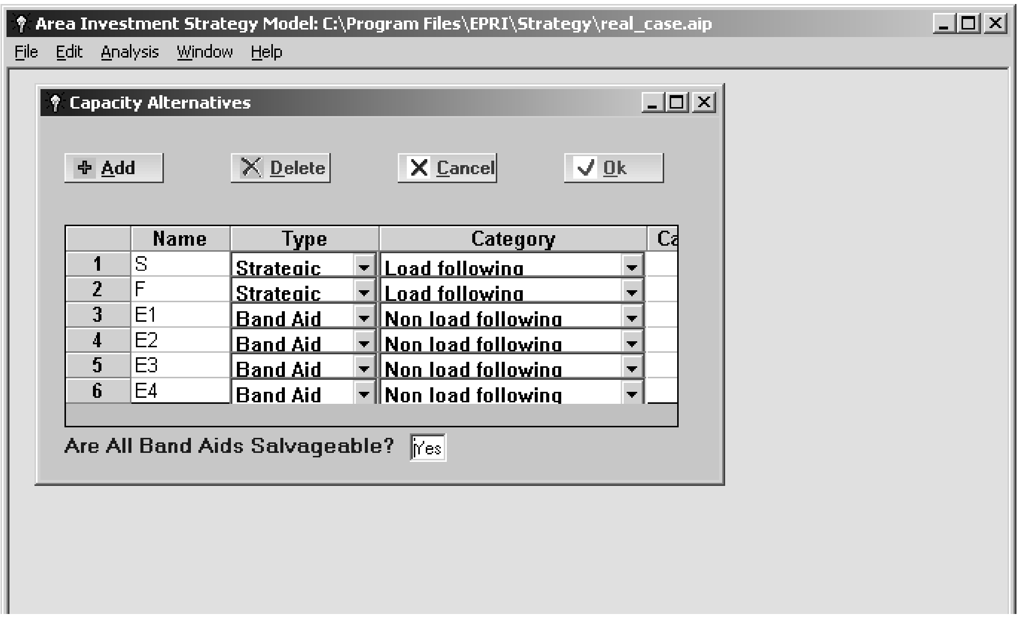

	Name	Type	Category	Ca
1	S	Strategic	Load following	
2	F	Strategic	Load following	
3	E1	Band Aid	Non load following	
4	E2	Band Aid	Non load following	
5	E3	Band Aid	Non load following	
6	E4	Band Aid	Non load following	

Are All Band Aids Salvageable? Yes

Figure 2.8 Capacity Alternatives

<u>Input</u>	<u>Explanation</u>
Name	A name used to label specific capacity alternatives. For example, S and F in Figure 2.8 are used to represent a substation and a feeder.
Type	The user must select one of three types of alternatives: *Strategic*, *Bandaid or Do Nothing*. *Strategic* is used to represent major capacity additions, such as substations. *Bandaid* is used to represent smaller alternatives such as capacitors, distributed resources and DSM programs. *Do Nothing* is used to represent the decision to delay any capacity additions until a given amount of load growth has occurred. Bandaid alternatives (B1, B2, B3, and B4 in Figure 2.8 are constrained to be installed in the sequence that they appear in the input form. There are only two numerical inputs associated with *Do Nothing:* (1) the capacity which defines the amount of the delay, and (2) variable O&M (see the operating cost screen, Figure 2.9). For *Do Nothing*, variable O&M is used to define the cost of unserved energy.
Category	The user must select one of four alternatives: Load Following, Non-Load Following, Load Following DSM, or Non-Load

	Following DSM. This parameter is used for operating cost calculations. See Section 6.7 for an explanation of the operating cost calculations.
Capacity (kW)	The peak capacity (in kilowatts) added by the alternative.
Capital ($000)	The current year capital cost in thousands of dollars.
Cost Escalation	This is an escalation factor relative to the rate of inflation. A value of 1.0 means that the cost of the alternative escalates at the rate of inflation. Values less than 1.0 mean that the escalation is less than the rate of inflation. The factor must be in the range 0.01 to 2.0.
Lead Time (yrs.)	This is the time in years between making the decision to invest in an alternative and the time at which it is placed in service. Lead times can be caused by such things as licensing, siting delays and construction times.
Life (yrs.)	This is the physical life of the asset (between 1 and 50 years).
Depreciation Schedule	This should be specified when the user wishes to select the after tax cash flow accounting method. Clicking on this box activates an input screen with four alternatives: NONE, Straight Line, MACRS15, and User Defined. See CHAPTER 6, Section 6.4 for details on modeling after-tax cash flows.
Description	A user-provided description of the capacity alternative.
Are All Band Aids Salvageable	The user must select yes or no. This allows for removal and reuse of some types of investments such as distributed generation. If a technology will never be removed and reused, the user should select no. See CHAPTER 6 Section 6.6 for technical details.

Note: We repeat a previous warning. If the model detects that there is not enough capacity to meet capacity needs under all potential load growth trajectories, it reports an error because then there is no feasible solution. You can get past this error in four ways: 1) adding additional alternatives, 2) decreasing load growth, 3) decreasing the planning horizon, or 4) modifying the constraints.

Operating Costs

Operating cost parameters include fixed O&M, fuel, heat rate, variable O&M, and system-energy costs. Operating cost inputs are entered using the Operating Cost form shown in Figure 2.9. Note that system energy costs are defined as costs avoided by the operation of local generation or DSM programs. For traditional T&D options, avoided system energy costs should be zero.

Using the form shown in Figure 2.9, the user can also enter emission costs associated with each of the capacity alternatives. Entering emissions costs is straightforward. It is done by clicking on the emissions tab and filling in the form. The form allows entries for NOX, SOX, CO2, and other. Emissions are entered in tons per kWh and the cost of emissions in dollars per ton.

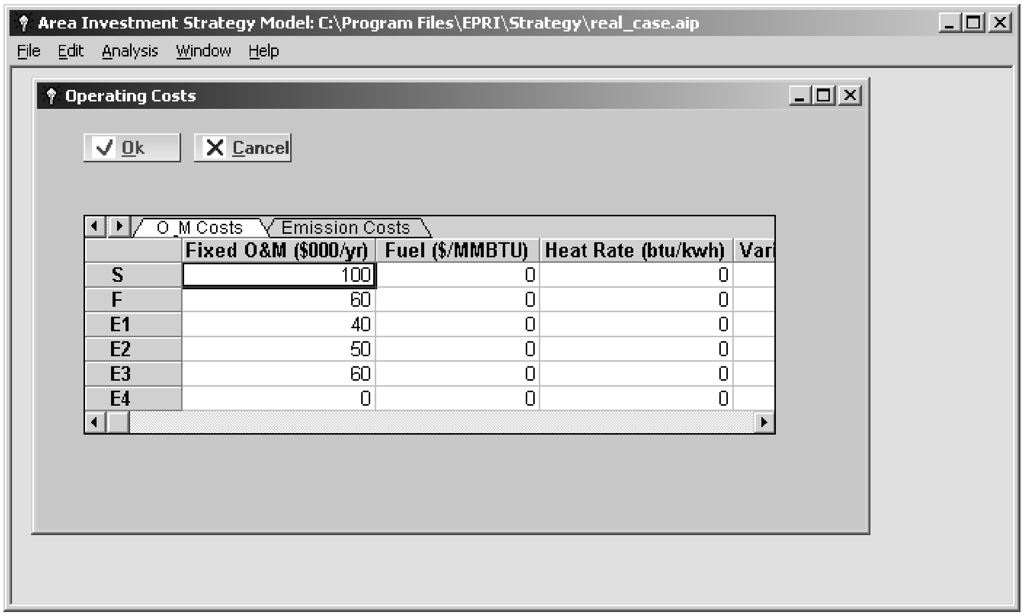

	Fixed O&M ($000/yr)	Fuel ($/MMBTU)	Heat Rate (btu/kwh)	Vari
S	100	0	0	
F	60	0	0	
E1	40	0	0	
E2	50	0	0	
E3	60	0	0	
E4	0	0	0	

Figure 2.9 Operating Costs

Input	Explanation
Fixed O&M ($000/yr.)	Fixed operations and maintenance costs in thousands of dollars.
Fuel ($MMBtu)	Fuel costs in dollars per million Btu (for alternatives that use fuel).
Heat Rate (Btu/kWh)	Heat rate for capacity alternatives that use fuel (in Btu per kWh).
Variable O&M ($/kWh)	Variable operations and maintenance costs (in dollars per kWh). Note that for traditional T&D options, such as circuits and substations, the user should include the cost of energy from the bulk system.
System Avoided Cost $/kWh)	Avoided system energy costs in dollars per kWh. If the user inputs a value greater than zero, the model treats the technology as *NOT* load following. If zero is entered, the model treats the technology as load following. Thus zero should be entered for all load-following devices (such as circuits, and substations). See CHAPTER 6, Section 6.7 for details.

Losses and Unserved Energy

Losses and unserved energy are input using the Losses and Unserved energy forms. The user should read Section 6.8 for a detailed explanation of these inputs. The input forms are shown in Figure 2.10 and Figure 2.11.

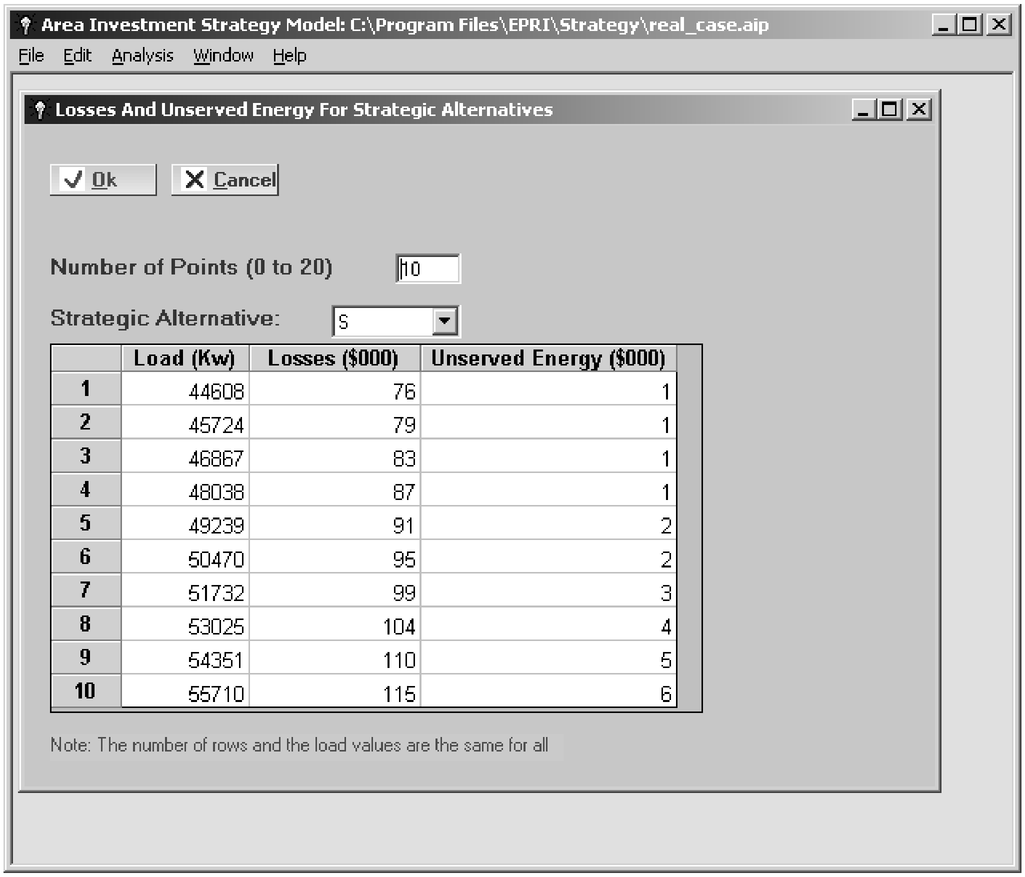

Number of Points (0 to 20) 10

Strategic Alternative: S

	Load (Kw)	Losses ($000)	Unserved Energy ($000)
1	44608	76	1
2	45724	79	1
3	46867	83	1
4	48038	87	1
5	49239	91	2
6	50470	95	2
7	51732	99	3
8	53025	104	4
9	54351	110	5
10	55710	115	6

Note: The number of rows and the load values are the same for all

Figure 2.10 Losses and Unserved Energy for Strategic Alternatives

These inputs allow users to associate the dollar value of losses and unserved energy with operating at specific load levels with specific technologies in the infrastructure technology mix. The default values for these inputs are zero. Thus, if these costs are zero for a particular study, the user does not need to provide the inputs.

<u>Input</u>	<u>Explanation</u>
Number of Points	The number of load levels at which losses and unserved energy needs to be specified. This number defines the number of rows in the losses and unserved energy table. This number of rows will be used for all strategic alternatives where losses and unserved energy inputs are specified.

Strategic Alternative	The user can specify losses and unserved energy for each of their strategic alternatives. For each table of inputs, the user must select one of their strategic alternatives from the drop-down menu. Once selected the table of losses and unserved energy can be filled in. The table of losses and unserved energy apply when that alternative is part of the technology mix. In addition to user-specified strategic alternatives, there is the "Do Nothing" alternative. If the user selects this alternative, the losses and unserved energy costs are applied in the case where none of the specified strategic alternatives are in the technology mix.
Load (kW)	The kilowatt load level for each row in the losses and unserved energy matrix.
Losses ($000)	The dollar value associated with the energy losses given the load level for each row (in thousands of dollars).

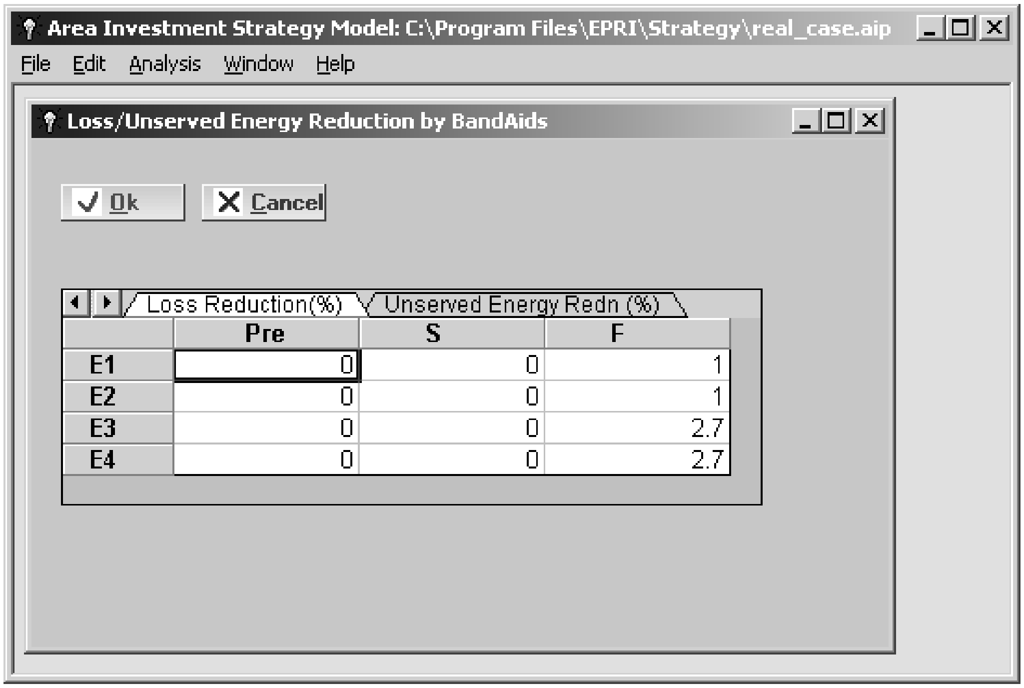

	Pre	S	F
E1	0	0	1
E2	0	0	1
E3	0	0	2.7
E4	0	0	2.7

Figure 2.11 Losses and Unserved Energy Reduction by Band-Aids

Unserved Energy ($000)	The dollars associated with unserved energy, given the load level for each row (in thousands of dollars)..
<u>Input</u>	<u>Explanation</u>
Loss Reduction (%)	The percent reduction in the cost of losses. This can be specified for bandaid installations prior to any of the strategic alternatives (“Pre” column) and for installations with each of the strategic alternatives. The user can provide values for each bandaid alternative. Default values are zero. Note that the values are cumulative.
Unserved Energy Reduction (%)	The percent reduction in the cost of unserved energy. This can be specified for bandaid installations prior to any of the strategic alternatives and for installations with each of the strategic alternatives. The user can provide values for each bandaid alternative. Default values are zero.

Installation Constraints

A key feature of the strategy model is the ability to place constraints on the number and ordering of the capacity alternatives. This is illustrated in the Installation Constraints form - Figure 2.12. In this data set, there can only be one S (substation), one F (feeder), and the feeder cannot follow the substation.

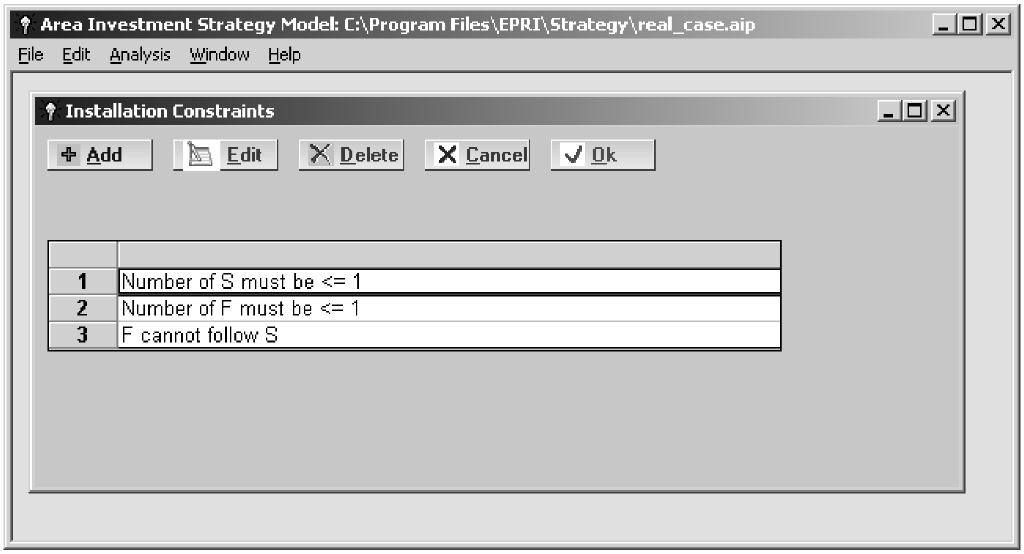

1	Number of S must be <= 1
2	Number of F must be <= 1
3	F cannot follow S

Figure 2.12 Installation Constraints

An input form allows installation constraints to be added, edited and deleted. Figure 2.13 illustrates the form used for specifying constraints. On the left side of the form the user must select a capacity alternative that will appear on the left side of the constraint condition. The middle of the form contains the set of constraint conditions that can be selected. On the right side the user can set quantitative limits on the number of occurrences of specific alternatives. The user can also use the drop-down menu on the right side of the form to select one or more capacity alternatives in order to apply precedence constraints (for example, F cannot follow S)

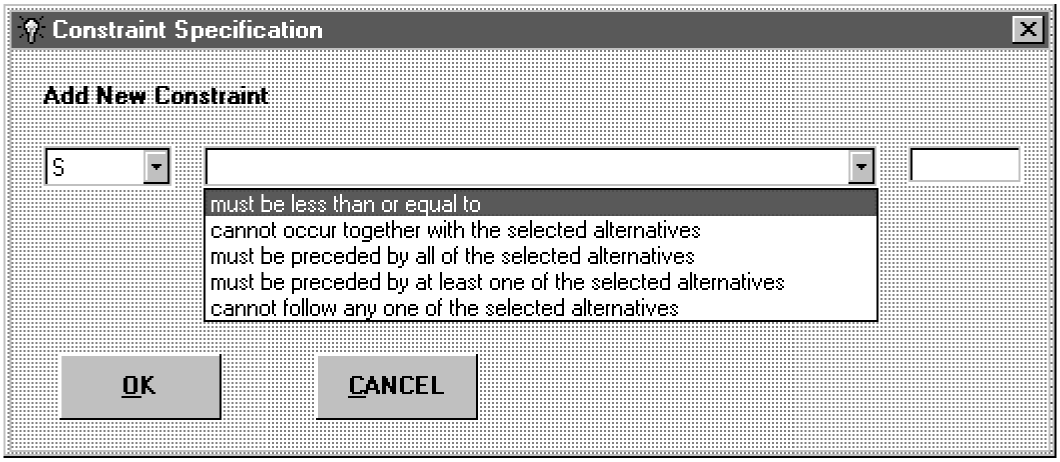

Figure 2.13 Specifying Installation Constraints

Terminal Value

Two options are allowed for terminal value. These are selected using the Terminal Value form. This input data is explained in Section 6.2: "Modeling Planning Period End-Effects."

The first option, terminal lottery on capacity price, assumes that all capacity added during the planning study is sold at the end of the planning period. The sale price is determined by the remaining life of the asset and the input variable Price of Capacity at Terminal Time. The second option, Cost-to-Go, assumes that the assets are kept until they reach the end of their life and then replaced with new assets. The cost of the new assets is the input variable Price of Capacity at Terminal Time.

The input form is Figure 2.14.

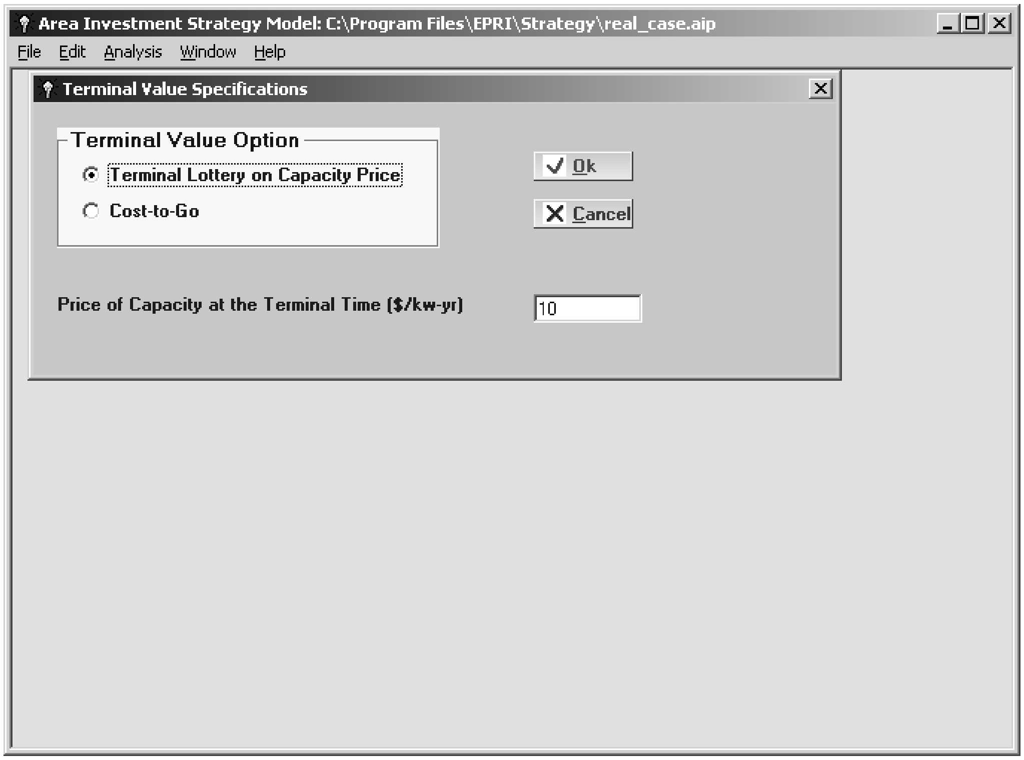

Terminal Value Option

● Terminal Lottery on Capacity Price

○ Cost-to-Go

Price of Capacity at the Terminal Time ($/kw-yr): 10

Figure 2.14 Terminal Value

<u>Input</u>	<u>Explanation</u>
Terminal Value Option	Select either Terminal Lottery on Capacity Price or Cost-to-Go. The User should read Section 6.2 for technical details. The second option, *Cost-To-Go*, should not be selected unless the user is an expert model user and has discussed the details of this option with the EPRI project manager Steve Chapel.
Price of Capacity at Terminal Time	Specify the cost per kW-yr. of buying capacity at the end of the planning period. This should be in current year dollars (dollars as measured in year 0 of the planning study).
Escalation on Price of Capacity	This is an escalation factor relative to the rate of inflation. A value of 1 means that the price escalates at the rate of inflation. Values less than 1 mean that the escalation is less than the rate of inflation. The factor must be in the range .01 to 2.

Operating Cost of Capacity ($/kWh)	This is the average operating costs for all capacity purchased after the end of the planning period. This should be in current year dollars (dollars as measured in year 0 of the planning study).
Escalation on Operating Cost	This is an escalation factor relative to the rate of inflation. A value of 1 means that the cost escalates at the rate of inflation. Values less than 1 mean that the escalation is less than the rate of inflation. The factor must be in the range .01 to 2.
Final Load Growth Rate	This is the long term average load growth for the area. This is the rate that is expected to occur after the end of the planning period. The allowed range is 1.00001 to 1.5.

Advanced User

The final input form is Advanced User - Figure 2.15. This form controls the number of nodes (between 50,000 and 9,999,999) allowed in the decision tree and the coalescence option. These options determine the size of the problem that the model will solve, the speed of execution, and have some effect on numerical results.

Exercising the coalescence option can dramatically improve the run time for large problems (it can actually increase the run time for small problems) but the method is an approximation and can have some effect on the numerical results. Because of this is an approximation, the model is constrained to not allow coalescence before stage 5 of a decision tree. The user should contact EPRI user support for an explanation of these options and their potential effects.

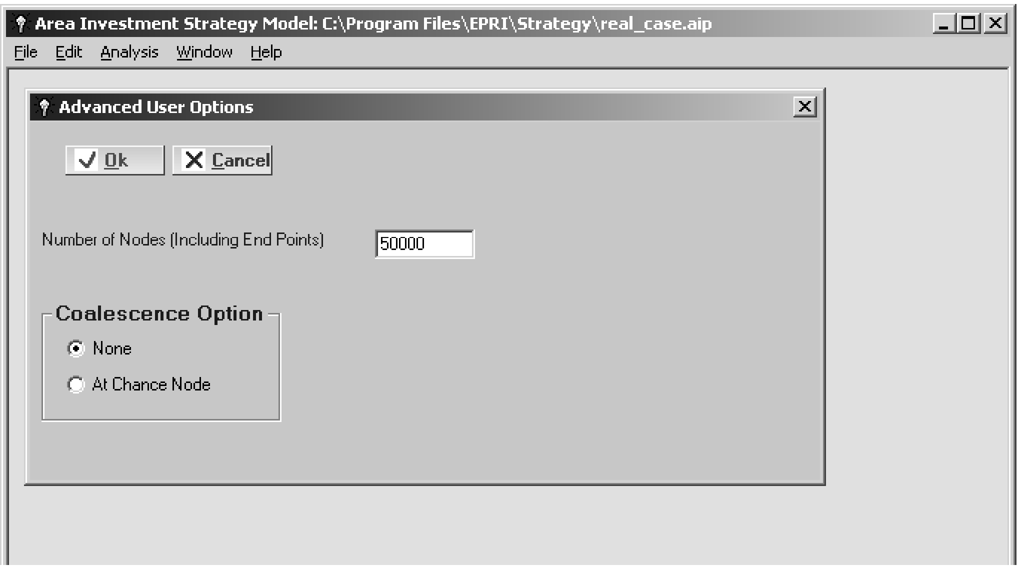

Number of Nodes (Including End Points) 50000

Coalescence Option

● None

○ At Chance Node

Figure 2.15 Advanced User

2.3.3 Running Model and Viewing Results

The Analysis menu allows the user to run the model and view output reports. The analysis menu is shown in Figure 2.16. The *Analysis* menu options are described next.

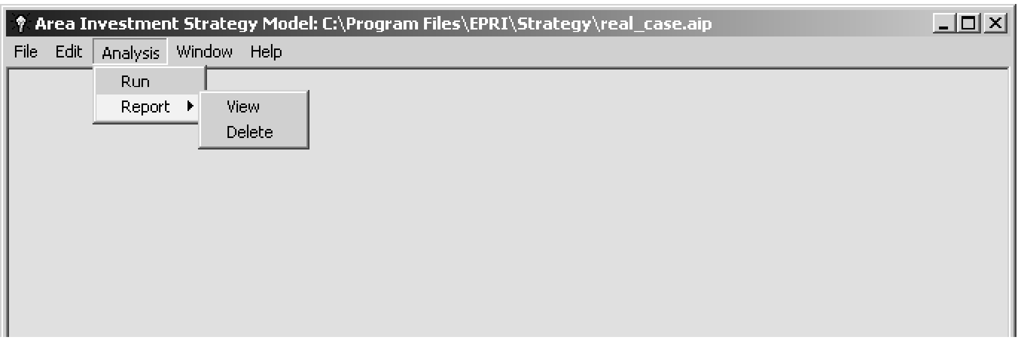

Figure 2.16 Analysis Menu

Run

If the user selects run form the Analysis menu, the model is run and the cursor changes to an hourglass. As soon as the calculations are complete a message box is displayed followed by a summary output dialog box. This summary output is shown below, Figure 2.17.

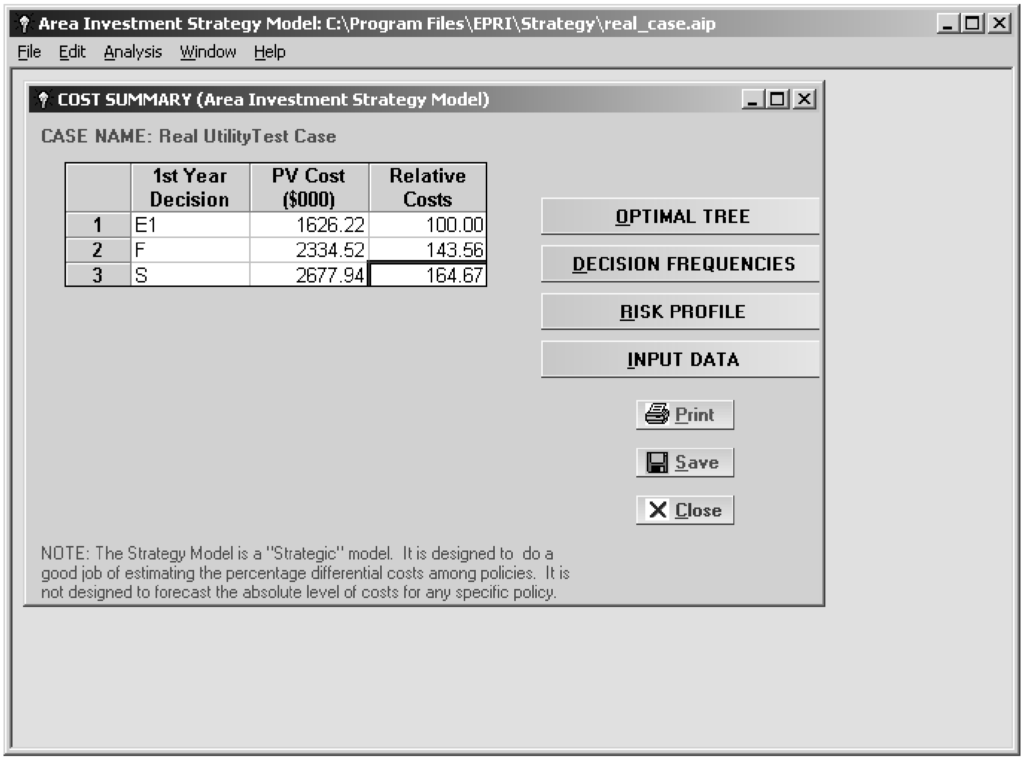

CASE NAME: Real UtilityTest Case

	1st Year Decision	PV Cost ($000)	Relative Costs
1	E1	1626.22	100.00
2	F	2334.52	143.56
3	S	2677.94	164.67

NOTE: The Strategy Model is a "Strategic" model. It is designed to do a good job of estimating the percentage differential costs among policies. It is not designed to forecast the absolute level of costs for any specific policy.

Figure 2.17 Summary Output

This dialog box shows present value of costs for the alternatives that could have been used as the first year capacity decision. These are ordered from low cost to high cost. The first item in the table is the optimal (least cost) first year decision and is the basis for the optimal strategy. From this form, the user can select four output reports. These four additional reports can be viewed by clicking on the buttons on the Summary Output window. The four reports are described below.

Optimal Tree Report

If the user clicks on the button "Optimal Tree" the least-cost decision strategy is shown. This is illustrated in Figure 2.18. There are several tabs at the top of this report. The first tab, "Optimal Decisions," shows only the investment decisions over time. The other tabs show greater detail on costs and load uncertainty. The information contained the Optimal Tree Report is explained in greater detail in CHAPTER 4 and CHAPTER 5, (ANALYSIS TUTORIAL, and REAL UTILITY EXAMPLE).

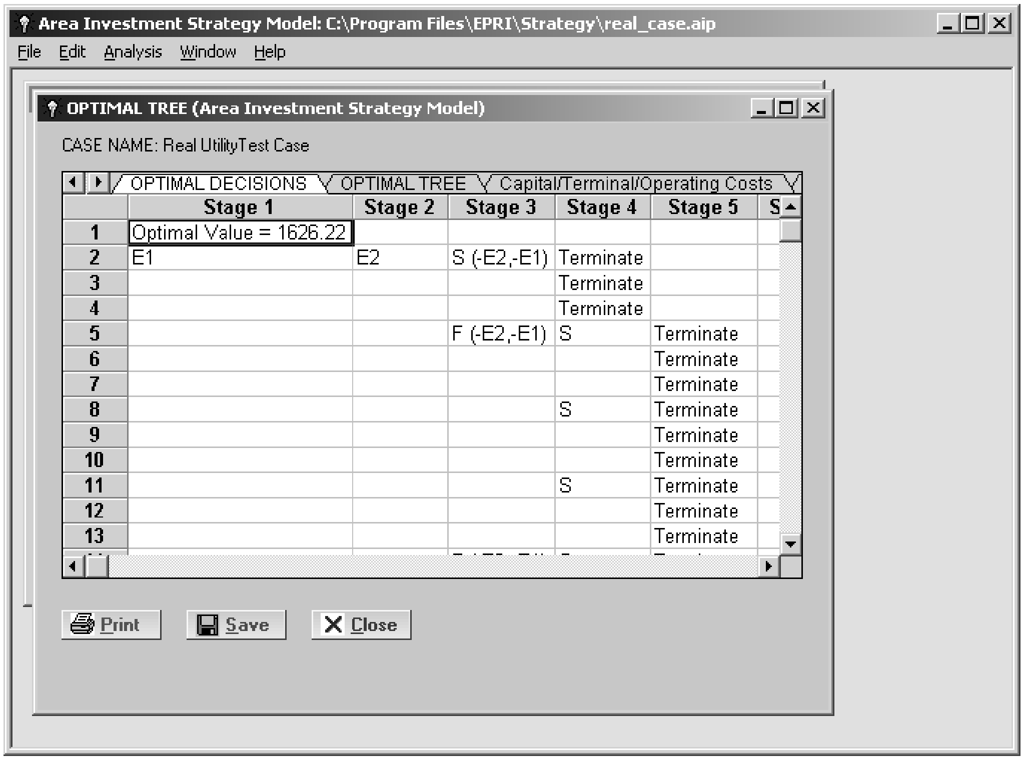

CASE NAME: Real UtilityTest Case

	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5
1	Optimal Value = 1626.22				
2	E1	E2	S (-E2,-E1)	Terminate	
3				Terminate	
4				Terminate	
5			F (-E2,-E1)	S	Terminate
6					Terminate
7					Terminate
8				S	Terminate
9					Terminate
10					Terminate
11				S	Terminate
12					Terminate
13					Terminate

Figure 2.18 Optimal Tree

Decision Frequencies

Figure 2.19 shows the output report that results when the user selects “Decision Frequencies.” This report shows for the years listed across the top, the proportion of time that each investment option is installed. For example:

1) at time 0-1 option E1 is chosen,

2) at time 2-3 option F is chosen 17% of the time and no investment is made 83% of the time,

3) at time 3-4, option S is selected 6% of the time and no investment is made the remaining 94%.

Multiple possible investments in future years reflects the fact that load is uncertain and can thus follow many possible growth trajectories. The Strategy Model enumerates the possible trajectories and identifies the least expected-cost investments conditional on current and future load growth.

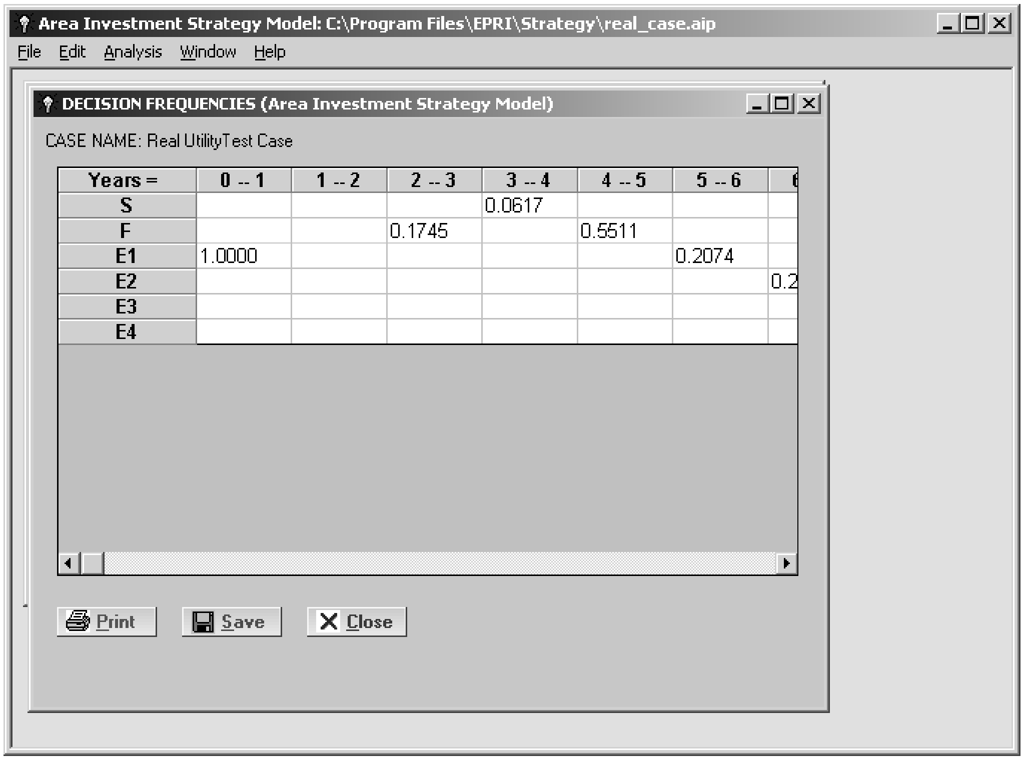

CASE NAME: Real UtilityTest Case

Years =	0 -- 1	1 -- 2	2 -- 3	3 -- 4	4 -- 5	5 -- 6	
S				0.0617			
F			0.1745		0.5511		
E1	1.0000					0.2074	
E2							0.2
E3							
E4							

Figure 2.19 Decision Frequencies

Risk Profile

Figure 2.20 shows the output report that results when the user selects "Risk Profile." This report shows the probability distribution on cost outcomes for the least cost investment policy. The first column gives cost-ranges and the second column give the probability of occurrence of these cost-ranges.

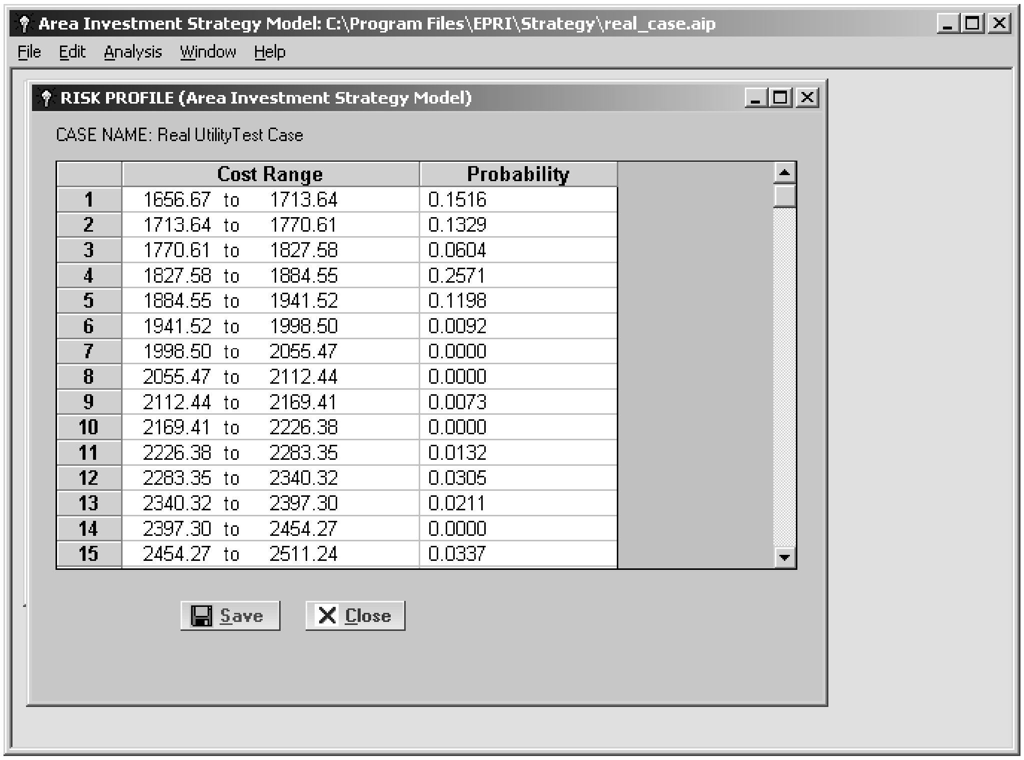

CASE NAME: Real UtilityTest Case

	Cost Range	Probability
1	1656.67 to 1713.64	0.1516
2	1713.64 to 1770.61	0.1329
3	1770.61 to 1827.58	0.0604
4	1827.58 to 1884.55	0.2571
5	1884.55 to 1941.52	0.1198
6	1941.52 to 1998.50	0.0092
7	1998.50 to 2055.47	0.0000
8	2055.47 to 2112.44	0.0000
9	2112.44 to 2169.41	0.0073
10	2169.41 to 2226.38	0.0000
11	2226.38 to 2283.35	0.0132
12	2283.35 to 2340.32	0.0305
13	2340.32 to 2397.30	0.0211
14	2397.30 to 2454.27	0.0000
15	2454.27 to 2511.24	0.0337

Figure 2.20 Risk Profile

Input Data

The final output report is obtained by clicking on the button labeled “Input Data.” That report lists the user supplied inputs. The report is not shown here. The user can print this report by selecting the print button. Due to a bug in a Visual Basic control, the user should always select the option to print in portrait mode. If they select landscape mode, data will be lost. As an alternative, the data can be copied into a word processor document and printed.

Report

On selecting the report menu option you have two choices: view or delete.

If Report View is selected, a standard windows dialog box allows the user to select output reports that were created by previous runs. All output reports have the form "*.Report". The user can also select Report Delete to delete existing report files.

2.3.4 HELP Menu

The help menu items are described below.

Contents

If the user selects contents from the help menu, the contents of the on-line manual are provided.

About

This provides copyright information and the version number of the model.

2.4 Using The Load Assessor Tool

2.4.1 Introduction

Load Assessor is designed to provide estimates of the load growth trends and the transition probabilities required by the Investment Strategy Model – the inputs shown in Figure 2.6. The estimating methodology requires data obtained from a set of questions about future potential load conditions. The approach and question were developed in collaboration with utility distribution planners. In developing the specific questions, the objective has been to base the methodology on a practical set of questions that the model user's can answer.

The purpose of the Load Assessor tool is to create a better base of load uncertainty information for developing and evaluating investment strategies. The tool provides probabilistic information of future load conditions. Two factors motivated the development of the approach for characterizing load uncertainty:

1) First, for distribution planning, a key issue is at what point in the future will load growth result in a given level of demand. This determines the time to the next investment decision, since distribution capacity requirements are based on load levels. Thus, a complete description of potential load trajectories over time is required in order to specify the time to the next decision.

2) Second, load growth follows trends. For example, area load typically might grow at a low, steady rate for several years and then transition to rapid growth for a time. This suggests that modeling future load conditions should start with characterizing key parameters associated with the possible trends -- the average duration's of the trends, and the likelihood's of shifting among the trends.

For a more detailed discussion of the load uncertainty methodology see Section 6.5. The user may also want to read the LoadDynamics™ User's Manual. LoadDynamics™ is a standalone load forecasting model that incorporates the same load uncertainty methodology as the Area Investment Strategy Model. The Load Assessor tool is also built into the LoadDynamics™ Model

---

2.4.2 Using The Load Assessor User Interface

When the load assessor starts the window below (Figure 2.21) replaces the investment strategy main window. The user must then use the File Menu to open a new or existing case. All Load Assessor files are of the form **.drs*. *assessor.drs* is the example case distributed with the software. Note that the File menu also allows the user to save load assessor cases and to copy load assessor results directly to the Strategy Model input screen, "Copy Results to Growth Scenarios."

Note: The data in the input screens shown here are for the data file *assessor.drs*.

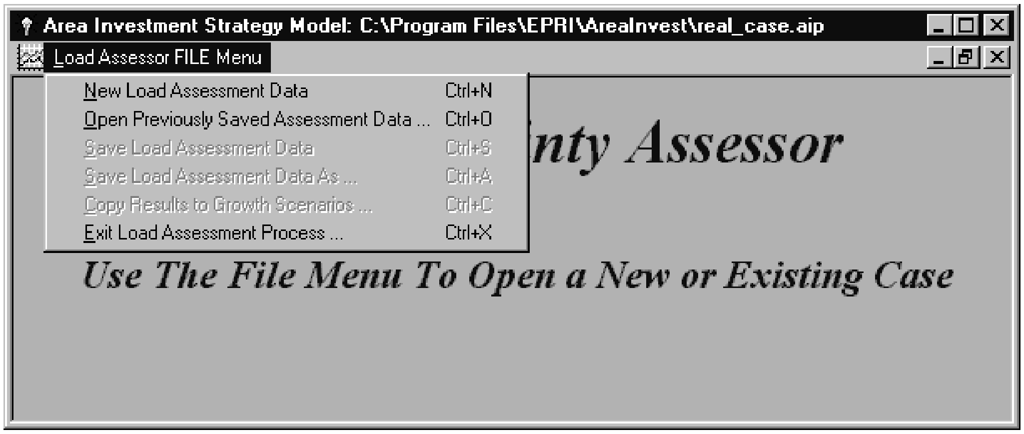

Use The File Menu To Open a New or Existing Case

Figure 2.21 Load Assessor Main Screen

Home Screen

As soon as a case is opened, a series of steps are listed in the home window along with the status of completion for each step (To Do, Done, Data Error, Incomplete). Figure 2.22 shows the Home Screen for an existing case where all of the data have been input, and the model parameters estimated. There are five steps in the assessment process:

1) Specify current load and growth

2) Specify load growth scenarios

3) Specify growth rate holding times
4) Specify maximum area load potential
5) Calculate load model parameters

When a case is opened, new or existing, the status of data inputs are indicated on the home screen. With a new case the user is constrained to proceed through the assessment steps in sequence – the model will not allow skipping a step and valid data must be provided for each step before the model will allow the user to go to the next step. The assessment software performs error checking for each screen.

When the user is satisfied with their load assessor results they can 1) save the results, and 2) copy the results to the Strategy Model "Load Growth Specifications, input form. Copying the data to Strategy Model inputs can be done via either the File menu or the button at the bottom of the home screen.

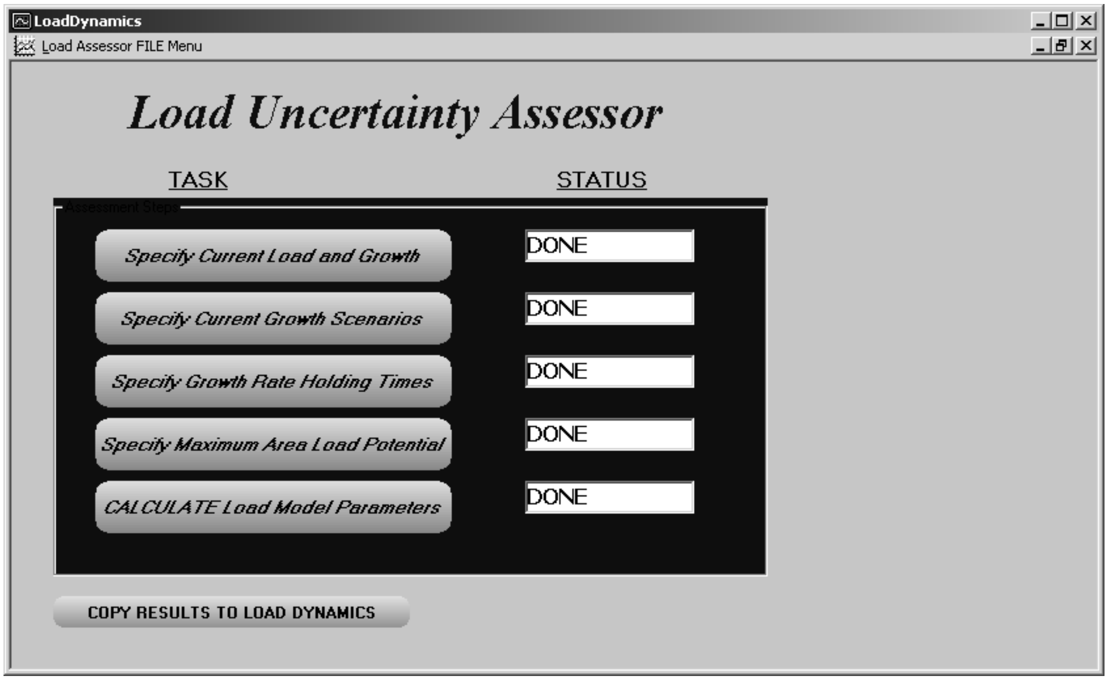

TASK	STATUS
Specify Current Load and Growth	DONE
Specify Current Growth Scenarios	DONE
Specify Growth Rate Holding Times	DONE
Specify Maximum Area Load Potential	DONE
CALCULATE Load Model Parameters	DONE

Figure 2.22 Load Assessor Home Screen

Specifying Current Load and Growth

The data screen for assessment step 1, Specify Current Load and Growth, is shown in Figure 2.23. Here the user must specify current area load and growth, and the time interval for the forecast assessment.

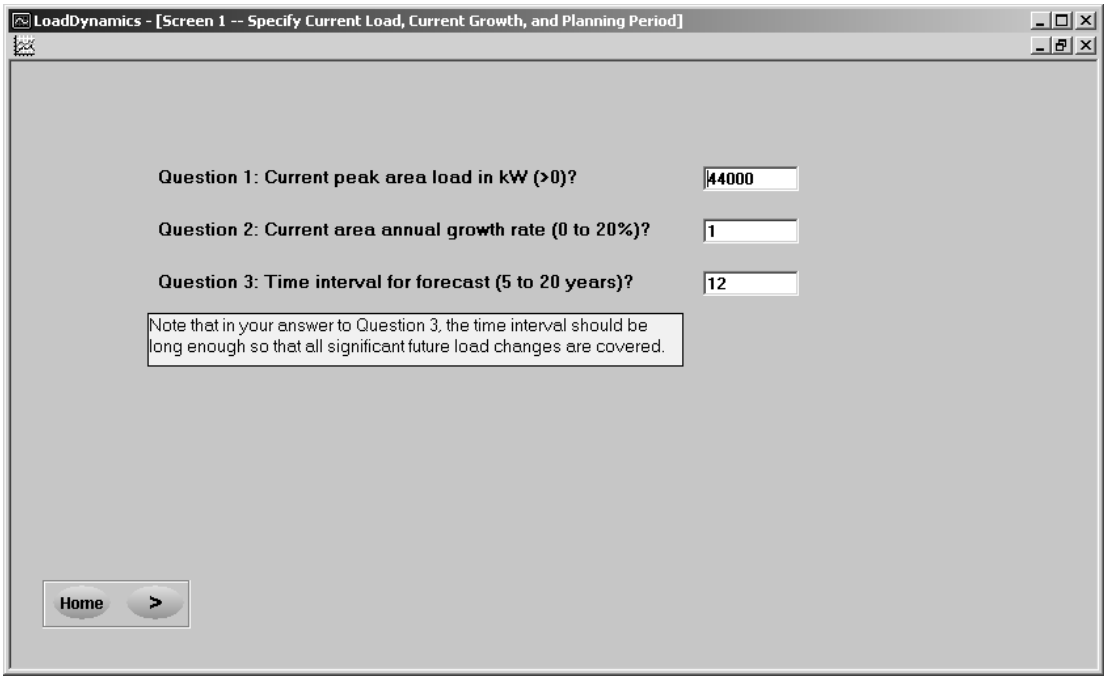

Question 1: Current peak area load in kW (>0)? 44000

Question 2: Current area annual growth rate (0 to 20%)? 1

Question 3: Time interval for forecast (5 to 20 years)? 12

Note that in your answer to Question 3, the time interval should be long enough so that all significant future load changes are covered.

Figure 2.23 Current Load & Growth

Specifying Load Growth Sceanrios

The data screen for the second assessment step, Specify Load Growth Scenarios, is shown in Figure 2.24. The purpose of the second screen is for describing potential future load growth scenarios. The basic idea is that the user must specify a set of alternative growth conditions that could occur over the planning period.

Before filling in the data required by this form, the user should read the instructions by pressing the *Instructions* button.

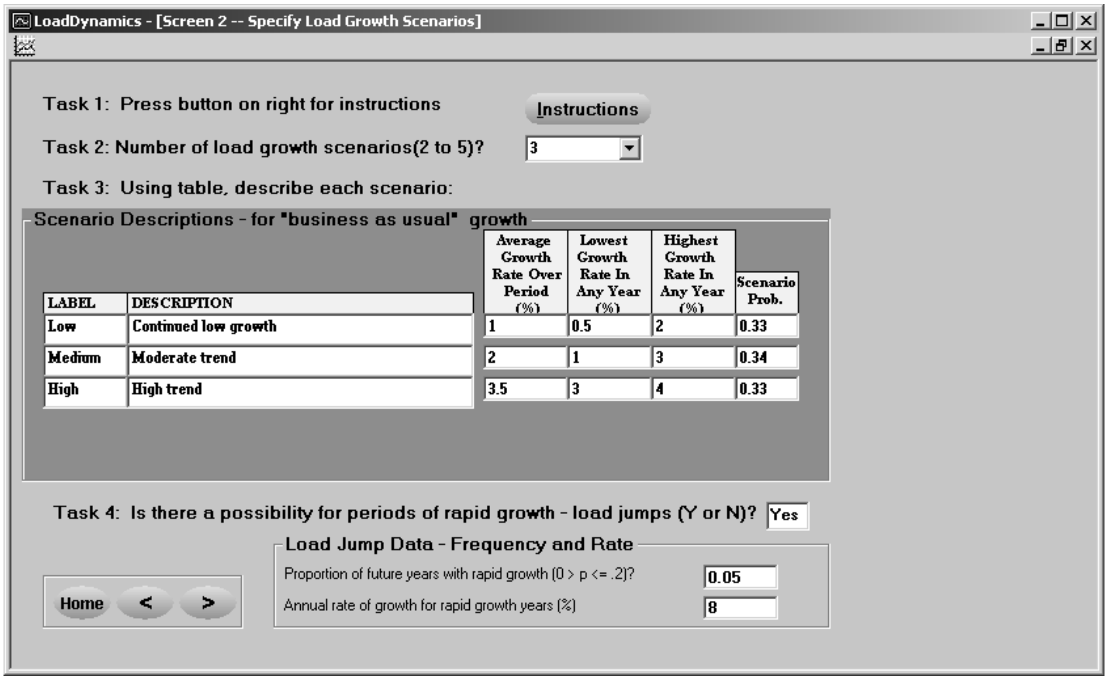

Task 1: Press button on right for instructions

Task 2: Number of load growth scenarios(2 to 5)? 3

Task 3: Using table, describe each scenario:

Scenario Descriptions - for "business as usual" growth

LABEL	DESCRIPTION	Average Growth Rate Over Period (%)	Lowest Growth Rate In Any Year (%)	Highest Growth Rate In Any Year (%)	Scenario Prob.
Low	Continued low growth	1	0.5	2	0.33
Medium	Moderate trend	2	1	3	0.34
High	High trend	3.5	3	4	0.33

Task 4: Is there a possibility for periods of rapid growth - load jumps (Y or N)? Yes

Load Jump Data - Frequency and Rate

Proportion of future years with rapid growth (0 > p <= .2)? 0.05

Annual rate of growth for rapid growth years (%) 8

Figure 2.24 Load Growth Scenarios

The data inputs for the scenarios are the following:

<u>Input</u>	<u>Explanation</u>
Number of Scenarios	The user must choose a number (2 to 5) from the drop-down box.
Label	A short label used to describe each scenario.
Description	A longer more descriptive characterization of each scenario.
Average Growth Rate Over Period (%)	For each scenario the user must specify the average growth rate that is expected for the entire forecast period.

Lowest Growth Rate In Any Year (%)	For each scenario the user must specify the lowest growth rate that can occur in any year for the forecast period. This growth rate must be lower than the average growth rate for the period.
Highest Growth Rate In Any Year (%)	For each scenario the user must specify the lowest highest growth rate that can occur in any year for the forecast period. This growth rate must be higher than the average rate for the period.
Scenario Probability	For each scenario the user must specify the probability that the scenario will occur. The probabilities for the scenarios must add to 1.0.
Is There a Possibility For Periods of Rapid Growth?	If there is a possibility of having short periods (one year or less) of rapid growth, enter Y.
Proportion of Future Years With Rapid Growth	If the user specifies Y to the question *Is there a possibility for periods of rapid growth*, they must indicate the proportion of future years where rapid growth is expected to occur (0 to 0.2)
Annual Rate of Growth For Rapid Growth Years	If the user specifies Y to the question *Is there a possibility for periods of rapid growth*, they must indicate the associated annual percentage growth rate (this rate must be at least twice as large as the highest rate entered in the ``highest growth rate in any year" column of the *Scenario Descriptions* table).

Specifying Growth Rate Holding Times

The data screen for the third assessment step, Specify Growth Rate Holding Times, is shown in Figure 2.25. Before entering any data, the user should click the *Instructions* button and get specific information about the data inputs.

This screen shows three one-year growth rates and the user must supply "average holding times for these rates. These rates come from user supplied inputs from the previous screens – the minimum and maximum one-year rates from the scenario inputs, and the current growth rate from the first input screen.

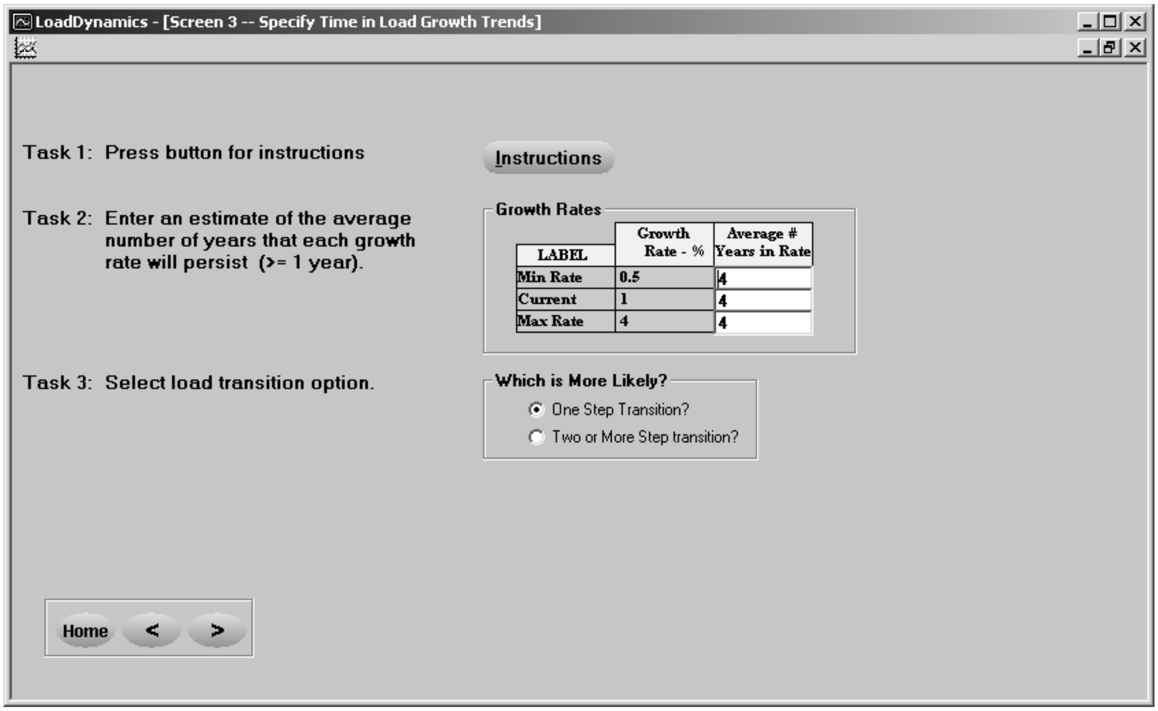

LoadDynamics - [Screen 3 -- Specify Time in Load Growth Trends]

Task 1: Press button for instructions

Task 2: Enter an estimate of the average number of years that each growth rate will persist (>= 1 year).

Task 3: Select load transition option.

Growth Rates

LABEL	Growth Rate - %	Average # Years in Rate
Min Rate	0.5	4
Current	1	4
Max Rate	4	4

Which is More Likely?

- ● One Step Transition?
- ○ Two or More Step transition?

Figure 2.25 Growth Rate Holding Times

The user must provide the following information:

Input	Explanation
Average Years in the Rate	This is what we have termed "holding time." Given that load is growing at a rate, this is the time that load is expected to continue to grow at that rate. This measures the tendency for load to stay in a particular growth trend.
Load Transition Option	When load growth shifts from one trend to another, the user must specify which of two conditions is more likely. (1). A transition of one trend-step versus, (2), a transition of two or more trend-steps. For example is it more likely that load will shift from a low rate to a medium rate or from low to a high rate, bypassing the medium rate?

Specifying Maximum Area Load Potential

The data screen for the forth assessment step, Specify Maximum Area Potential Load, is shown in Figure 2.26. Before entering any data, the user should click the *Instructions* button and get specific information about the data inputs.

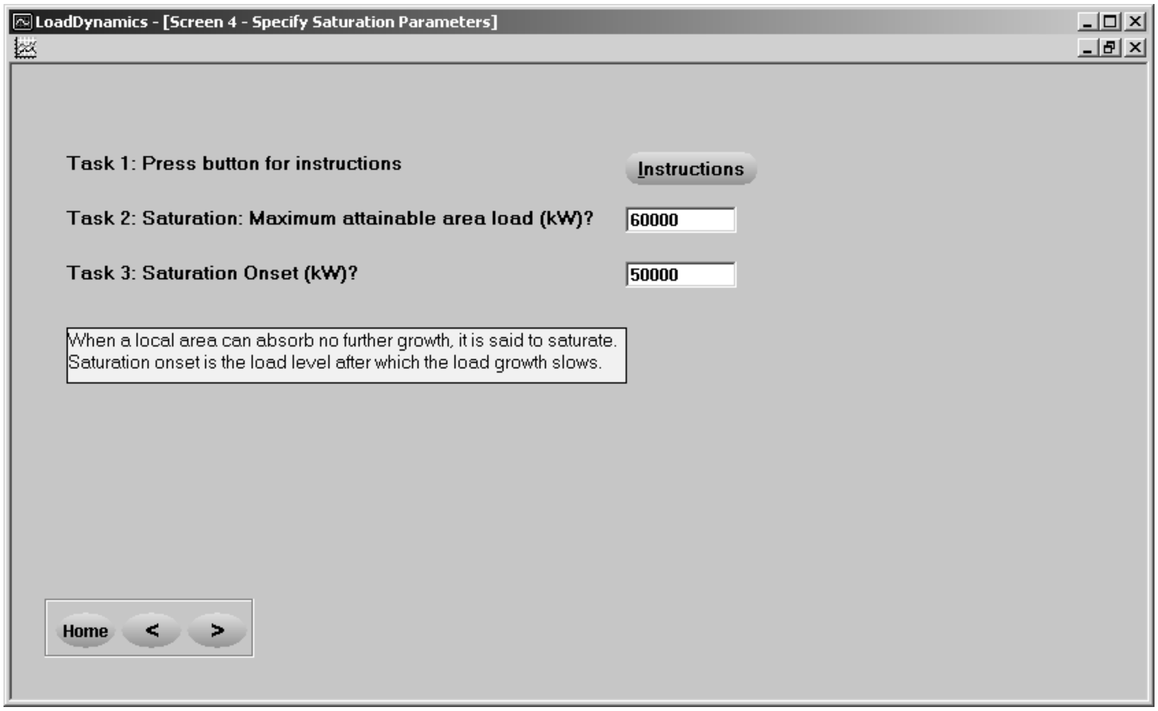

Task 1: Press button for instructions

Task 2: Saturation: Maximum attainable area load (kW)? 60000

Task 3: Saturation Onset (kW)? 50000

When a local area can absorb no further growth, it is said to saturate. Saturation onset is the load level after which the load growth slows.

Figure 2.26 Maximum Area Potential Load

The user must provide the following information:

Input	Explanation
Maximum Attainable Area Load (kW)	This is the load level when the area is completely built-out. It is when the area can absorb no further growth.
Saturation Onset	When saturation begins to occur, load growth begins to slow. Saturation onset is the load level after which load growth slows.

Calculate Load Model Parameters

The data screen for the final assessment step, Calculate Load Model Parameters, is shown in Figure 2.26. This screen shows the results after the *Estimate Parameters* button has been clicked.

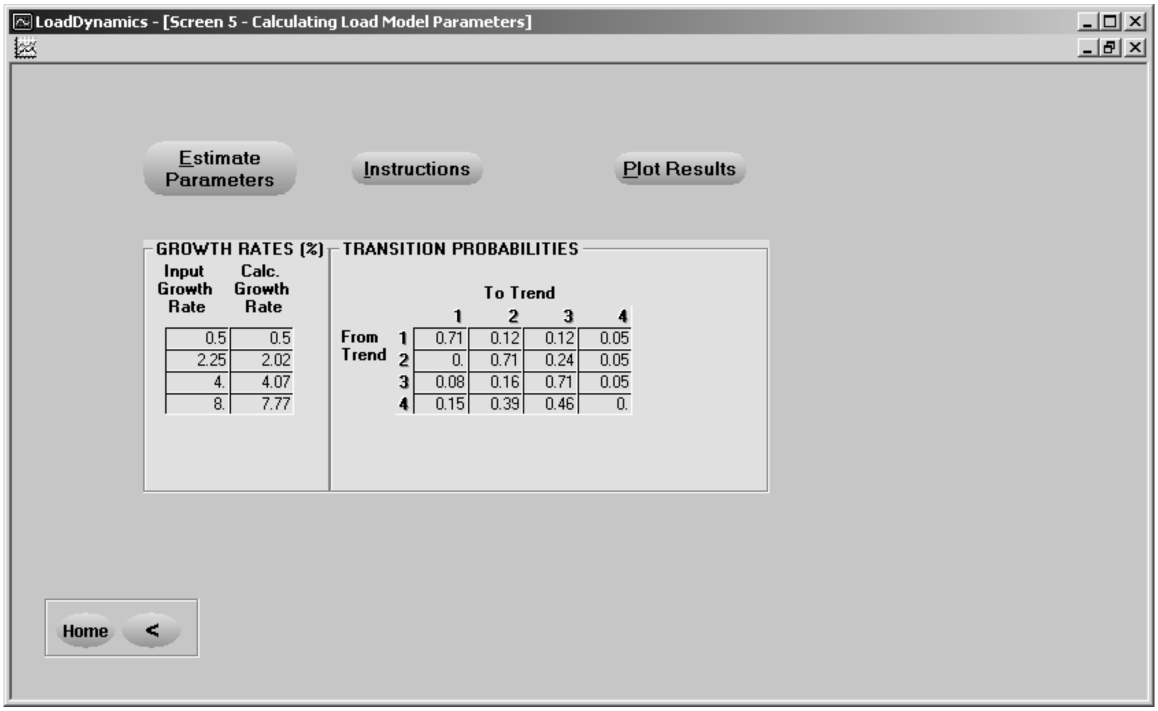

GROWTH RATES (%)

Input Growth Rate	Calc. Growth Rate
0.5	0.5
2.25	2.02
4.	4.07
8.	7.77

TRANSITION PROBABILITIES

From Trend	To Trend 1	2	3	4
1	0.71	0.12	0.12	0.05
2	0.	0.71	0.24	0.05
3	0.08	0.16	0.71	0.05
4	0.15	0.39	0.46	0.

Figure 2.27 Calculate Load Model Parameters

The functionality of the form for calculating the load model parameters is summarized below:

Button	Explanation
Estimate Parameters	If this button is clicked, the estimation model is run to produce a new set of growth trends and transition probabilities. If a new or modified data exist in any of the screens, the transition matrix will not be shown until the model is run.

Plot Results	If this button is clicked a plot of the probabilistic load forecast is produced. This plot is shown below -- Figure 2.28
Instructions	This provides information on how the growth trends and transition probabilities are estimated.

View Probabilistic Forecast

If the user clicks the *Plot* button in the *Calculate Parameters* screen, the probabilistic forecast is plotted. This is shown in Figure 2.28. For each year in the planning period, the maximum, minimum and average load are plotted.

If this load forecast does not accurately reflect the user's assessment of the possible loads for the area, the *Options* button should be clicked. Clicking the *Options* button brings up a dialog box with the following information:

1) If the upper bound on load is too high (low), then you should try reducing (increasing) the average growth rate for the high growth scenario. You can also change the maximum possible load level, saturation level.

2) If the lower bound on load is too low (high), then you should try increasing (reducing) the average growth rate for the low growth scenario.

3) If the expected value for load is too high (low) then you can reduce (increase) the probabilities for the higher growth scenarios and / or increase (decrease) the probabilities for the lower growth scenarios. You can also change the highest yearly growth rate for the high growth scenario.

4) If load is growing too slowly (too rapidly), you should try increasing (reducing) the highest yearly growth rate for the high growth scenario.

This completes Chapter 2, Model User's Guide.

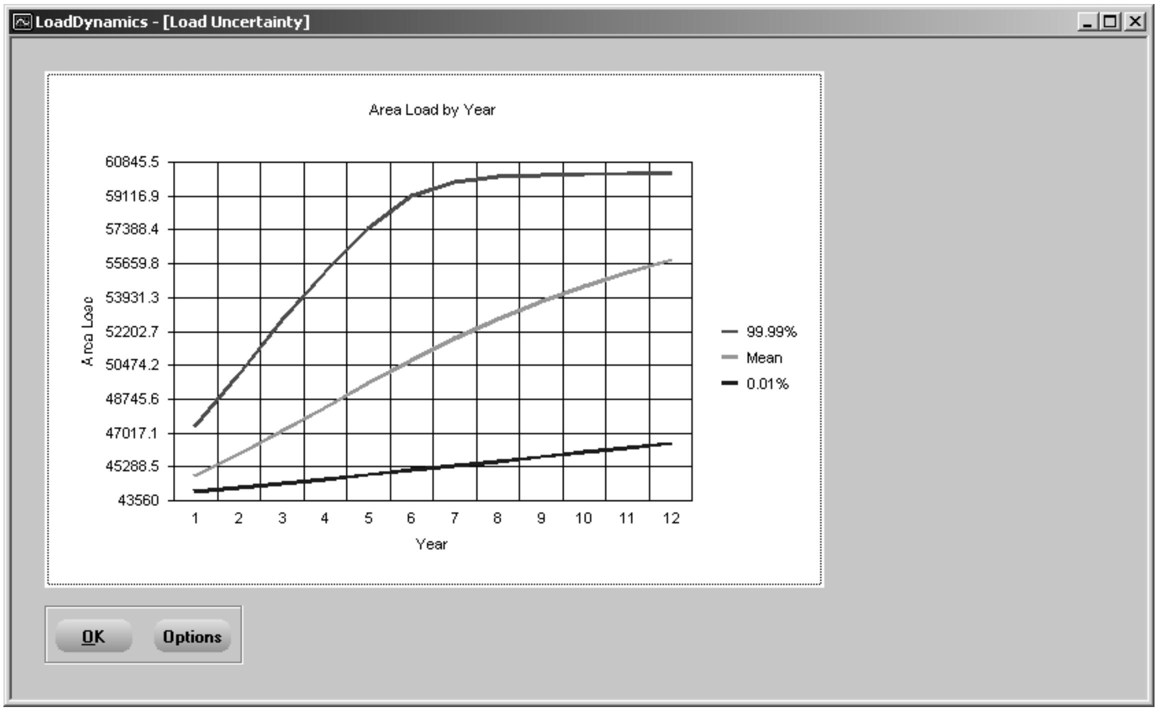

Figure 2.28 Plot of Probabilistic Forecast

CHAPTER 2 MODEL USER'S GUIDE

CHAPTER 3 BASIC CONCEPTS

3.1 Introduction to Strategic Planning for Distribution Investments

This section outlines the basic principles underlying strategic planning for distribution infrastructure investments. The target audience for this section includes the planners that operate the computer model and perform the analysis, the decision makers who select and implement investment policies, and the regulators who review policies and performance. For the Area Investment Strategy Model to be used effectively, all three of these groups need to understand the basic principles.

Four subsections follow. The first subsection discusses the nature of the investment strategy problem. The second subsection describes how the Area Investment Strategy Model relates to more traditional capacity planning methods. The third subsection describes the necessary reformulation the objective and the analysis methods The forth subsection summarizes the analytical dimensions of the problem.

3.1.1 The Investment Strategy Problem

Distribution assets comprise over 40 percent of total electric utility investment. Although it would appear natural for distribution planning to focus on long-term investment issues, at present a distribution planner's main objective is maximizing system reliability within a fixed budget. Distribution planning has traditionally focused on meeting near-term customer demand by adding transformers and feeders. Little attention is given to alternative technical solutions, the appropriate scale for investments, or the long-term implications of resource commitments.

Maximizing reliability within a fixed budget had been achieved through application of planning criteria rules of thumb that minimize the effects of delivery problems and unexpected increases in demand (supply and demand disruptions). During normal conditions, outages have not been acceptable. During disruptions (due to equipment failures or rare high demand levels), application of the rules of thumb have limited the outages to a few hours per year.

This purely technical infrastructure planning approach was successful in a particular business environment. That environment no longer exists. We note the following changes:

- Industry restructuring, already under way, is bringing about the disaggregation of the vertically integrated electric utility.
- For many companies, the funds allocated to the distribution side of the business have been reduced. Distribution engineers are being forced to build business cases around all significant projects.
- New customer needs, which vary by market segment, are beginning to emerge. In particular, an increasing number of customers are demanding higher reliability and improved power quality.
- Technology advances in information processing, communications, power electronics, materials, and distributed generation are providing new distribution investment options for meeting customer needs.
- The industry is increasingly affected by uncertainties in load growth and resource costs.

3.1.2 An Alternative to Traditional Capacity Planning

As mentioned above, traditional electric distribution infrastructure planning addressed short-term problems with installations having long-term implications. Responding to this imbalance, recent methods have taken a longer-term perspective, in which the planner is asked to specify the future load growth and capacity plans in order to determine the impact of deferring the plan using smaller, modular technologies (distributed generation and load control programs). While this type of approach does address one important piece of the problem – the long-term implications of large investments – it neglects another, equally important, piece, which is the uncertainty that accompanies any attempt to project more than a year or two into the future.

The recent approaches are sometimes characterized as *scenario planning* because they are based on one or a few selected snapshots or scenarios of the future in which all future variables are assumed to be known and fixed over the planning period. Since the future is uncertain, this type of deterministic planning does not reflect the actual utility decision environment and process. Thus, it is a very poor basis for developing investment strategy.

In the scenario-planning approach, current investment alternatives are evaluated based on their impact on some fixed future installation strategy. To define this strategy, the planner must fully specify the type and installation dates of future T&D, generation, and distributed resource investments. This, of course, requires specification of a fixed future load pattern over the planning horizon, which is typically 10 to 20 years long. Current decisions are then compared with respect to how they affect the timing and cash flow of the pre-specified installation strategy. For example, the decision to install a primary feeder at the beginning of the planning period is evaluated in terms of its impact on delaying future capital investment in a substation.

Not only is the assumption of future certainty not realistic over any reasonable planning horizon, but it places an artificial and uncomfortable burden on the planner. We have found that planners are extremely reluctant to specify a future investment plan when they know that such a plan is only one of perhaps hundreds of future realistic possibilities. The actual future investment plan will depend on how load evolves and on the resolution of other uncertainties, such as siting constraints, regulatory policies, environmental impacts and the structure of the business environment.

In such situations, the decision process must be dynamic and responsive to change. Good decision makers understand this well. They know that they must decide what to do today, and as the future evolves, they will have to react to the changing situation and make further decisions and commitments.

In order to be useful, planning models must capture the essence of the actual decision environment and process. The Area Investment Strategy Model not only identifies what action to take today, but also how to respond as key uncertainties evolve. The approach reflects the dynamic nature of the actual decision process that adjusts investment strategy when the future needs and costs are revealed. It simultaneously determines the best policy and the cost of that flexible policy. Moreover, the Area Investment Strategy Model does not require prior specification of a future plan. The value of an investment decision is not measured by a so-called “deferral credit” based on an arbitrary future investment. Instead, only actual cash flows, for both capital and operating costs, are used to determine the best strategy.

The model can be used to perform traditional deterministic scenario analysis as well as the more appropriate analysis under uncertainty. In 5 we provide an example in which the model is used to compare the results of scenario analysis to the results of a full analysis under uncertainty. The example demonstrates that addressing issues such as management flexibility and learning can fundamentally affect policy. Scenario analysis will often produce the wrong policy and, in so doing, incorrectly estimate the costs of all policies.

3.1.3 Reformulation of Distribution Investment Planning

The new approach to investment planning requires reformulation of the investment objective and the analysis methods. The key aspects of this reformulation are described below.

Structure the Problem Recognizing The System Nature Of Distribution Planning

Electric distribution systems are complex structures. They are composed of interconnected transmission and distribution circuits and associated transformers, switches and relays. Because the systems are composed of interconnected components, identifying alternatives for increasing capacity and expanding the systems over time requires identifying specific load-growth locations and the bottlenecks in the system that limit capacity.

Problem structuring involves specifying the capacity problem being solved and identifying feasible alternative solutions. The inputs to the problem include the load growth that will be experienced, the current structure of the distribution system, and the alternatives available for installation. The constraints include the capacity bottlenecks. The solutions are sequences of investments that provide the required capacity as load develops. Our experience is that problem structuring is fundamentally an issue of doing careful thinking about the problem and far less a matter of doing engineering analysis and running load flow models.

Remove Deferral Bias

The objective of distribution investments is the same as that of any other investment in infrastructure: minimize the cost of service subject to meeting appropriate reliability and obligation to serve constraints. Meeting this objective may entail the delay of more traditional infrastructure investments, but such delay is in no way the investment objective.

The EPRI approach to distribution investment planning is designed to remove the deferral bias that is present in the current view of using small, modular investments (for example, DR or distributed resources) to defer larger infrastructure capacity investments. Instead of selecting modular investments to defer planned capacity installations, the new approach selects the smaller investments to *minimize* the cost of service. This removes the deferral bias. In order to accomplish this selection, the investment strategy problem is formulated as an optimization problem. Given that formulation, economic analysis correctly identifies the least-cost strategy., Problem formulation is fundamental. If the problem is formulated incorrectly, the wrong answer will be found, or at least the best solution will be found only by accident.

Base Calculations on Actual Cash Flows

The first change in economic analysis is that the actual cash flows, capital plus operating, are used in the analysis. One should eliminate the use of arbitrary marginal avoided costs wherever possible.

Select Modular Investments to Hedge Against Future Uncertainties

The main benefit of modular investments (portable substations, engines, and other small options) is their ability to delay large infrastructure investments until the need for those latter investments is more clearly understood. Indeed, it is possible that the need will never materialize, but it is almost certainly not true that their main benefit comes from deferring indefinitely, thus replacing, substations or feeders. Modular investments are not a substitute for infrastructure investments. Their real value is that they permit the planner to wait, or fill a gap, without making a risky, large, long-lived capital commitment.

This objective is not inconsistent with minimizing the cost of service. In fact, in the context of finding the optimal investment policy, the hedging possibilities are only selected when they provide least-cost solutions.

The physical nature of the modular investments is what permits them to be used as hedges and what prevents them from being used as substitutes for infrastructure. By definition modular investments come in relatively small capacity increments. They are flexible, since they can be relatively easily sited and, in many cases, moved from one location to another in the distribution system.

The value of modular investments as hedging investments can be found by specifying the uncertainties in a distribution area and finding the best investment policy. The EPRI-supported analysis is designed to do that. It is important to note that the value of modularity increases as the uncertainties increase. The optimal investment policy is less likely to include modular options if there is relatively little uncertainty. This is because the additional cost of modularity and flexibility is not justifiable if the future can be treated as if it were known.

Explicitly Incorporate Uncertainty Into The Analysis:

Load Growth Uncertainty: For distribution planning, a key issue is at what point in the future will load growth result in new capacity requirements. A complete description of potential load trajectories over time is required in order to specify the probability distribution on the time that new capacity is required.

The uncertain load growth should be characterized dynamically. Most distribution planners believe that load growth can be usefully described in terms of multiple possible trends that can persist for uncertain durations. A simple model, such as a linear fit to load data, misses the dynamic behavior of the trends. Further, it is essential that the variance in the load on the distribution system not be underestimated. Again, a simple model, such as a linear fit to observed load data, will consistently underestimate the variance in the load compared with a dynamic probabilistic model.

A successful dynamic model will have minimal data requirements while preserving the appropriate degree of accuracy and variation that a load forecast ought to exhibit. The inputs to the model should be based on user assessments. The assessment procedure should be simple and transparent, which means that the user should find it both easy to use and easy to understand.

Siting Uncertainty: Among the important uncertainties that must be addressed by distribution planners is siting uncertainty. A critical issue is whether sites will be available for infrastructure investments. The relative likelihood of site availability for substations, feeders, and other investments is a strategic issue. Whether or not a site must be secured before it is actually needed becomes a strategic decision, depending on whether it is easy or not to obtain a site. Further, the effects of regulatory and environmental considerations must be addressed since these effects can influence investment strategy. There are several ways to address these issues. Perhaps the simplest is to permit the probability of obtaining a site to vary with time. Another approach is to translate site uncertainty into a cost of delay in obtaining a site.

Operating Cost Uncertainty: Operating costs of various technologies is an effect that must be considered by planners. For example, operating local generation technologies in uncertain conditions makes them somewhat risky. To the extent that different fuels are used by alternate technologies, the differential amounts of uncertainty in future prices can have an effect on choice of investments. The simplest approach to incorporating fuel price uncertainty is to assign a probability distribution to the future cost of fuels. Perhaps the expected value and the variance would be sufficient for most applications.

Technology Uncertainty: The uncertainty in technology, either in what is available or what it may cost in the future, is an important aspect of the strategic assessment of

DR. For example, although the current capital costs of both photovoltaics and fuel cells are sufficiently large so that neither of these technologies can be generally adopted, the future situation may be different. The consequence of using current DR technologies to delay while waiting for some cost reductions to occur in these other technologies can be discovered.

User-specified Uncertainties: In general, any user-specified uncertainty should be brought into the analysis. Some cases require modeling the effect of weather on the demand for energy. In other cases, the planners may need to analyze the effects of a single customer's decisions, about which they are uncertain. Although the fundamental uncertainties are those listed above, others are worth analyzing if the planner believes that such uncertainties have important effects on system performance and requirements. Any modeling approach or analysis technique must be flexible enough to permit such considerations to be incorporated.

Integrate Modular Investments Into Existing Expansion Plans:

If there is an existing plan for a local area, it will be of some interest to see whether integrating modular alternatives can yield a modified plan that is superior. A typical approach to determine how best to integrate modular considerations would be to consider whether it would be least-cost (or lower cost) to delay the investments in the plan. Some delicacy is required in such an instance, however. It is important to recognize that deferring assets may be beneficial but need not be optimal. Although improvements to existing plans may be achieved, there is no theory that indicates that such improved plans are the best that can be found.

Since many distribution areas do not have a plan, and since distribution investing appears to be reactive, it is more likely that the modular investments would be part of an evolving expansion plan. The EPRI approach to this problem suggests two innovations. First, the distribution system should be recognized as a subsystem that requires its own planning function. This consideration applies not only to the integrated utility but also to any other business organizations that may emerge in future. As the industry restructures and if the "poles and wires business" emerges as a separate entity, such recognition would be a natural result. Second, planners and other utility (or distribution company) managers ought to recognize that a plan is a policy that describes not just what to do next year but also what to do as the need for capacity evolves over time.

Identify Dynamically Optimal Investment Strategy:

The investment strategy that best utilizes modular investments is a dynamically optimal strategy that minimizes the expected net present value of the costs of acquiring and operating the investments over their useful lifetimes. The dynamically optimal strategy is a plan of what to do currently and what to do as uncertainties resolve over time. While the main focus may be on finding an immediate solution to a capacity problem, the least-cost near-term decision cannot be found without taking into account future possible conditions and decisions. Special modeling approaches and algorithms must be applied in order to determine such a strategy. These approaches have been built into the Area Investment Strategy Model.

3.1.4 The Analytical Problem

This subsection describes the economic valuation principles and the analytical questions that must be addressed when solving the investment strategy problem. The analytical questions are: (1) scale economies of the investment alternatives, (2) limited scope for modular investments (distributed resources and load control programs), (3) uncertain future load and costs, and (4) the need to include load control programs as part of the portfolio of investment alternatives.

Investment Valuation Principles:

Investment capacity planning is guided by a set of economic valuation principles. Four key principles are stated below. These underlie value estimation for the new DR and distribution planning methodology.

1. Deferral of capital projects has direct economic value. This is the result of the opportunity cost of financial capital – the discount rate when doing net present value calculations. The higher the cost of financial capital the greater the value of deferral. If the cost of capital were zero there would be no pure capital cost savings due to deferral. Large investments providing over capacity in the near-term can be good investments. The value of large investments depends on the cost of alternative smaller investments and on the eventual load – they can be ok if the capacity is eventually needed.

2. There is a tradeoff between economy of scale and flexibility. Big resources are generally cheaper per unit capacity but provide limited future decision flexibility. Small investments defer big investments and provide the option to revisit the big decision. This option to delay allows for the potential for learning before deciding.

3. The value of being able to revisit a large investment decision depends on the nature of load uncertainty. If there is no uncertainty, there is no potential for learning and no value associated with revisiting the decision. Even with

uncertainty, if the uncertainty is not reduced over time, there is no value associated with delaying to revisit the decision.

4. Independent of uncertainty, modularity has value. Small increments of capacity track load more closely and can be easier to site.

Scale Economies - The "Two-Edged Sword:"

In most cases, there are scale economies associated with larger capacity investments. If this were not so, the solution to the investment problem would be obvious: make small capacity investments that match the load as closely as possible. Scale economies suggest that it may make sense to invest in a large increase in delivery capacity if that capacity will eventually be needed. Hence, the main question is when to invest in the large-scale technologies. Two kinds of errors can result from an incorrect analysis of the benefits of economics of scale. The first error is to avoid a big investment because the first cost is large but, in fact, the big investment is the least cost strategy. The second error occurs if a big investment is made because the unit cost is relatively small but that investment, when correctly evaluated, is not the least cost choice.

Limitation of Scope - Modular Investments & DSM:

Interest in incremental investment strategies is increasing. Planners are starting to recognize that local, small-scale generation and load management options, also known as distributed resources, can provide at least a temporary alternative to traditional upgrades. The investment question is under what conditions should traditional upgrades be avoided by exercising modular options to meet uncertain customer needs for reliable and economic service? The answer to this question must be based on two facts. First, the marginal cost of those investments tends to increase, such that if modular investments are pursued aggressively, the cost of providing a fixed increment of capacity increases. This is illustrated in Figure 3.1.

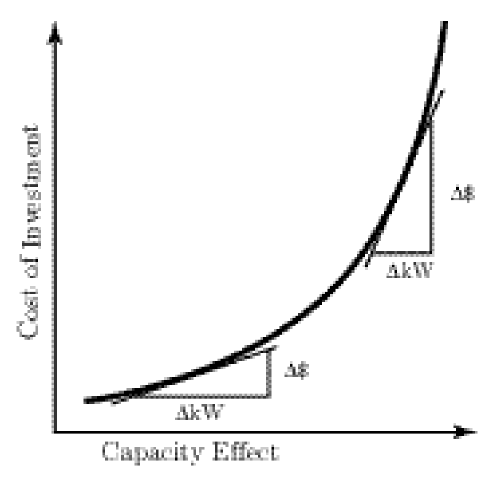

Figure 3.1 Limitation of Scope

Second, there is a physical limitation on the effect of modular and DSM investments. It may not be possible to locate the distributed assets where they are needed. It is possible that the capacity reduction effect of these investments will be limited. This situation is illustrated in Figure 3.2.

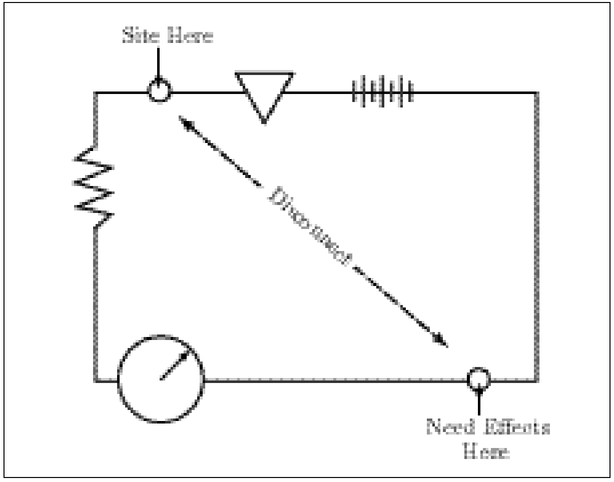

Figure 3.2 Physical Limit on Capacity

Uncertainty in Value of Investments:

The need for investments, and thus their value, can be very uncertain; current assessment of value depends on predictions of load, costs, and technology far into the future. In most cases, future load is an important uncertainty. Distribution infrastructure investments tend to come in large sizes, and once installed, last for a very long time. Because of their longevity, the investments are exposed to the risk that the forecasts of load, upon which the acquisition decisions were based, will not be accurate.

Load uncertainty makes the scale of distribution investments an important strategy issue. Smaller scale units provide a hedge against the risk that loads do not materialize. Further, because learning may occur over time, small-scale investments provide the opportunity to delay and revisit large-scale investment decisions. However if load does eventually materialize or if the cost of the infrastructure increases as areas develop, there may be significant cost penalties associated with deferring distribution upgrades using small-scale distributed technologies. Figure 3.3 suggests that the preferred choice between a Large or Small capacity investment depends on the forecast of load growth. For slowly growing load, the lower trajectory, a series of small capacity increments is best. If load is forecast to grow rapidly, then the large capacity alternative is superior.

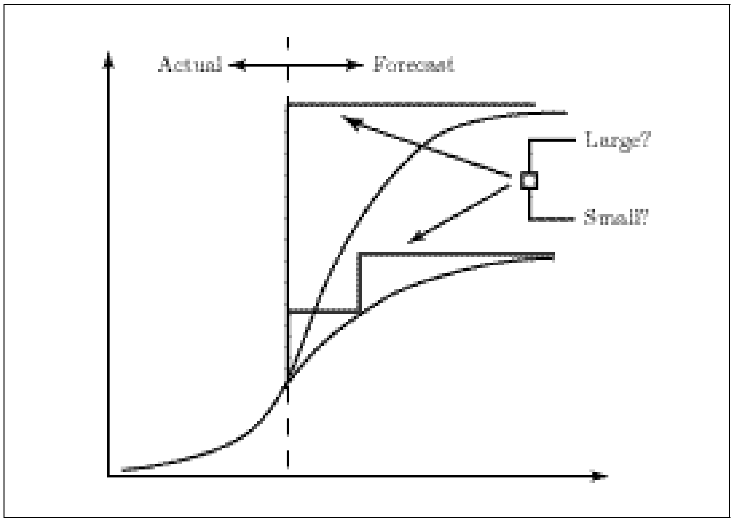

Figure 3.3 The Dilemma of Uncertainty

Load Management Programs

A fourth source of complexity is the existence of load management programs. System engineering combined with market intelligence can (1) establish some appropriate level of unplanned outages (unserved energy), and (2) identify investment alternatives that provide the commensurate level of service. It may also make sense to have selective outages for some customers (planned unserved energy). In some situations, such load management programs may be one of the more cost-effective alternatives for providing peak capacity relief. If this is true, these programs should be treated as an infrastructure investment alternative. Load management programs also give the planner the opportunity to fine-tune the level of service provided to individual customers as part of a complete investment strategy.

3.2 Model Description

EPRI, in conjunction with a number of utilities, has developed a new analytical approach for distribution infrastructure investment planning. The Area Investment Strategy Model is an integral part of this new approach. In this subsection we summarize the new approach and the associated analytical tools.

The new strategy development process builds on current reliability-focused planning. The approach is a two-step process. First, system engineering is used to establish the capacity of the existing local area distribution network and to identify alternative investment sequences that provide enough capacity to maintain a given service level. The Area Investment Strategy Model is then used to identify optimal sequences of investments that are conditional on future load conditions and other uncertainties. The sequences describe which investment to make today and how to respond as load and other uncertainties are resolved. The sequences are optimal since they provide the pre-defined level of service at least cost.

The Area Investment Strategy Model captures the uncertainty of future load and costs in the design of plans for infrastructure additions. It provides the cost and risk information utilities need to build business cases for least-cost investment strategies.

The model provides an analytic framework for users to specify alternatives, characterize important uncertainties, and develop strategies that minimize the expected cost of meeting future customer needs. The Area Investment Strategy Model is the only distribution-planning tool currently available that includes analysis of future load and cost uncertainties, and develops least-cost plans given the future uncertainties. The model allows joint consideration of T&D upgrades and distributed resources, including load management programs.

The Area Investment Strategy Model produces probabilistic descriptions of future load conditions, given a simple set of user inputs. This part of the model is also available on a stand-alone basis. Users interested in load forecasting can use the Load Dynamics Model. The load forecasting model provides: (1) the probability of being at a given load level for any year in the planning period, and (2) the probability distribution on the time that it will take load to reach a given level.

3.2.1 Purpose of the Area Investment Strategy Model

The purpose of the Area Investment Strategy Model is to enable the user to determine local distribution area expansion plans that are least-cost under uncertainty.

An expansion plan is a timed sequence of investment decisions that is contingent upon the future occurrence of various states of nature. For example, if load in an area were to begin to grow rapidly, the optimal expansion plan would recognize the need for capacity expansion in the area sooner than it would if load were growing less rapidly. The expansion plan itself is described in a contingent manner: e.g., install capacity amount A of type X at time t_1 if load growth is rapid, install capacity amount B of type Y at time t_2 if load growth is moderate, and install capacity amount C of type Z at time t_3 if load growth is slow.

Least-cost under uncertainty refers to the fact that the optimal expansion plan will be the one whose expected net present value is smallest compared to all feasible expansion plans considered in the analysis. Uncertainty means that the planning context is one in which the user’s lack of surety about future conditions of load growth, fuel costs, regulatory environment, technological progress, site availability, and possibly other important variables, is a major influence on the investment decisions actually made.

The appropriate investments when future load on the distribution system is certain are very different than those when future load is uncertain. It is that difference that motivated the construction of the strategy model. When future conditions are uncertain, then future cash flows are uncertain, and the optimal policy minimizes the expected present value over all future cash flows, where by expected we mean the probability-weighted average of the uncertain cash flows.

Unlike other models that rely on very complex and detailed data inputs, the strategy model requires relatively little in terms of data. This is because the design of the strategy model replaces the need for data with a sophisticated model structure that can extract, from the data, sufficient information to determine optimal investment strategies. The fundamental model design tradeoff is between a primitive model structure driven by exhaustive data sets compared to a mathematically robust model, with internal optimization logic, that requires an easily assembled data set.

3.2.2 Data Requirements

There are eight categories of data required in the strategy model.

- Basic Planning Data. The basic data includes (a) a specification of the planning period over which specific decisions will be made, measured in years; (b) the appropriate discount rate for future cash flows, expressed as a decimal; (c) the applicable inflation rate, expressed as a decimal, and (d) the tax rate, expressed as a decimal.

- Load Growth Specifications. The load growth is characterized using the parameters of the Load Model. These parameters include a collection of load growth rates, an initial peak load value, an initial load growth rate, and a matrix that captures the probability of making a transition between any two growth rates in each year of the planning period.

  The Load Model describes the dynamics of peak load in the distribution planning area by specifying the way load growth rates change over time. The model treats load growth as a random process such that each year a particular growth rate applies but the load growth for the next year is selected randomly, conditional on the current growth rate.

  As many as five distinct growth rates may be specified by the user. For each growth rate, the probability distribution on the next growth rate—the one that will apply in the next year—is required. This distribution can be thought of as a row in a matrix of probabilities. Each element in the matrix is the probability of making a transition from the initial growth rate (specified by the row) to the final growth rate (specified by the column). The sum of the probabilities in each row must equal one. The collection of rows defines the transition matrix.

  The initial load (kW) on the distribution system must also be reported. This provides the base for future growth. The initial load growth rate must also be provided by the user. This value must be bounded by the largest and smallest growth rates specified.

- Load Duration Curve. The load duration curve for the local distribution area is required. The curve is approximated by a piecewise linear function. The user provides up to 20 points on the curve, with coordinates (hour, % of peak load). As load grows, according to the Load Dynamics Model, the peak load is increased but the percentages are preserved, so that the load duration curve's relative shape is preserved.

- Capacity Alternatives. This data describes the alternatives available to a distribution planner to achieve capacity needs. The strategy model treats two kinds of alternatives, strategic alternatives and so-called "bandaids." The strategic alternatives are typically large-scale capacity additions that can be selected in any order along a trajectory. The bandaids have typically smaller capacities and can be selected in a fixed order. That order is implicitly specified by the user when the bandaids are entered into the model; i.e., the model assumes that the bandaids are entered in the sequence in which they will be installed. That sequence is independent of any strategic investments that may be installed along the trajectory.

  Typical strategic alternatives are feeders or substations. Capacitors, small engines and other distributed technologies such as batteries and DSM can be thought of as bandaids. The model is able to analyze distributed generation and storage located throughout the distribution system as either strategic alternatives or as bandaids. The user must specify a sequence of distributed alternatives if those alternatives are to be considered as bandaids.

  Each alternative must be named; the capacity (measured in kW) must be specified; the purchase cost (in thousands of dollars) and cost escalation (other than inflation, specified as a multiplier) are required; the lead time for each alternative, measured in years, is that time between commitment to install a new asset and when the capacity is actually available; the life is the useful life of the asset.

  It is important to note that the strategy model interprets the capacity value specified for each alternative as the peak-load-relieving capacity provided by the alternative. Capacity additions are required when peak load equals the total installed capacity. Therefore, when an investment is made, the total installed capacity increases by the capacity of that investment. Thus, each capacity addition satisfies load growth until the peak load increases by exactly the capacity that was added.

- Operating Costs. Each of the capacity alternatives described in 4, above, is also characterized with respect to operating costs. The model requires the fixed operating and maintenance (O&M) costs (measured in \$/yr.); appropriate fuel costs; the heat rate, as applicable, of each alternative; variable O&M, as appropriate; and the value, per kWh, of system energy cost for each of the alternatives, which is used as an energy credit if appropriate.

- Installation Constraints. This data describes restrictions on alternatives within a policy. A policy can be thought of as a sequence of investments; an investment provides capacity in response to load growth. Therefore, policies vary with respect to type of investments, ordering of investments as the policy extends over time, and timing of investments in response to load growth. The constraints restrict the possible sequences with respect to quantity and precedence of investment alternatives. The quantity constraints indicate how many times an alternative may be repeated within a policy; e.g., there can be at most two substations installed in the area. The precedence constraints permit flexibility in combining alternatives. For example, it is straightforward to restrict policies such that, for alternatives A, B, and C, A cannot occur with B, A must be preceded by C, A cannot follow B and C, etc. This flexibility is a unique and powerful feature of the Area Investment Strategy Model.

- Salvage Value. The salvage value option places a value on the remaining life of installed assets that are present at the end of the planning period. There are two options. The first option permits the user to accept a lottery on the price or value of installed capacity at the end of the planning period. The user specifies the expected value of future capacity as a rental price ($/kW-yr.). A simple equation converts this expected value into the salvage value. This option seems reasonable if the planner anticipates that a competitive market for distribution capacity will develop. The second option called the "cost-to-go" estimates the actual policy cost over the remaining infinite horizon. To apply the "cost-to-go" option, the user must specify four additional parameters, (i) the average capacity cost in the future, (ii) the average operating cost in the future, and (iii, iv) the real escalation rates for each. Then, the infinite horizon cost computation is based on the additional assumption that an investment is replaced when it comes to the end of its useful life. Thus there is an infinite stream of replacements of each investment. A sensitivity analysis can be performed to observe how changes in the salvage value/cost-to-go affect the optimal policy.

- Advanced User. The user may also provide additional data that determines the way the strategy model actually performs its operations. First, in order to match the computer memory allocated to the problem with the size of the problem, the user may select the number of nodes allowed in the model. An advanced user will recognize that the investment strategy problem is formulated using a decision tree and that increasing the number of nodes in that decision tree increases the size of the decision problem that the model will solve. However, increasing the number of nodes also uses up more of the computer's memory and tends to slow down run times. The advanced user may prefer to perform the tradeoff of size and speed rather than accept the default values in the model.

  Second, the model can perform its work using the idea of coalescence. Coalescence means that all nodes in the decision tree for which the future description of the decision problem is identical (i.e., nodes at the same total load level for which the same collection of technologies has been installed) are treated identically with respect to future computations. When the coalescence option is enabled, the model only evaluates the earliest repeated node in the decision tree and uses the results of this evaluation to attach a value on all other similar, coalesced, nodes. Exercising the coalescence option can dramatically improve the run time, but the method is an approximation and can have some effect on the numerical results.

The advanced user should contact EPRI user support for a detailed explanation of these options and their potential effects.

3.2.3 Overview of the Operation of the Model

Formally, the Area Investment Strategy Model is a dynamic optimization model. The capacity expansion problem is formulated as a stochastic optimal control problem over an infinite horizon. A frequently applied representation of such a problem is based on the structure of a decision tree. The tree contains a sequence of nodes; each node is either a decision node or a chance node. At a decision node, a choice of paths can be made. At a chance node, the path that actually occurs is not chosen but rather is governed by a probability distribution. All the uncertain variables in the problem are modeled using probability distributions on chance nodes. A policy is a complete path through the tree, such that the choice made at each decision node is specified conditionally with respect to the sequence of resolutions of the prior chance and decision nodes, and the combination of choices and uncertainties determines the cost along each path. The optimal policy is the one that minimizes the expected cost through the tree.

An example of a decision tree is provided in Figure 3.4 below. Beginning at the left, the first decision is to choose between A and B. If A is selected there follows a chance node. With probability p, no further decisions need be made, which is symbolized by the triangle at the end of the branch emanating from the chance node. With probability 1-p, the next decision is to choose between C and D. Regardless of which one is selected, no further decisions are required. If B is selected initially, there follows a chance node such that with probability q, the next decision is to choose between E and F. Regardless of which one is selected, no further decisions are required. Alternatively, after B is selected, with probability 1-q, the next decision is to choose among E, F, and G. Notice that the alternatives can change depending upon the condition achieved.

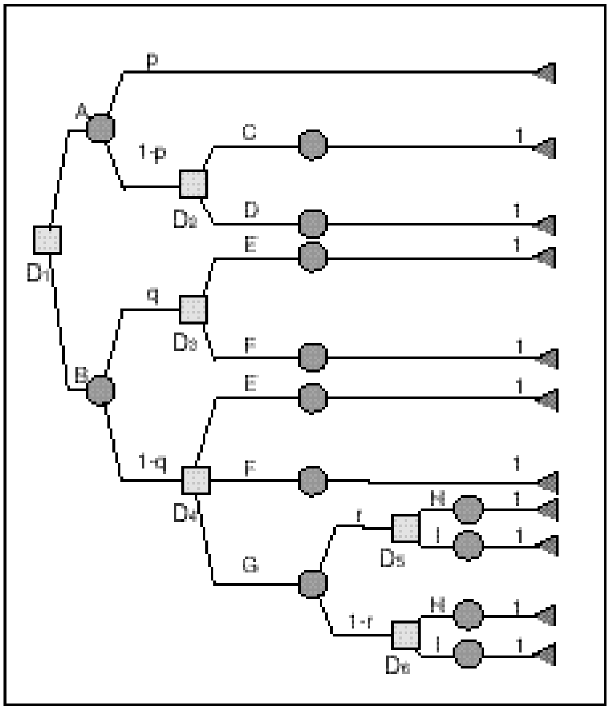

Figure 3.4 A Simple Decision Tree Structure

Choosing E or F means that no further decisions are required. Choosing G leads to another chance node, such that with probability r, the next decision is between H and I. Similarly, with probability 1-r, the next decision is between H and I. No further decisions are required.

A policy is a complete setting of the decision nodes for every possible resolution of the chance nodes. Thus, a policy is a set of contingent actions that specifies what to do in any state of the world. For example, one policy is: select A at the first decision node and if the lower branch of the chance node occurs, select C. Here is an alternate policy: select B, followed by E if the upper event occurs and G if the lower event occurs; after G, choose H in either case. The optimal policy is the path that is least cost. The algorithm that determines the optimal policy has only two rules: (1) replace all chance nodes by their expected cost, and (2) replace all decision nodes by the alternative with minimum cost over all possible alternatives.

The Area Investment Strategy Model begins by setting up the tree structure, using the constraints specified by the user to eliminate infeasible paths. The salvage value specifications provided by the user permit the model to compute the cost at the end of the tree. Once costs at the end of the tree have been determined, it is straightforward to compute the expected value of each alternative at each decision node. That expected value is based on the probabilities of uncertain events, such as load growth, as specified by the user. After computing all such expected values, the model searches the tree to find the path that is least cost. Results of the Area Investment Strategy Model analysis are presented in a format that is consistent with the tree structure used to solve the problem. Illustrative examples of the tree structure and the operation of the model are given in Section 3.2.3, below.

3.2.4 Model Outputs

The model currently produces four output reports.

1. Decision Frequencies. This report indicates which alternatives are installed in each year of the planning period under the optimal policy. The report is expressed as a matrix, with rows corresponding to the alternatives and columns listing the years. The number in each cell is the likelihood that a particular alternative will be installed in a particular year. If the value is 1.000, the alternative is always installed in that year. Values less than unity indicate that there is some chance that the alternative is not installed in that year. Empty cells mean that the alternative is never installed in that year. Note that neither the rows nor columns must sum to one, since it is not mandatory that every alternative be installed sometime, nor is it mandatory that some alternative must be installed in each year. The purpose of the report is to summarize the installations along the optimal policy, without regard for the contingent aspects of the installation sequence.

2. Risk Profile. This report describes the probability distribution of the cost of the optimal policy. The cumulative distribution is specified by a collection of points that describe the probability that the cost of the policy is less than or equal to the lower limit of each cost range. The purpose of the report is to provide a graphical description of the risk of the optimal policy.

3. Optimal Tree. This report presents the contingent optimal policy in the form of a tree. The sequence of decision nodes and chance nodes that describe the optimal policy is shown. At each decision node, the optimal alternative is noted. At each chance node, the probability of each event is listed, followed by the time that the next decision occurs. If an alternative has sufficient capacity to last until the end of the planning period, that is noted by the word "Terminate." The purpose of the report is to present the optimal policy in a simple graphical format that highlights the contingent nature of the policy. (See Section 3.2.3, below, for examples of the optimal tree output report.)

5. Input Data. This report echoes the input data provided by the user..

CHAPTER 4 ANALYSIS TUTORIAL

4.1 Introduction

This section discusses the modeling process and provides an example of how the Area Investment Strategy Model can be used to determine the best investment policy for a local area and to provide insights into the key issues affecting strategic policy. A sequence of tasks is described, designed to implement the Area Investment Strategy Model within an organization and to develop the capability to support local area decision making.

To illustrate the process of getting started with the Strategy Model, the investment-planning problem of a fictitious utility, National Power Company (NAPCO) will be analyzed. After describing how to set up the model, an initial base case will be developed from historical data. This initial case will be unrealistically simple; its purpose is to familiarize the new user with the structure of the problem and the handling of input data sets. The results of the first case will most likely be too simplified to use for policy recommendations. An iteration through the analysis cycle will make the model's representation of the electrical system more realistic and complex. A new input set will be developed to represent more accurately the utility's electrical problem and demonstrate many of the features of the Strategy Model. The tasks outlined in this section are designed to increase the user's knowledge of how the model works and how to determine what aspects of a specific problem are important.

To obtain the software, you must be a member of the Distribution Resources Target or the Distribution Business Area. A copy of the Strategy Model along with test data sets is available through:

EPSC/Electric Power Software Center
Suite 120
11025 North Torrey Pines Road
La Jolla, CA 92037
1-800-763-EPSC
(1-800-763-3772)

If you have any problems, you may call Steve Chapel, EPRI, (650) 855-2608.

4.2 Analysis Cycle

Before assembling input data and starting to run the Area Investment Strategy Model, it is important to examine what the result of the modeling process will be. The process by which it is used is as important as the model itself. The model is not a "black box" that will provide a single answer for the lower cost investment policy. While the Area Investment Strategy Model will identify the least-cost policy for a specified set of input parameters, support of investment planning decisions involves a great deal more than a single run of the model. Proper analysis is accomplished by many model runs, cycling through the process depicted in Figure 4.1.

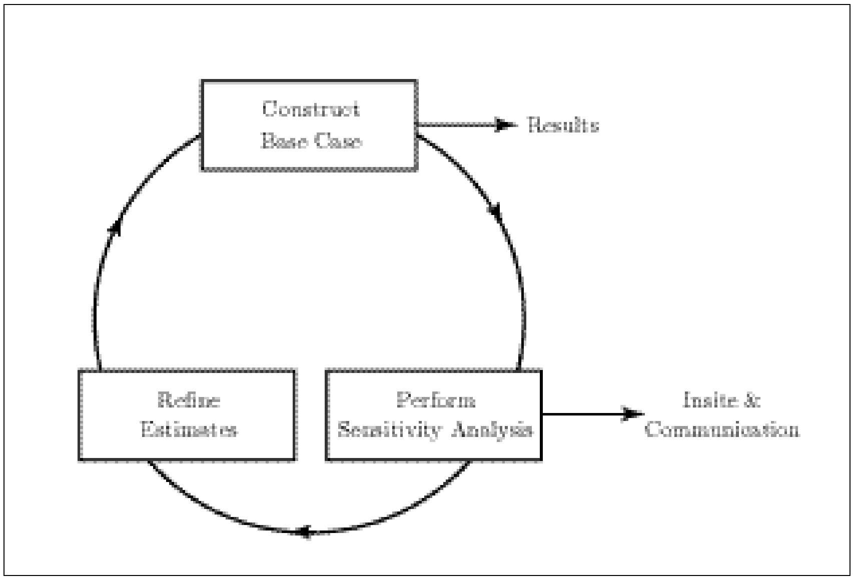

Figure 4.1 Analysis Cycle

The analysis cycle constantly refines one's description of the electrical system, testing the sensitivity of results to changes in inputs and identifying the key issues. Examination of the effects on model results when input data are varied provides a quantitative understanding of the critical factors that determine capacity requirements. The refined input description will not only determine the best investment policy, but also indicate how costs vary with policy and illustrate the consequences (e.g., capital and operating costs, investment risks and future load probabilities) of following a policy.

4.3 Using the Strategy Model

The Strategy Model requires a personal computer and a planner with experience in distribution planning. The planner installs the model on the host computer. The planner gathers the model data, operates the model, interprets the model results, and communicates results and insights to policy and decision-makers. Version 1.6 is relatively fast, taking seconds or minutes for a typical case on current PCs.

4.3.1 Work with Test Cases

The Strategy Model is shipped with three input data sets discussed in this section, along with the resulting output sets. The input files are **case1.aip, case2.aip**, and **case3.aip**. The output files are **case1.Report**, etc. These test data sets serve three purposes.

They:

1. demonstrate that the model has been correctly installed on the computer,
2. provide data sets for the tasks described in this section, and
3. provide templates for future data sets.

4.3.2 Structure the Problem

It is helpful to analyze investment-planning problems in general before studying the specifics of one’s own situation. This provides a context for the problem and identifies the issues appropriate for further investigation. It also highlights the items that will be candidates for sensitivity analysis following initial model runs.

NAPCO Example

As discussed earlier, there are five fundamental types of issues that must be evaluated in order to determine the least-cost expansion strategy for a local area. The first three basic issues are economic issues that exist under certainty. The last two issues exist because the future is uncertain.

- Value of Lower Life-Cycle Costs
- Value of Deferral
- Value of Tracking Load
- Value of Flexibility
- Value of Learning

Value of Lower Life-Cycle Costs:

Under certainty, the issue of which technology to use to satisfy load is a classical economics problem. The objective is to minimize the present value of the cost stream over time among the various alternatives. If two investments have the same size, the solution is determined by evaluating the tradeoff among capital costs, operating costs, and investment life. It is easy to demonstrate that the best alternative will be the one that has the lowest annualized, or levelized, cost.

Value of Deferral:

Deferral of an investment has a direct economic value, which is the opportunity cost of money. To see this, consider the example in Figure 4.2, which shows two alternatives for meeting load over some time period t. One alternative is to install the larger investment (capacity C, capital cost $X) at time 0, which would serve load until time t. The other alternative is to install the smaller investment (capacity C/2) at time 0 and then again at time t/2. Suppose the operating costs are zero and that the capital cost of the smaller investment is half the capital cost of the larger investment. In other words, there are no economies of scale. Also, for simplicity, suppose that t = 2 years. The present value capital cost of the large alternative is $X. The present value capital cost of the smaller alternative is $X/2 plus $X/2(1+r), where r is the discount rate. The discount rate is the cost of capital; for example, if the cost of capital is 10%, this means one must pay $1.10 in order to finance $1 of investment for a year. If $X = $1,000 and r = 0.1, the present value capital cost of the large investment is $1,000, whereas the present value capital cost of the small investment is $955. The small investment has an advantage of $45 simply because it defers the capital expenditures.

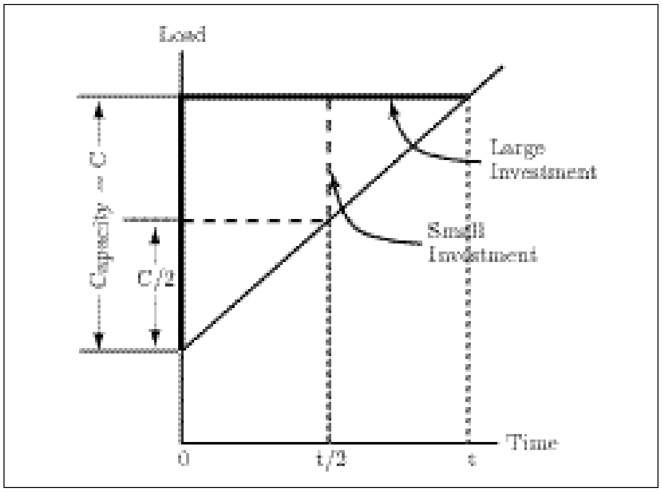

Figure 4.2 Benefits of Deferral

The cost of money, or the time value of money as it is sometimes called, is more than just an abstract economic concept. Note that if the cost of capital (real discount rate) were zero, there would be no capital cost savings due to deferral.

Value of Tracking Load:

Continuing the example of Figure 6, it is clear that the smaller investment tracks load more closely than the larger investment (i.e., the average amount of excess capacity is smaller with the modular investment). This can generate several additional types of economic value. For example, suppose that a disruption occurs somewhere else on the system. If a larger unit is installed the average amount of unserved energy will be less, since there is more excess capacity. On the other hand, if the investments themselves are unreliable, failure of the large-investment technology would have twice the negative capacity impact as failure of one of the smaller technologies. Perhaps the most direct benefit of tracking load comes with must-run generation units. Note in Figure 6 that any operating cost associated with the excess capacity of the larger unit is twice the operating cost of the excess capacity of the smaller unit (compare the areas above the load curve). Perhaps more subtle, but at least as important for investment planning, are two principles related to the value of modular investments under uncertainty.

Value of Flexibility:

Big investments are generally cheaper on a per kW basis, but they provide less flexibility than small investments. Small investments can defer big investments and provide a potentially valuable option to revisit the question of whether to make a big investment. For example, if load growth is uncertain and it is unclear whether a local area will saturate (i.e., hit a maximum load level) at a high or low level, it may be far less expensive to invest in small investment alternatives and then invest in the larger alternatives if and only if load grows sufficiently high. The test cases below will demonstrate (1) that the value of this type of flexibility can be very high, and (2) that it is impossible to estimate this value without addressing uncertainty explicitly.

Value of Learning:

The value of flexibility is intimately related to the value of learning about the investment uncertainties. If an investment uncertainty, such as load growth, is modeled as an independent process, then, by definition, there is no learning. For example, if load in period 2 is independent of load in period 1, then load being high in period 1 tells us nothing about the level of load growth in period 2. However, load growth is a good example of a process in which uncertainties are generally not independent. For example, it may be more likely that load growth will continue to be high after three years of high load growth than it would be for load growth to be high after three years of very low load growth. In this case, we can update our state of information about future load growth by observing past load growth and “tailor” our future decisions accordingly. For example, it may be appropriate after observing several years of high load growth to invest in a large alternative because the future load growth necessary to justify such an investment is much more likely

4.4 NAPCO Example

National Power Company wants to analyze the area investment policy for Thapa County. The Thapa County system currently consists of one small and aging substation. The basic decision Thapa County's planners face is whether to replace the old substation with a new one or whether to install a smaller-capacity feeder and then possibly use local generators over time if needed to serve load.

4.4.1 Setting up an Initial Base Case

The first base case constructed will take an intentionally simplified broad-brush approach to the problem. It will ignore uncertainty in load and differences in investment sizes. It will provide a context for setting up the basic model inputs and will provide an initial planning policy under certainty. While simplistic, this type of case allows the new user to get the model working and to begin to understand how the model works. For the experienced user, starting with a simplified case helps to identify areas in which further refinement of the data is necessary, and provides a basis for sensitivity analysis.

For ease of understanding, we shall build up a progression of cases, starting with the simplest possible case and adding complexity as we go. This is the recommended process in a real application. Table 4.1 summarizes the key planning assumptions common to all cases. Table 4.2 shows the model inputs for the simplest possible base case.

Note these tables list the key input parameters but not every input parameter for the tutorial cases. The user can examine all of the data by using the model interface. To do this open the *File* menu, select *Open Case* and select the appropriate input file from the dialog box. The inputs can then be examined using the *Edit* menu.

Note: Do NOT attempt to build the cases from scratch using the data in Table 4.1 and Table 4.2. These are not complete data sets. Use the input files distributed with the software.

COMMON ASSUMPTIONS FOR ALL CASES		
Time Horizon	20 years	
Discount Rate	5%	
Inflation Rate	0%	
Accounting Method	Before Tax Cash Flow	
Initial Load	100,000 kW	
Terminal Value Specifications		
- Price of Capacity at Terminal Time	$10/kW-yr.	
Variable O&M Costs	$0	
Emissions Costs	0	
Load Shape	*Time (hrs)*	*% of Peak*
	0	100%
	8760	0%
Load Growth Trends		
	Growth Rate	
Low	1.005 (0.5%)	
Medium	1.01 (1.0%)	
High	1.05 (5.0%)	

Table 4.1 Common Assumptions

Note: The load growth trends used by the model may be slightly different than the user input trends. The model algorithm requires an approximation of the user inputs. The approximation is very accurate and is reported in the model output report "*Input Data*."

BASE CASE			
Technologies (Capacity Alternatives)	*Size (kW)*	*Cost ($1000)*	
Technology S	25,000	$2,500	
Trend Transition Probabilities (Load Growth Specifications)			
	Low (0.5%)	*Medium (1%)*	*High (5%)*
Low (0.5%)	*1.00*	0.00	0.00
Medium (1%)	*1.00*	0.00	0.00
High (5%)	*1.00*	0.00	0.00
Initial Load Growth Rate		"Low" 0.5%	
Result (NPV Cost $1000)		$ 1,326	

Table 4.2 Key Assumptions For The Base Case

Case Description

The input file for this case is **case1.aip**. The base case includes a single technology, call it technology S (think of it as a new substation), with capacity 25,000 kW having a capital cost of $2,500,000 (entered as $2,500 since we will assume cost units in thousands of dollars). We assume load grows at a fixed and certain rate of 0.5% per year starting from a base of 100,000 kW. Note that the user can run deterministic cases by assigning a probability of 1.0 to one load-growth trend (in this case, the initial load growth rate is 0.005 and there is probability 1.0 of staying at the initial rate).

For simplicity in this case and in the following cases, we assume the substation has no operating costs of any type, and no losses and unserved energy. We wish to plan over a time horizon of 20 years using a 5% real discount rate (no inflation). Finally, we will use the *Terminal-Lottery-on-Capacity-Price* method for salvage value (see CHAPTER 2 for an explanation of this method) and assume that the price of capacity at the end of the time horizon is $10/kW-yr.

Case Results:

The optimal policy based on this case is shown in Table 4.3. This is a simple example of the Optimal Tree output report. (Section 3.2.4, Model Outputs, describes how to read the Decision Frequencies, Optimal Tree and Risk Profile reports.) Here, for brevity, we will focus on the optimal policy only. For this simple case, the results are not surprising. There is only one technology, the substation S, and it is large enough to serve the certain projected load (at a certain load growth of 0.5%, the load grows by 10,490 kW over 20 years). Since S is the only alternative, the optimal policy is simply to install S at the beginning of the time horizon. The value of the optimal policy, which is the discounted present value of the cash flow of costs, is $1,326, which is equal to the capital cost of S plus the discounted salvage value.

Decision (Stage 1)	Chance	Decision (Stage 2)
Optimal Value = 1325.78		
S at t=0.00, L=100000	p=1.000, t=44.74, g=1.0050	Terminate at t=20, L=110490

Table 4.3 Optimal Strategy for the Single-Technology Case

Insights and Observations

Although the base case is very simple, nevertheless we can begin to build insight by performing some simple sensitivity analyses, changing assumptions and seeing how the case changes. The reader is encouraged to run some cases with different inputs. For example, try running the same case for non-zero settings of the salvage value to see how the total cost changes. Also, try both the alternative salvage value definitions.

Another useful exercise would be to add realism by introducing the different components of operating costs to see how they affect the optimal value. Finally, add some losses and unserved energy costs to get a feeling for their impacts on the bottom line.

4.4.2 Introducing a Modular Technology

Case Description

We next investigate two subcases that introduce an additional modular technology, a feeder F that has 1,250 kW capacity, or half that of S. The user should modify **case1.aip** to create the inputs for this case. The first subcase is summarized at the top of Table 4.4 (Introducing a Modular Technology [no economy of scale]). This subcase is identical to the base case except that there is an additional capacity alternative, Technology F, that is exactly half the size of Technology S and has exactly half the capital cost ($1,250). Thus, both alternatives have the same unit cost of $100/kW. The second subcase, summarized at the bottom of Table 4.4 (Introducing a Modular Technology [with economy of scale]), introduces a perhaps more realistic Technology F, which is half the size of Technology S but with twice as high a unit cost of $200/kW.

INTRODUCING A MODULAR TECHNOLOGY (NO ECONOMY OF SCALE)				
Technologies (Capacity Alternatives)	*Life*	*Description*	*Size (kW)*	*Cost ($1000)*
Technology S	40	Large	25,000	2,500
Technology F	40	Modular	12,500	1,250
Trend Transition Probabilities (Load Growth Specifications)				
		Low (0.5%)	*Medium (1%)*	*High (5%)*
Low (0.5%)		1.00	0.00	0.00
Medium (1%)		1.00	0.00	0.00
High (5%)		1.00	0.00	0.00
Initial Load Growth Rate			"Low" 0.5%	
Result (NPV Cost $1000)			$663	

INTRODUCING A MODULAR TECHNOLOGY (WITH ECONOMY OF SCALE)				
Technologies (Capacity Alternatives)	*Life*	*Description*	*Size (kW)*	*Cost ($1000)*
Technology S	40	Large	25,000	2,500
Technology F	40	Modular	12,500	2,500
Trend Transition Probabilities (Load Growth Specifications)				
		Low (0.5%)	*Medium (1%)*	*High (5%)*
Low (0.5%)		1.00	0.00	0.00
Medium (1%)		1.00	0.00	0.00
High (5%)		1.00	0.00	0.00
Initial Load Growth Rate			"Low" 0.5%	
Result (NPV Cost $1000)			$1,326	

Table 4.4 Introducing a Modular Technology

Case Results

The results are quite interesting. Table 4.5 shows that the optimal policy for the No-Economy-of-Scale case is to use the smaller Technology F. Note that the NPV cost of $663 is lower than the $1,326 cost associated with Technology S, even though both S and F have the same cost per kW. This cost reduction occurs because of the smaller technology's ability to follow load more closely and defer a portion of the investment (see the Value of Tracking Load section above).

Decision (Stage 1)	Chance	Decision (Stage 2)
Optimal Value = 662.89		
F at t=0.00, L=100000	p=1.000, t=23.62, g=1.005	Terminate at 20, L=110490

Table 4.5 Optimal Strategy for the No-Economy-of-Scale Case

Now consider the optimal policy for the Economy-of-Scale case shown in Table 4.6. We see that the least-cost policy is the same as in the base case: install the large Technology S immediately to serve load over the time horizon. Thus, we see that the economy-of-scale benefit of Technology S is worth more than the load-matching benefit of Technology F.

It is interesting to note, with the low load growth of 0.5%, we had to increase the cost of technology F to $180/KW before the economy-of-scale benefit outweighed the load-matching benefit.

Decision (Stage 1)	Chance	Decision (Stage 2)
Optimal Value = 1325.78		
S at t=0.00, L=100000	p=1.000, t=44.74, g=1.0050	Terminate at 20, L=110490

Table 4.6 Optimal Policy for the Economy-of-Scale Case

Insights and Observations:

The model explicitly trades-off economies of size as captured by load matching ability versus economies of scale as captured by capital costs. Try running some cases with different capital costs and different technology sizes. You will find that there is no general rule-of-thumb for anticipating which effect will dominate strategy in any particular situation. Indeed, a key benefit of this type of model is its ability to make economic tradeoffs of these and much more complex situations.

But we are not done with our evaluation of the modular technology. In fact, we have left out a key element of the value of a small modular technology; that is, its ability to provide management flexibility. We shall investigate this feature in the next case.

4.4.3 Uncertainty and Management Flexibility

Case Description

We now examine the implications of adding load growth uncertainty into the case. The input file for this case is **case2.aip**. We shall see that uncertainty has profound implications for choosing the best strategy as well as for accurately valuing the benefits of a modular technology. Table 4.7 shows inputs for a new case, titled Uncertainty and Management Flexibility. The new case is identical to the previous case except: 1) the load transition matrix now has probabilities other than one or zero, and 2) we assume that technology F costs $1.60 million ($128/kW). This means that load is now modeled probabilistically – the model reflects explicitly the fact that load is uncertain. The "steady state probability" is the long run average fraction of time that load will be in any of the three growth states. The steady state probability and the reason for including results with S and F only are explained below.

Uncertainty and Management Flexibility			
Technologies (Capacity Alternatives)		*Size (kW)*	*Cost ($1000)*
Technology S		25,000	2,500
Technology F		12,500	1,600
Trend Transition Probabilities (Load Growth Specifications)			
	Low (0.5%)	*Medium (1%)*	*High (5%)*
Low (0.5%)	0.80	0.20	0.00
Medium (1%)	0.80	0.10	0.10
High (5%)	0.10	0.10	0.80
Steady State Prob.	0.74	0.17	0.09
Initial Load Growth Rate		"Low" 0.5%	
Result (NPV Cost $1000)		$979	
-- with S only		$1,317	
-- with F only		$1,089	

Table 4.7 Management Flexibility Case

The table of load growth probabilities is called the transition matrix. The transition matrix contains the probability of various load growth rate levels in the next period, given the current load growth trend. For example, if the load is currently growing at

0.5% (*Low*), then the probability that next period's trend will be *Low* is 0.8; the probability that next period's trend will be *Medium* (1.0%) is 0.2; and there is no chance of jumping from the lowest trend to the highest trend in a single period. However, if the current trend is *Medium*, there is a 0.8 chance of jumping to Low, and a 0.1 chance of either jumping to *High* or remaining in *Medium*. Finally, if the current trend is *High*, there is a 0.1 chance of jumping to *Low* or *Medium* and a 0.8 chance of staying at *High*.

The user specifies the levels associated with each trend, the trend name and the probabilities defining the uncertain evolution of the load as model inputs.[1]

Case Results

Table 4.8 shows the optimal tree through Stage 5.. This is a somewhat more complex example of the output report provided by the model. The tree indicates now that the optimal strategy is what we call a contingent policy. Start with the smaller Technology F. If load growth is fast enough so that new capacity is needed before the end of the 20 year period, then, depending on the rate of growth of load, install either Technology F or S. If load growth has been very low (probability 0.619 of average growth of 0.57%), F will cover any load growth to the end of the planning horizon. But, in the fast-growth contingency, if load continues to grow at a high rate after F is installed, additional new capacity may be needed in about two and one-half more years, or when eight years have elapsed since the beginning of the planning period. If so, S should be installed. The substation will satisfy load growth in all but the high growth case. If subsequent load growth is high, then additional installations of S will be required.

---

[1] **Note**: A methodology and computer program is available to assist users in developing their own load model parameters. The program, LOAD ASSESSOR, is tool that is part of the Strategy Model.

Decision (Stage 1)	**Chance**	**Decision (Stage 2)**
Optimal Value = 978.90		
F at t=0.00, L=100000	p=0.619, t=20.78, g=1.0057	Terminate: L=112005
	p=0.271, t=14.68, g=1.0081	S at t=14.68, L=112500
	p=0.110, t=5.40, g=1.0220	F at t=5.40, L=112500

Chance	**Decision (Stage 3)**	**Chance**
p=0.546, t=34.64, g=1.0058	Terminate, L=116019	
p=0.306, t=22.01, g=1.0092	Terminate, L=118087	
p=0.148, t=7.18, g=1.0283	Terminate, L=130513	
p=0.470, t=18.59, g=1.0057	Terminate, L=122202	
p=0.229, t=11.81, g=1.0090	S at t=17.22, L=125000	P=0.550; t=31.56; g-1.0058
		P=0.296; t=20.03; g=1.0091
		P=0.154; t=6.25; g=1.0296
p=0.301, t=2.53 g=1.0425	S at t=7.93, L=125000	P=0.243; t=30.93; g=1.0059
		P=0.295; t=17.82; g=1.0103
		p=0.462; t=4.34; g=1.0429

Decision (Stage 4)	**Chance**	**Decision (Stage 5)**
Terminate; L=127025		
Terminate; L=128205		
Terminate; L=135572		
Terminate; L=134214		
Terminate; L=141427		
S at t=12.27; L=150000	p=0.236, t=26.41, g=1.0059	Terminate, L=156921
	p=0.263, t=15.30, g=1.0101	Terminate, L=162134
	p=0.501, t=33.53, g=1.0446	S at t=15.80, L=175000

Table 4.8 Optimal Strategy with Management Flexibility

Thus, in this case, under the optimal (least-cost) strategy there three sequences of installations that optimally serve load for 20 years: F or F-S or F-F and one or more S. Moreover, the actual installation sequence will depend on the resolution of the load uncertainty over time.

Insights and Observations

Part of the value of a modular technology is the ability it affords to "wait and see" before making a big commitment. The optimal strategy under uncertainty takes advantage of management's ability to react. We call this management flexibility. One of the features of the model is we can use it to attach an economic value to this management flexibility.

Determining the Economic Value of Management Flexibility: Note that the overall cost of the optimal strategy is $979. This has little meaning until we run some comparative cases. We ran two additional cases with the same load forecast, one assuming only Technology S was available, and one assuming only Technology F was available. (These cases are easy to run using the Installation Constraint feature, which is an option under the model's main Edit menu). The results are shown at the bottom of Table 4.7.

For this case there are two inflexible strategies S only and F only. F is a more expensive technology ($128 per kW versus $100 per kW) but follows load more closely and is better able to match the uncertain load growth. The benefit of F over S is significant: $1,317 - $1089 or $288. However having the flexibility to use both technologies provides an even better result. The optimal strategy that uses both technologies has 10 percent lower cost that F alone and over 25 percent low cost than S alone.

Learning

A key feature of the model is that it explicitly takes into account the ability of the utility to learn about load growth tendencies as load grows over time. A flexible strategy can take advantage of the ability to learn about future load tendencies from observing current load trends. Learning can greatly influence the optimal strategy and its costs. This case demonstrates that combinations of alternatives working together in a dynamically optimized strategy creates real economic value. The strategy and economic value can only be discovered by explicitly modeling the uncertain environment in which the alternatives operate.

As a side note, the model's ability to analyze dynamically changing strategies can help utility planners by giving guidance as to what to monitor as time goes by.

4.4.4 Pitfalls of Scenario Planning

As mentioned in the introduction, a common heuristic used to develop strategy under uncertainty is the technique of scenario planning. Scenario planning is simply the study of alternative deterministic plans. Unfortunately, scenario planning does not get either the policy or the costs right. It is supported by deterministic models and thus cannot reflect the dynamic nature of the decision process and help the decision maker decide how to react as the world evolves. We have developed an example to illustrate how unrealistic and misleading deterministic and scenario planning can be.

Scenario planning has several serious pitfalls. A typical procedure would be to run a model under several deterministic scenarios, determine the best strategy under each scenario, and then concoct a mixed strategy from the results. This often produces suboptimal (not least cost) results. Try running the model with a deterministic low-load trajectory over the planning period and with a deterministic high-load trajectory over the planning period. You will find that the best “low strategy” is to install a single feeder with an optimal value of $2,871, The best “medium strategy” is to install a single substation with an optimal value of $3,579 and the best “high strategy” is to install a series of substations with an optimal value of $10,389. There is no way to deduce from these three scenario runs even an approximation to the actual least cost strategy, which is to “mix and match” S and F over time depending on load conditions. This is a general observation. In fact, given a single load trajectory, deterministic economic analysis will always select the same technology, based on size and cost, and repeat it as needed over the planning period. This is because the critical tradeoff in scenario analysis is between scale economy and cost of excess capacity. Learning and flexibility are ignored. Hence, no mixed strategy emerges from such an analysis.

A scenario planner might also suggest that we approximate the cost of the optimal strategy by weighting the three scenario costs with the relative probabilities of each scenario. How to figure the three scenario weights when in fact there are hundreds of potential future load paths is problematical. But, let’s give the planner the benefit of the doubt and weight the scenarios with the probabilities of all the future paths with loads similar to those in the three scenarios (this gives low, medium, high scenario probabilities of 0.5, 0.3, 0.2). The result is an expected value of $4,587. This is quite different than the actual expected cost of the optimal strategy, $4,009. In fact, there is no reason to expect the scenario planning approach to provide an accurate estimate in general. Most important, the scenario analysis does not provide information or insight about the actual least cost policy.

Two additional points should be made. First, the scenario approach makes it clear that, since the answer depends on the scenario analyzed, the optimal answer MUST depend both on the relative likelihoods of the different scenarios and on the potential for shifting from one scenario to another during the planning period. These shifts are captured in the description of the probabilistic behavior of load growth rates. Thus, there is no way to avoid a probabilistic analysis if one wants to determine the truly least-cost policy. Second, it is important to state that, while deterministic scenario analysis is not useful for developing the optimal strategy, it can be useful in providing insight into how the model works and into the effects of different assumptions

CHAPTER 5 REAL UTILITY EXAMPLE

5.1 Introduction

The data and results in this subsection are based on an actual utility case study. National Power Company wants to develop and investment plan for the Teco substation area. The Teco area currently consists of one small and aging substation. Currently, load is growing about 1% per year but development is starting to take place. Peak load is close to existing capacity. Significant growth is likely to occur in the area some time in the next five to ten years. The extent and timing of that growth is very uncertain. .

The alternatives for meeting future capacity needs are to add a feeder from an adjacent substation or replace the existing substation with a new and larger substation. If the substation is added first, it will meet future capacity needs and the feeder will not be required. If the feeder is added first, the substation can be added later as load dictates the need for additional capacity.

As an alternative to the substation and feeder, small generators could be used to meet small increments in load and to defer the larger investments. Based on engineering and siting work, the utility determined that as many as four distributed generators could be placed in the area. However because of limits on where generators can be sited, each generator's effective unit capacity is lower than its nameplate capacity, and the cost per kilowatt increases as the generators are added.

The EPRI strategy model team worked with the utility to develop the model inputs and especially to create a characterization of the required load parameters (as mentioned in an earlier section, a tool now exists, *The Load Assessor*, that can be used directly by a utility to estimate the load uncertainty parameters). The load behavior for the area tends to follow strong trends. That is, when load is growing slowly it tends to continue to grow slowly and when growing at a rapid rate, it tends to continue to grow at that rate for a number of years.

Introduction

Table 5.1 summarizes the key planning assumptions for the planning study. Three illustrative cases are reported here. Case 1 assumes that the local generators, once installed, are not removed (salvaged) when the larger traditional investments are made. Cases 2 and 3 assume that the generators can be salvaged if removal makes economic sense. Case 3 also assumes that there is “no-learning” associated with the load uncertainty behavior (for an explanation see the “Learning” case reported in an earlier section). Table 5.2 shows the model inputs for cases 1 and 2.

To replicate the results in this chapter the user should use the input file **real_case.aip** as the starting point. This file contains the inputs for Case 1 – “No Salvage”. Following the instructions in this chapter, the user can modify this input file to produce the inputs for Cases 2 and 3.

The input and output screens shown in CHAPTER 2 are for the input file **real_case.aip** modified to be “Case 2 – “Salvage.” The user may want to refer back to CHAPTER 2 to verify the inputs.

Note: Do NOT attempt to build the “Real Utility” cases from scratch using the data in Table 5.1 and Table 5.2. These tables contain key input assumptions but are not complete data sets. Use the input file distributed with the software, *real_case.aip*.

.

COMMON ASSUMPTIONS FOR STUDY		
Time Horizon	12 years	
Discount Rate	5.77%	
Inflation Rate	4%	
Accounting Method	Before Tax Cash Flow	
Initial Load	44,608 kW	
Maximum Area Load	70,000 kW	
Saturation On-Set Load	60,000 kW	
Terminal Value Specifications		
- Price of Capacity at Terminal Time	$10/kW-yr	
Variable O&M Cost – S & F	$0.02/kWh	
Variable O&M Cost – Engines	$0.05/kWh	
System Avoided Costs	$0.02/kWh	
Emissions Costs	0	
Load Shape	*Time (hrs)*	*% of Peak*
	0	100%
	88	95%
	264	90%
	8759	25%
	8760	0%
Load Growth Trends		
	Growth Rate	
Low	1.01 (1%)	
Medium	1.0201 (2%)	
High	1.051 (5%)	

Table 5.1 Common Assumptions for Study

ASSUMPTIONS – CASES 1 & 2			
Technologies (Capacity Alternatives)	*Life*	*Size (kW)*	*Cost ($1000)*
S: Substation	40	20,000	$2,000
F: Feeder	30	6,000	$900
E1: Engine 1	30	3,000	$1,500
E2: Engine 2	30	1,500	$750
E3: Engine 3	30	3,000	$2,250
E4: Engine 4	30	3,000	$2,500
Trend Transition Probabilities (Load Growth Specifications)			
	Low (1%)	*Medium (2%)*	*High (5%)*
Low (1%)	0.75	0.25	0.00
Medium (2%)	0.125	0.75	0.125
High (5%)	0.00	0.25	0.75
Initial Load Growth Rate		“Low” 1%	

Table 5.2 Key Assumptions for Cases 1 & 2

5.2 Case 1 “No Salvage”

The least cost policy is shown in Table 5.3. When the engines are constrained to remain in place, the best policy is to install the feeder followed by the substation. This meets all potential future load that is anticipated for the area. This policy is not surprising given that the cost of the substation is \$100 per kW while the engines cost a minimum of \$500 per kW.

Decision (Stage 1)	Chance	Decision (Stage 2)
PV Cost = 2334.52		
F at t=0.00, L=44608	p=0.182, t=11.44, g=1.0111	S at t=11.44, L=50608
	p=0.595; t=7.82; g=1.0163	S at t=7.82; L=50608
	p=0.223; t=4.39; g=1.0292	S at t=4.39; L=50608

Table 5.3 Least-Cost Strategy for “No Salvage” Case

5.3 Case 2 “Salvage”

The least cost policy is shown in Table 5.4. E1 and E2 are engine investments, and T indicates that the end of the planning period has been reached. In this case engines can be removed when large capacity is added. The least-cost policy is very different. Now, even though engines cost far more per kilowatt capacity than either the substation or feeder, they are part of the strategy. The fact that they can be removed when the larger investment is made, reduces their contribution to present value costs. It now makes sense to use the small modular investments to delay the traditional investments until load dictates that the larger investments are needed. It is also interesting that only two engines enter the optimal policy. This is due to two factors: (1) after the first two engines, the cost per kilowatt increases substantially and (2) compared with the feeder and substation alternatives, the capital cost per kilowatt of engines is relatively high. The engines are best used to delay large investments. Engines are not an economically efficient choice for providing large amounts of capacity.

Table 5.4 shows only part of the optimal policy through “Decision Stage 4.” The reader is encouraged to obtain the model, experiment with this case, and to explore the detailed results.

Case 2 “Salvage”

Decision (Stage 1)	**Decision (Stage 2)**	**Decision (Stage 3)**	**Decision (Stage 4)**
PV Cost 1626.22			
E1	E2	S(-E1, -E2)[2]	T
		F(-E1, -E2)	S
		F(-E1, -E2)	S
	F(-E1)	S	T
		S	T
		E1	S(-E1)
	F(-E2)	S	T
		E1	S(-E1)
			E2
			S(-E1)
		S	T

Table 5.4 Least-Cost Strategy for “Salvage” Case (through Decision Stage 4)

---

[2] (-E1, -E2) means that the two engines are removed (salvaged) and replaced by the feeder and substation investments in stage 3. This allows the same engines to be used in subsequent stages.

Note: this table provides the results through Decision Stage 4.

5.4 Case 3 "No Learning"

The no learning case assumes that there are no trends in load growth – that the load growth in the next period is independent of the current growth rate. This means that you learn nothing about longer-term growth rates by waiting to observe the load rate in the next period. This behavior is modeled by changing the transition probabilities so that they are the same for all three trends. Table 5.5 shows the transition probabilities assumed for this case.

Note that these transition probabilities represent the long-term average for the chance of being in one of the three growth trends. This is known as the steady-state probability distribution. These steady-state probabilities were derived from the transition probabilities assumed for Cases 1 and 2.

ASSUMPTIONS – CASES 3

Trend Transition Probabilities
(Load Growth Specifications)

	Low (1%)	*Medium (2%)*	*High (5%)*
Low (1%)	0.25	0. 50	0. 25
Medium (2%)	0. 25	0. 50	0. 25
High (5%)	0. 25	0. 50	0. 25
Initial Load Growth Rate		"Low" 1%	

Table 5.5 Assumptions for Case 3

The least cost policy given these load uncertainties is similar to that for Case 2. However the cost, $1941.62, is much higher – almost 20 percent. If trends in growth exist, you can develop polices that take advantage of the trend information. The resulting policies can, in some cases, have significantly lower costs.

Case 3 “No Learning”

CHAPTER 6 TECHNICAL FOUNDATIONS

The purpose of this chapter of the Guide is to discuss the technical foundations of the Strategy Model. Each section describes a particular aspect of the investment strategy problem. The reader may find it useful to consult the introductory overview in each section. This overview describes the aspect of the problem addressed in the section and the basic modeling approach. The overview is not meant to be technically descriptive. Additional technical details are appended so that the interested reader may discover how each issue is actually modeled. These appended details are often expressed in mathematical language. The user need not master the technical details in order to use the model successfully. This volume is provided mainly for reference.

6.1 Problem Formulation and Modeling Issues

6.1.1 Overview

This section discusses the methodology underlying the Strategy Model. The discussion is somewhat abstract, since the purpose of the section is to present the fundamental modeling ideas employed in the formulation of the distribution system investment-planning problem. In this section, we have made no attempt to relate these ideas to the actual model implementation. Users interested in reading more about the theory applied to the distribution system investment planning problem may be interested in this section. Users with little interest in theory may safely skip this section without any loss of ability to use the methodology.

We begin by noting that there is at present no generally accepted methodology for distribution system investment planning. Recently, distribution investment planning methods have focused on integrating distributed resources into conventional transmission and distribution (T&D) expansion plans. With relatively minor variations, four common assumptions characterize the methodologies used to investigate distribution system investment planning and measure the value of distribution system investments, including investments in distributed resources:

- Each methodology requires the prior specification of a conventional deterministic expansion plan for the T&D system. Each investment in the expansion plan is assumed to be made at a specific time in the future.
- The future peak load on the system is assumed known with certainty over the entire planning period.
- The capital costs and operating costs associated with future investments are assumed known with certainty over the entire planning period.
- The benefit of investing in distributed resources is assumed to be achieved by deferring the capital investments in the conventional expansion plan.

Practice has revealed that such assumptions are insufficient both for distribution system investment planning in general and distributed resources integration in particular. The shortcomings include the following. First, distribution planners rarely have a prior expansion plan that can be specified for an arbitrarily long period of time (typically ten, fifteen, or twenty years in the future). Second, the peak load on the system is quite uncertain; it cannot be predicted with certainty even for a short period in the future. We have found utility planners to be uncomfortable with the accuracy of long term deterministic forecasts. Third, the future costs of technology are uncertain. Fourth, contrary to a commonly-held belief, deferring a given conventional expansion plan for the longest time possible *does not* yield the least-cost solution to distribution investment planning with distributed resources.

The Strategy Model methodology overcomes these difficulties. It is designed to be a distribution system investment planning tool. The methodology does not require prior specification of an expansion plan, nor is the objective to maximize the deferral of conventional investments. Rather, the methodology determines an optimal mixture of conventional investments and distributed resources that minimizes the expected net present value of total system costs of serving load. The methodology recognizes that the future is uncertain and allows the user to characterize the important uncertainties that govern the investment problem. Moreover, the methodology explicitly recognizes and attaches an economic value to the management flexibility inherent in distributed resources options.

We also emphasize that the purpose of the Strategy Model is to guide the decision maker to what appears to be the best *current* investment in distribution assets. Because assets typically have long lives, the best current decision cannot be determined without considering the future consequences of that decision. Therefore, the Strategy Model plans for the indefinite future in order to determine what is best to do in the short term. As information and other conditions change, the long-term strategy found by the Strategy Model may be expected to change. But at any point in time, the current investment policy provided by the Strategy Model is optimal with respect to current conditions and current assessments of the effects of future uncertainties.

The remainder of this section describes the modeling approach in greater detail.

Users not interested in technical details may skip the remainder of the section without loss of ability to apply the Strategy Model.

6.1.2 Problem Formulation

The objective of the distribution system investment planning problem is to meet capacity and energy needs over the indefinite future at the lowest expected present value of all future costs. There are several aspects of this problem statement that require further discussion. It is convenient to develop the ideas in stages.

Deterministic Distribution Investment Planning Problem

We begin with a deterministic formulation assuming continuous time and infinitely divisible assets, and then extend the formulation to consider lumpy investments under uncertainty. The minimum cost objective may seem innocuous, but in fact represents an essential difference between our methodology and other methodologies. Other approaches find a solution that delays the construction of new T&D facilities by investing in distributed resources until the total cost of the distributed resources investments exceeds the total benefit of T&D delay. This solution is not least cost. The differences in the approaches are shown in Figure 6.1. The total cost is minimized when the difference between the cost savings due to deferral of T&D and the cost of this deferral, the cost of distributed resources, is maximized. This occurs when the *marginal* cost of deferring equals the *marginal* benefit. Since the maximum T&D deferral time occurs at the point at which the *total* deferral cost equals the *total* benefit of deferral, the minimum cost solution is earlier and requires less investment in distributed resources. Stated simply, maximal deferral is not the least-cost strategy and, in fact, undervalues the benefits of distributed resources while increasing investment costs with insufficient compensating benefits.

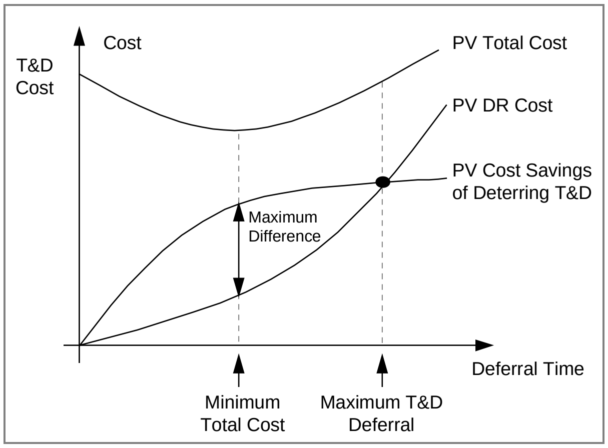

Figure 6.1 Relationship between Least Cost and Maximum Deferral Solutions

The solution to the deterministic distribution system investment planning problem is a sequence of capital investments in technology that provides the capacity necessary to meet the area load over time. It is natural to consider that the time is composed of periods of possibly variable duration. The key idea embedded in the mathematical formulation is that time is not indexed by a fixed interval, say year by year, which would define a problem of enormous dimension (e.g., 10 possible technology choices over a 10 year planning period implies 10^{10}, or 10 billion separate strategies that would have to be evaluated).

Further, we assume, as is the case under current regulatory and economic conditions, that there will be no cost savings or benefits associated with early installation. Thus, since there is no value to over-capacity, it follows that capacity decisions need only be made when capacity is needed (e.g., if the technologies each supplied two years of load growth, there would be only 10^5, or 100,000 separate strategies that would have to be evaluated). The resulting deterministic planning problem is in the form of a discrete optimal control problem.

A key unaddressed issue in the deterministic formulation is the behavior of the load from period to period. We address this issue by incorporating a stochastic-process model of load uncertainty into the formulation, and explicitly modeling the ability of the distribution system planner to adopt strategies that respond flexibly as the uncertainty resolves over time.

Probabilistic Distribution Investment Planning Problem: Value of Management Flexibility

Management flexibility in the face of uncertainty about the future is possibly the greatest benefit provided by modular distributed resource technologies to utility planners. Before proceeding, it is useful to consider a simple example. Suppose a planner is faced with a certain load growth of 10 units this year. However, the load next year is uncertain: there is a 60% chance of a 90 unit increase due to a potential major industrial electricity user moving into the area, and a 40% chance of no increase. In either case, no subsequent load growth is expected. The planner has two alternatives: a large conventional technology with 100 units of capacity at a capital cost of \$100M or a modular technology with 10 units of capacity at a capital cost of \$20M (twice the large technology's cost per unit of capacity). The traditional deterministic approach would assume that the load in the second year will attain the most likely value, 90 units. With a discount factor of 0.9 (i.e., \$1.00 in a year is valued at \$0.90 today), the least cost alternative is to install the large technology immediately to satisfy the load over the two year planning horizon (present value [PV] of \$100M). This cost is superior to that of the alternative strategy, which is to install the modular technology followed by the large technology (PV = \$110M = \$20M + 0.9x\$100M). Compare this with a "Wait and See" strategy of installing the modular technology, and then installing the large technology only if load grows. The expected present value of the Wait-and-See strategy is \$74 M = \$20M + 0.6 x 0.9 x \$100M + 0.4 x \$0M. Thus, the modular alternative in conjunction with the flexible Wait-and-See strategy saves the utility planner \$26 M in expected present value. The value of the modular alternative comes not from its cost characteristics—it is more expensive per unit than the large technology—but from the management flexibility it affords the utility planner. Leaving uncertainty out of a planning analysis systematically undervalues modular distributed resource alternatives.

Therefore, we assume that the underlying load dynamics is governed by a particular stochastic process. Two fundamental factors motivated our approach for characterizing load uncertainty. First, for distribution planning, a key issue is at what point in the future will load growth result in new capacity requirements? Thus, a complete description of potential load trajectories over time is required in order to specify the probability distribution on time to the next investment decision. That is, since there is no benefit to overcapacity, the time to the next investment decision is the time required for load to grow sufficiently so that the current installed capacity can no longer meet the demand. Second, based on extensive work with utility planners, we found that it is desirable to describe load growth in terms of multiple trends that persist for uncertain durations. For example, area load typically can grow at some steady rate for several years and then transition to a rapid growth spurt for a period of time. This suggests that a good way to model future load conditions is to characterize likelihoods of the possible trends and their durations.

A simple yet robust mathematical representation of such a phenomenon that captures the complexity and but requires relatively few parameters to be estimated is a Markov chain. (See also Section 6.5 for further details of the Markov model.) The Markov model is illustrated Figure 6.2. The top half of the figure shows that, over time, load may follow a sequence of different growth trends with varying durations. The bottom half of the figure illustrates how transition probabilities can be used to represent the uncertainty in the level of the next trend and the uncertainty in the time to the next trend. In this case, there are three trend states. For example, if load is growing at the medium rate, there is a 0.8 chance that the medium growth trend will persist into the next period, a 0.1 chance of a transition to the low trend and a 0.1 chance of a transition to the high trend. This characterization of load also represents the tendency to stay in a given trend, since the average time in a trend may be expressed as:

Average Time In Trend = 1 / (1 - p(staying in same trend)).

For example, for the given probabilities, the average time that the load growth will persist in the low state is 10 periods.

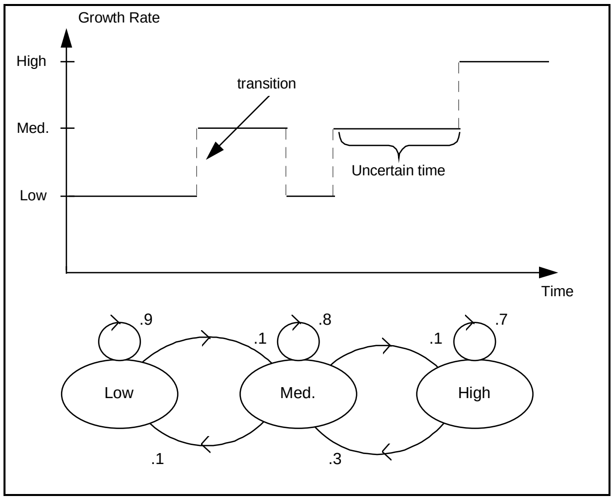

Figure 6.2 Probabilistic Dynamics of Load Growth

The Markov model differs from the more traditional approaches to load modeling. Those approaches are elementary, and are based on an estimate of average growth rate obtained by fitting a straight line through historical data. In many cases, no uncertainty in that estimate is provided. If uncertainty is incorporated into the estimate, the uncertainty is based on the value of the variance found from regression residuals. Such approaches ignore trends or correlations in the data and, in so doing, systematically underestimate the future load uncertainty. Moreover, the regression approach presupposes that additional observations provide no information. In the Markov setting, the likelihood of future load behavior depends on the current conditions, so observations modify forecasts.

The Markov model also responds to the fact that the distant past is rarely representative of the future, especially for local area planning. Growth trends are driven by real events such as changes in zoning laws and shifts in the local economic conditions. If the future is believed to be different than the past, the model allows the user to specify appropriate parameters that capture those beliefs.

Modeling Uncertainty in Time to the Next Decision

The Markov load model is used to determine the dynamic state probabilities that characterize the chance that the peak load will grow from any level to any higher level over a period of fixed length. From these state probabilities, the probability distribution of first passage times is found. The first passage time is the time required for the peak load to reach a specified higher level, starting from a given level, for the first time. The uncertainty in the time to the next investment decision is given by this probability distribution, based on the capacity added by the current investment. In practice, the model approximates the actual distribution with a discrete three-branch distribution that matches the first five moments of the actual distribution. This approximation is based on Gaussian quadrature.

Dynamic Programming Solution Technique

The solution procedure incorporates a principle known as dynamic programming. Dynamic programming succeeds by viewing a complex problem as a sequence of simpler problems. The decision variables in the problem are the various investment alternatives. As described above, the timing of the decisions is related to the load uncertainty and the incremental capacity provided by each investment.

We emphasize that dynamic programming is not an algorithm. Indeed, successful application of the principle is based on the art of mathematical modeling rather than on numerical aspects of problem solution. The unifying concept in applying dynamic programming to a problem is the notion of *state*. This notion is critical if application of the principle can avoid the so-called "curse of dimensionality": as the problem increases in complexity, the number of computations required for a solution grows exponentially. The investment strategy model uses a particular definition of the state to overcome this difficulty.

Formally, a *state* defines an *equivalence class* of trajectories. (A trajectory is a specific sequence of actions and uncertainty resolutions. For example, a trajectory could be: install two 500 kW engines at time zero; observe that the time to the next decision is 2.4 years; at that time, install a 20 MVA substation, which lasts for the next 12.6 years (and beyond), the end of the finite planning period of 15 years.) The desirable property of such an equivalence class is such that for all trajectories that have the same state, the future (i.e., the set of all decisions, costs and probabilities) appears identical viewed from any member of the class. Therefore, the future optimal decisions and cost computations need to be specified only once for all the members of the equivalence class. This greatly reduces the number of computations required to solve the problem.

We define the first component of the state as the collection of all investment decisions made through a given stage, independent of the order in which they were made. This makes sense as a state descriptor because electric utilities economically dispatch their technologies on hand so as to minimize operating cost, independently of their capital cost or when they were installed.

We define the second state component as the total capacity of the investments made along the trajectory. Since capacity is not added until it is needed, the total load in a state is equal to the total capacity installed, so that we need not define total load as a separate state component.

We define the third state component to be the most recent load growth rate. The Markov load model implies that the current load state is sufficient to describe future load growth behavior.

An example will illustrate the notion of state and how it can dramatically reduce the number of computations in a dynamic programming solution. Consider a problem with 3 stages, such that at each stage one of three decisions $\{d_1, d_2, d_3\}$ may be selected, followed by one of three possible load growth rates {l, m, h}. For simplicity, in this example, let us assume that the load levels are independent. An example of a complete trajectory is the sequence $\{d_2, h; d_2, l; d_1, m\}$. There are $(3 \times 3)^3 = 729$ unique trajectories for this problem. The total number of computations that would be required to solve the problem, without considering the effect of a state-based equivalence class, is 9 + 81 + 729 = 819 path-cost computations. Now consider what happens when a state variable, the total amount of capacity installed, is introduced. The first stage is as before. Entering the second stage, there are three unique capacity states, the result of the first stage decision, so that 3 x 9 = 27 computations are required. Entering the third stage, there are only six unique capacity states, so that 6 x 9 = 54 computations are required. The six unique capacity states are the pairs $\{(d_1, d_1), (d_1, d_2), (d_1, d_3), (d_2, d_2), (d_2, d_3), (d_3, d_3)\}$. Thus, using the concept of state, the total number of computations has been reduced to 9 + 27 + 54 = 90. For larger problems, the effect is far more striking. In a recent application, the application of the state concept reduced the size of the problem from roughly 100 million path calculations to several thousands.

This concludes our description of the highlights of the model. The remaining sections in 0 describe the different model elements in more detail.

6.2 Planning Period End-Effects

6.2.1 Overview

The purpose of this section is to discuss some aspects of the issue of terminal value or salvage value in the investment strategy model.

Perhaps the most important aspect of this issue is to define carefully what we mean by salvage value and why it is present in the strategy model. We are trying to solve the problem of investment planning *over the indefinite future* (in the sense that there is no *a priori* finite time beyond which the future consequences are unimportant). To solve that problem, we have formulated a finite time horizon model rather than an infinite time horizon model. Indeed, even an infinite time horizon model is a mathematical abstraction of the notion “the indefinite future.” Since the planning horizon in the model is finite for mathematical convenience, the analysis stops at some fixed time with assets in place of various vintages (based upon when they were first put in service) such that each may have some useful life remaining. The question we ask is how to value the remaining useful life of each of the assets in place at the end of the planning horizon.

The Strategy Model contains two different approaches to measure the value of the remaining useful life of installed assets. The user may select either approach for any particular analysis. The first approach is to value each asset as if it were sold at a market price. That price, which may be uncertain, is an input to the model, assessed by the user. This input, referred to as a Salvage Value or a “Terminal Lottery on the Capacity Price,” can be though of as the liquidation value of the assets of the utility, assuming that the utility were going out of business.

The second approach is called “Cost-to-Go.” This value is the cost associated with serving the load over the infinite horizon, assuming that the utility will stay in business for the indefinite future. The data required for this approach are the costs and real escalation rates of both capital investment and operations. The measurement procedure is based on two assumptions. First, we assume that any existing capacity will be retained for its lifetime and then replaced by future assets of the same capacity, but with capital and operating costs given by the user-supplied values. The replacement timing is independent of load growth, but the operating costs depend on the actual peak load. Second, we assume that any additional load growth beyond the capacity installed at the end of the finite planning period will be met by future investments at the time needed. These future investments are assumed to follow load perfectly, so there is no lumpy capital investment effect in the distant future. Capital and operating costs of this additional capacity investment are given by the user-supplied values.

The purpose of the Cost-to-Go approach is to simulate actual costs over the indefinite future, based on the user's beliefs about future conditions. This approach is designed to remove any artificial effects associated with choice of finite time horizon. (Compare this with the typical engineering-economic approach to investment planning, which either ignores the consequences of an arbitrary finite time horizon, or, what is far worse, is based on the erroneous but prevalent idea that expressing all capital costs in terms of annualized values deals with the problem correctly. It does not.) In particular, applying the infinite horizon cost-to-go should make the optimal policy costs for essentially the same policies be approximately equal, over all finite time horizons. Of course, that result depends on whether the user has selected the future costs and growth rates in a way that is consistent with the other cost and growth data supplied to the model.

The remainder of this section describes the theory and methods used in the model in greater detail.

Users not interested in technical details may skip the remainder of the section without loss of ability to apply either of the two approaches.

6.2.2 Definitions

We begin the detailed discussion of terminal or salvage value with a set of definitions. Perhaps the most familiar definition of salvage value is the following:

Definition 1. Salvage Value. The salvage value of an investment is the value of its remaining useful life.

The main issue raised by this definition is how to determine the value of the remaining useful life of an asset. The simplest, and apparently correct, interpretation is that the value of the remaining useful life of an asset is equal to the value having it provides over an infinite time horizon.

We argue that planning over an indefinite future is most accurately modeled as an infinite horizon problem. We do not solve that problem, however, and instead stop the planning period after some arbitrary, usually large, time has elapsed. We do this because it is computationally burdensome, if not impossible, to solve the infinite time horizon problem by our methods (without some heroic assumptions). To represent the behavior of the optimal investment strategy over the remaining indefinite future, which we do not model explicitly, we assign a salvage value to each asset that has remaining useful life at the end of the planning period. Therefore, the salvage value of an asset should be determined by solving the optimal investment problem over the remaining, infinitely long, time horizon. Of course, this is precisely what we cannot do. Therefore, we must produce methods that can estimate this value. Regardless of the method, however, the refined definition of salvage value under this argument is given by the following.

Definition 2. Salvage Value. The salvage value of an asset is its worth based on the subsequent optimal investment policy over the indefinite future.

The fundamental idea embodied in definition 2 is that having an asset of a particular vintage in place at the end of the finite time horizon influences the cost of the remaining “indefinite future”. This so-called “cost-to-go” depends on the bundle of assets held at the end of the finite planning period.

It is worth noting that the salvage value need not be unique. That is, one may express the salvage value of each of the bundles of assets as the sum of an arbitrary constant plus each bundle’s cost-to-go. Thus the salvage value can be changed (by changing the arbitrary constant) without changing the ordering of all policies.

The salvage value of the bundle can be easily decomposed and allocated to each of the individual assets, if that is necessary. In fact, it is actually the optimal cost-to-go over the indefinite future that we seek; salvage value is just a surrogate for that unknown cost. The allocation of the total cost-to-go among the assets in the bundle is completely arbitrary and is meaningless in the context of infinite horizon optimization.

Some analysts consider an apparently related alternative definition and express the following.

Definition 2′. Salvage Value. The salvage value of an asset is the value of the cash flows it is expected to generate over its remaining useful life. (EPRI Capital Budgeting Book, 11-24).

Definitions 2 and 2′ are not equivalent nor do they suggest the same evaluation procedure. As noted, Definition 2 requires (implicitly) the solution of an infinite time horizon problem to determine the salvage value. Definition 2′ asks less and requires only an estimate of the cash flows associated with the remaining life of the asset. We argue that Definition 2′ is not correct since it does not include the effect of the remaining life of an asset on future capital investments in other assets required in the indefinite future. The claim we are making is that the salvage value of an asset ought to include the effect the asset has on future investments. Of course, if the notion of “cash flows ...generate[d] over its remaining useful life” is widened somewhat, then definitions 2 and 2′ can be made equivalent.

It is also important to note that all the arguments advanced above are very different in spirit compared with the argument used to determine the residual value of a depreciated asset. Many analysts apply the following.

Definition 3. Salvage Value. The salvage value of an asset is the market value or trade-in value at the time of disposal; the market value is the actual amount for which the item can be sold at the time it is taken out of service.

This is the definition that one would find in an engineering economics textbook; it is often confused with the *book value* of an asset. The book value is based solely on the initial capital cost of an asset and a depreciation schedule to specify capital recovery. The EPRI Capital Budgeting Book (11-24) argues that the value derived from such considerations may not be a very good measure of the actual salvage value of an asset (as given by definitions 2 or 2′).

6.2.3 Measuring the Salvage Value of an Asset

The measurement problem is complex. There seems to be no universally accepted approach to determining salvage values under definitions 1, 2, and 3, above. We return to definition 2.

In fact, we are not trying to determine a so-called salvage value of an asset, but instead are trying to estimate the optimal cost-to-go over the remaining infinite horizon of an investment policy that begins at the end of the finite time horizon with a particular bundle of assets of various vintages in service. In other words, the detailed finite time horizon analysis investigated an investment policy that resulted in holding a specific bundle of assets at the end of the finite horizon. We know, from the Strategy Model, what the cost is to get to that point in time, following an investment policy that resulted in holding that specific bundle at that time. Next, we must estimate the cost that would result if we went further. That is, we must estimate how much an optimal investment policy would cost if that policy were designed to last over the indefinite future and if that policy were to begin with the specific bundle of assets in service at the end of the finite time horizon. Again, this is not the same as finding what most analysts might call a salvage value. Rather, we are trying to estimate the remaining cost of the infinitely-long-lived optimal policy. We do this because the dynamic model we are using to determine policies and costs is arbitrarily stopped at some convenient finite time and we wish to eliminate any effects of that arbitrary choice.

As one might expect, there is no generally accepted solution to this problem. There are two approaches entailing two different models that we can consider.

Terminal Lottery on Capacity Price.

The idea of this model is that at the end of the finite time horizon, we act as if we were going out of the electric utility business. We assume that there is some future value for units of capacity that have been chosen as investments, and we treat that value as an uncertain variable. We model the uncertainty and evaluate the final bundle of investments with respect to this value.

This model requires that the user specify a probability distribution or a point estimate on the price of capacity at the finite time horizon. This can be thought of as a rental price for capacity, expressible in $/kW per year. The terminal value (present value, year zero dollars) of an investment of capacity C, first put into service at time t, salvaged at time T, with useful life L is

$$V_T = C\ p(T)\ [\alpha^{T+1} - \alpha^{L+t+1}]/[1 - \alpha],\ \text{if } \alpha \neq 1$$

$$V_T = C\ p(T)\ [t+L-T],\ \text{if } \alpha = 1$$

where p(T) is the price of capacity (or the expected value of that price) at the terminal time T and the parameter α is the ratio [1+i]/[1+r], where i is the inflation rate and r is the discount rate.

The virtues of this model are that it is very easy to implement, it requires the user to estimate only one additional parameter, and it is transparent. The drawbacks of this model are that it assumes that the utility is going out of business, that capacity can be priced, and that future operating costs are irrelevant to the terminal value. To the extent that terminal value ought to be an approximation to the remaining infinite time horizon cost-to-go, these drawbacks can be considerable. Nevertheless, if utility planners believe that the world is changing and that a competitive market for capacity will develop, this approach may be the most appropriate model.

Terminal Value Model (Infinite Horizon Cost-to-Go).

This model estimates directly the cost-to-go over the remaining infinite time horizon, and responds directly to the drawbacks noted above with respect to the terminal capacity price approach.

In this model, we ask the user to specify the average capacity cost in the future, the average operating cost in the future, and the real escalation rate for each, such that each of these assessments is made with respect to the finite horizon time, T. (For example, one might ask what the cost of capacity will be in 20 years from now. A probability distribution would be an ideal response to this question. We would use the expected value of such a distribution.) This requires the user to supply (at least) four additional parameters. If the user cannot or chooses not to specify these values, the model can propose a set of default values that are averages of the existing alternatives.

The model requires that all investments be kept only for their separate lifetimes, compared with other approaches that assume that investments are replaced when they come to the end of their useful life. Such replacement will tend to accentuate the differences among the valuations of the different investment bundles that may not have been so pronounced over the finite time horizon. Also, identical replacements will tend to perpetuate specific choices that were made in the past for short-term considerations that will not necessarily be present in the future. It would seem to be unnecessary to repeat perpetually the investments made over the finite time horizon. Therefore, beyond that single lifetime of an investment, capacity is replaced with costs given by the assessed parameters.

The infinite horizon cost (present value) associated with a capacity investment made at time t, that has useful life L, such that t+L is greater than the finite horizon T, is equal to the present value of the investment made at time t, which includes capital cost and operating cost over the interval t to t+L, plus the present value of the cost (capital plus operating) of replacing that investment beginning at time t+L+1 over the remaining infinite horizon.

We are only going to report the cost-to-go until some fixed arbitrary time and until some fixed arbitrary level of load has been achieved. There are natural limits for both these values. The fixed arbitrary time is

$$\tau = T + L_{max},$$

where T is the finite planning period and L_{max} is the largest lifetime of any alternative. The fixed load level is

$$C^* = \min \{\max \{C(T)\}, S\},$$

where C(T) is the capacity installed on any trajectory at time T and S is the saturation level. (Note that if saturation is not present, $S = \infty$) Any replacement costs beyond these levels are identical for all trajectories and therefore do not affect the optimal policy. These common costs are large and tend to overwhelm the costs of the policy over the finite planning period, hence they have been eliminated from the output reports.

The expression for the capital cost is simple and straightforward. The analysis of the operating cost is somewhat more complex.

Capital Cost: We assume that the investment of capacity C lasts until time t+L, which exceeds the time horizon T. There are no capital costs incurred until t+L+1, using an end-of-period convention for discrete time payments. At that time, new assets must be rented annually, at the cost assessed by the user, say \$x /kW-yr., where x is assumed to be expressed in year zero dollars. (See Figure 6.3, below.) Let γ be the real growth rate of this rental price. Then the present value of the infinite stream of replacements is (α is the ratio [1+i]/[1+r], i is the inflation rate and r is the discount rate)

$$PW_{cap} = \sum_{k=t+L+1}^{\tau} \alpha^{k} \gamma^{k-T} Cx$$

$$= (\alpha\gamma)^{(t+L+1)} \gamma^{-T} Cx \, [1-(\alpha\gamma)^{\tau-(t+L)}] \, [1-\alpha\gamma]^{-1}$$

This value can be added to the capital cost incurred at time t when the capacity C is installed.

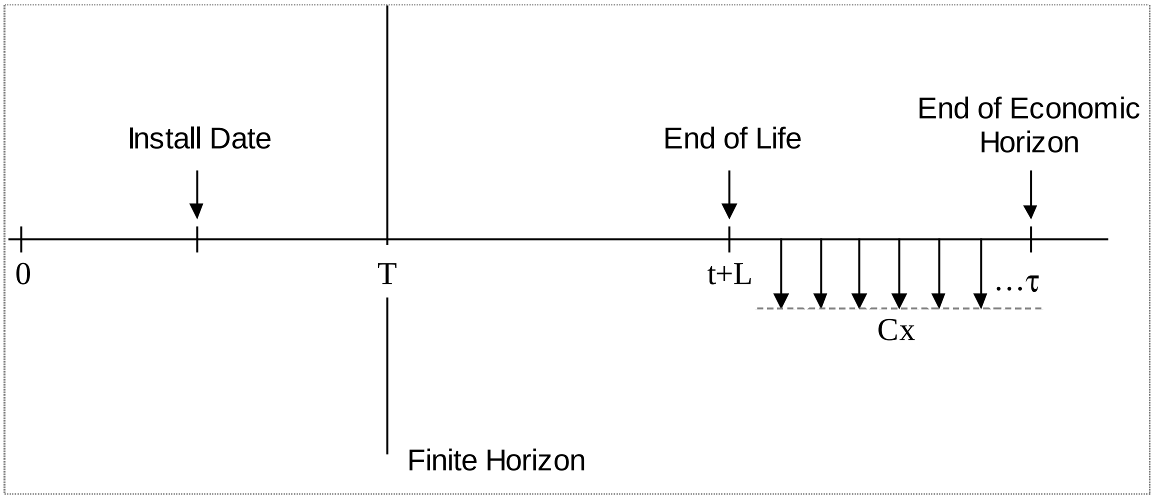

Figure 6.3 Capital Flows

Operating Costs: The objective is to determine a value that represents the future operating costs associated with the investment made at time t. This value can be added to the cost incurred at time t when the capacity C is installed.

There are three regimes over which the operating costs must be specified. (1) From time of installation, t, to the end of the horizon, T, the operating costs are found by the strategy model in its normal course of operation. (2) From the end of the horizon, T+1, to the time the investment's useful life expires, t+L, is treated immediately below. In this period, we shall assume that capacity is provided by only the investment made at time t and the fictitious asset characterized by the user. All other investments that may have been made subsequent to time t are not considered during this period. (3) From the time the investment's useful life has expired, t+L+1, to infinity is also treated below. This last period is somewhat simple, since we shall assume that capacity is provided by only the fictitious asset characterized by the user. This is consistent with the treatment of capital cost-to-go.

Load Growth

We must specify the anticipated load growth rate from time t, the time that capacity C was installed. Let

$L(t)$ = peak load at time t,

L_o = original peak load, which we interpret as the original wires capacity in the area,

$L_{avg}(T)$ = the expected value of the peak load at time T, given the peak load L(t) at t,

g = the future load growth rate, provided by the user, assumed to be in effect beyond T.

We can find the expected peak load at time T by running the load model forward until time T and taking the expectation. The peak load is assumed to grow at the user-supplied growth rate. Thus, the peak load at any time $\tau \geq T$ is

$$L(\tau) = L_{avg}(T)\,(1+g)^{\tau-T}.$$

Load Duration Curve

The operating costs are determined by the load duration curve. We will approximate the load duration curve over the interval (T+1, t+L) by a single straight line. We are only interested in that portion of the load duration curve that is above the original wires capacity, L_o.

At time T, the peak load is assumed to be $L_{avg}(T)$. The load duration curve is approximated by a straight line with slope equal to the slope of the first segment of the load duration curve specified by the user as part of the inputs to the model. (See Figure 6.4.) The input to the strategy model includes a specification of a piecewise linear load duration curve. We force the curve to pass through the point (0, 100%). The user then specifies the next point on the load duration curve as $(t_1, 1-p_1)$; e.g., 200 hours, 95% of peak load. The slope of that initial segment of the load duration curve is $m_1 = -p_1/t_1$. (See Figure 6.5.)

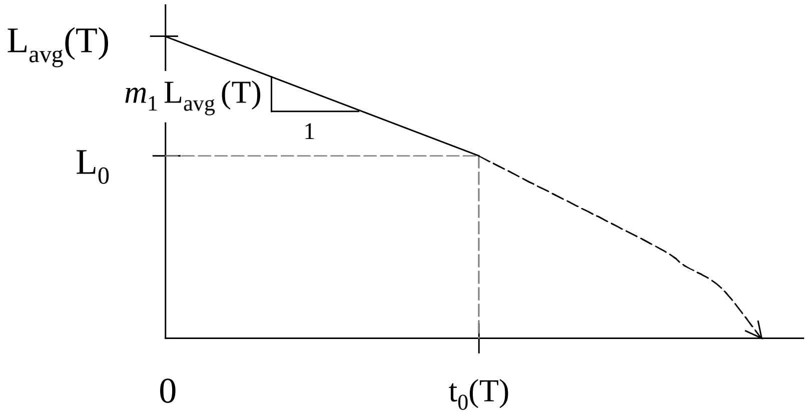

Figure 6.4 Approximate Load Duration Curve at T

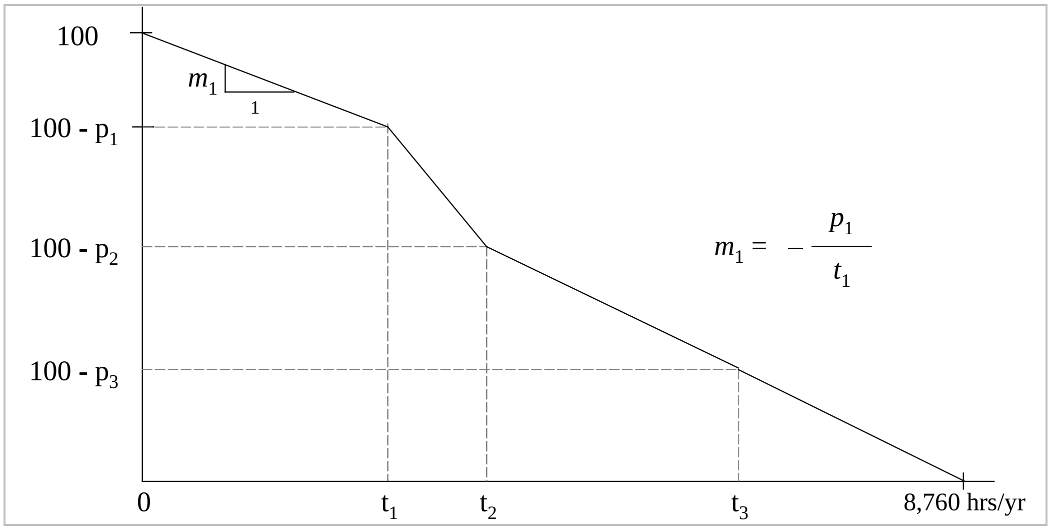

Figure 6.5 LDC Specified by User

This slope determines the amount of time the load is above the original wires capacity, L_o. Let this time be denoted by t_o. Then at time T,

$$t_o(T) = (1 - L_o / L_{avg}(T)) / m_1.$$

This equation applies for all other times $t \geq T$. In general, $t_o(t)$ is the intersection of the sloping line with the horizontal at L_o. (See Figure 6.4.)

Operating Costs

There are three operating costs that are required. Let

OC = the operating cost ($/kWh) of the investment installed at time t,

OC_{fut} = the operating cost ($/kWh) of the fictitious future asset described by the user,

OC_{sys} = the operating cost ($/kWh) of the bulk system; the cost of system energy.

Regime 2: from T+1 to t+L

In this regime, the operating costs are determined by integrating the approximate load duration curve. However, a closed form expression for the operating costs is difficult to find since there are several possibilities in each year. Therefore, the computation is expressed as a loop, using the above values, as follows.

Parameters required: m_1, g, L_o, $L_{avg}(T)$, OC, OC_{fut}, OC_{sys}, C.

```
STEP 0.  Set $TC_{T+1, t+L} = 0$.  Set $t = T+1$.

STEP 1.  COMPUTE  peak load $L(t) = L_{avg}(T) (1+g)^{t-T}$.

STEP 2.  COMPUTE  $t_o(t) = (L(t) - L_o) / m_1(t)$,  $m_1(t) = L(t)m_1$, where $m_1$ is the
         slope of the first segment of the load duration curve, expressed in percent
         change per hour.

STEP 3.  IF $C > L(t) - L_o$

IF C is load-following,

          COMPUTE  $TC(t) = (1/2)(OC)(L(t)-L_o)(t_o(t))$

IF C is non-load-following

          COMPUTE  $TC(t) = (OC)(C)(t_o(t)) - (OC_{sys})(t_o(t))[C + L_o/2 - L(t)/2]$

ELSE  ($C \leq L(t) - L_o$)

          IF  $OC > OC_{fut}$

          IF C is load-following,

                    COMPUTE  $TC(t) = (1/2) (C^2)(1/m_1)[OC]$
```

```
          IF C is non-load-following
                    COMPUTE  $TC(t) = (C^2)(1/m_1)[OC-(1/2)OC_{sys}]$
ELSE ($OC \le OC_{fut}$)
IF C is load-following,
          COMPUTE  $TC(t) = (C)(t_o(t))[OC] [1-(C/2)(1/[L(t)-L_o]) ]$
IF C is non-load-following
          COMPUTE  $TC(t) = (C)(t_o(t))[OC-(1/2)(OC_{sys}) (C) (1/[L(t)-L_o]) ]$
STEP 4.  SET $TC_{T+1, t+L} = TC_{T+1, t+L} + \alpha^t TC(t)$.
STEP 5.  IF t = t+L, STOP.
ELSE SET t=t+1 AND GOTO STEP 1.
```

The outer "if" statement in Step 3 tests whether the investment provides excess capacity. If so, then no future asset is required and running the investment penetrates below L_o. (See Figure 6.6.) If not, then capacity must be provided by some future asset as well. The inner "if" statement compares operating costs. If the operating cost of the installed asset is greater than that of the future asset, then relative to that future asset, the installed asset is a "peaker" and operates less than t_o. (See Figure 6.7.) Conversely, if the installed asset is cheaper to operate, it is relatively a "baseload" plant and operates for the full t_o, the time for which peak capacity is required. (See Figure 6.8.)

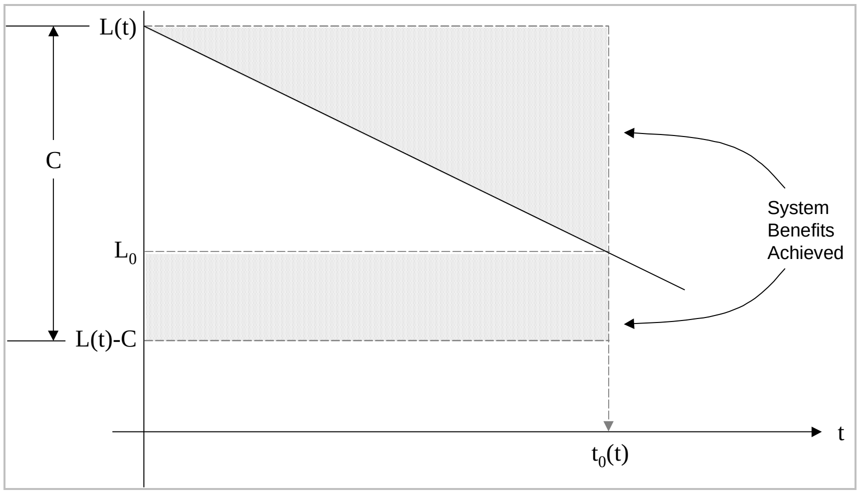

Figure 6.6 Operating Costs – IF C > $L_{(t)} - L_0$

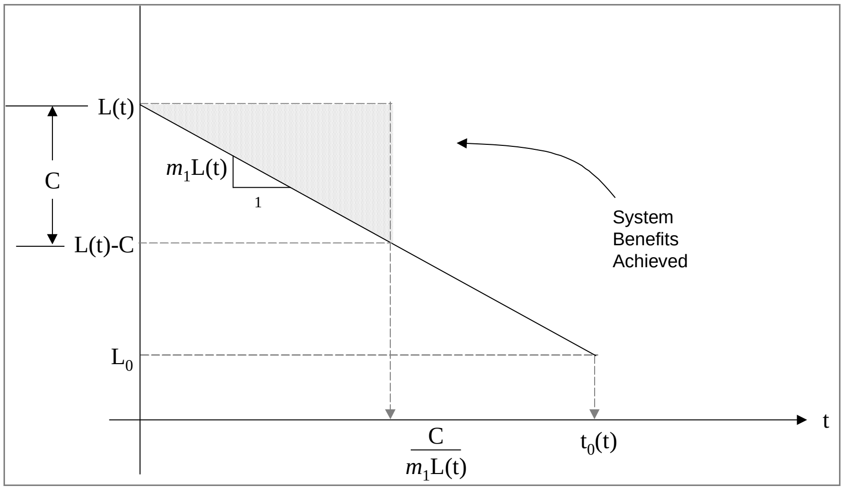

Figure 6.7 Operating Costs – IF "Peaker"

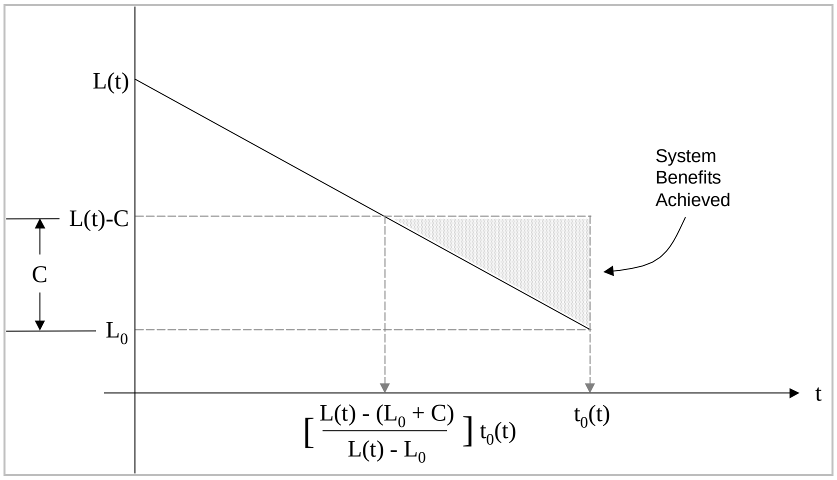

Figure 6.8 Operating Costs – IF "Baseload"

Regime 3: from t+L+1 to τ

In this regime, we assume that the capacity is provided by the fictitious future asset described by the user. We further assume that all fictitious assets are load following. The hours of operation of the fictitious asset will depend on whether the peak load is sufficiently large so as to exhaust the capacity C provided by the investment that is replaced by the fictitious asset. The logic is identical to that of STEP 3 in the computation over Regime 2: from T+1 to t+L. In this case, it is possible to find closed form solutions for the costs in all cases.

For $k \in [t+L+1, \tau]$,

```
    IF $C > L(k) - L_o$
        THEN $TC(k) = (OC_{fut})(1/2)(L(k) - L_o)(t_o(k))$
    ELSE $(C \leq L(k) - L_o)$
        THEN $TC(k) = (OC_{fut})(C)(t_o(k)) (1 - (C/2)(1/[L(t)-L_o]))$
```

where

$t_o(k) = (L(k) - L_o) / m_1(k)$, where $m_1(k) = L(k)m_1$.

It is possible to find k* such that

$C > L(k) - L_o, k<k^*$

and

$$C \leq L(k) - L_o, k \geq k^*.$$

Since $L(k) = L_{avg}(T)\,(1+g\,)^{k\,-T}$,

it follows that k* is such that

$$C = L_{avg}(T)(1+g\,)^{k^*-T} - L_o,$$

hence

$$k^* = \ln\,([C+L_o](1+ g\,)^T / L_{avg}(T)) / \ln\,(1+ g\,).$$

Then set

$t^* = \max\,\{\mathrm{int}(k^*), \mathrm{int}(t+L-1)\}$, where int(k*) is the greatest integer in k*+1.

If $t^* = \mathrm{int}(t+L-1)$, then the peak load is greater than $L_o + C$ when the installed capacity C has exceeded its lifetime. Else $t^* > t+L$, so that the load is growing very slowly, on average, load growth does not exhaust the installed capacity C at the end of its life, and the only installation during T was C.

We may then write

$$TC_{t+L+1,\tau} = \eta^{-T}\,OC_{fut}/(2m_1) \bullet$$

$$\Big\{ (L_{avg}(T)/(1+g\,)^T) \Big([\alpha\eta\,(1+g\,)]^{\mathrm{int}(t+L-1)} - [\alpha\eta\,(1+g\,)]^{t^*} \Big) / (1- [\alpha\eta\,(1+g\,)])$$

$$- 2L_o \Big([\alpha\eta\,]^{\mathrm{int}(t+L-1)} - [\alpha\eta]^{t^*}\Big) / (1- \alpha\eta\,)$$

$$+(L_o^2(1+g\,)^T / L_{avg}(T)\,) \Big([\alpha\eta\,/(1+g)]^{\mathrm{int}(t+L-1)} - [\alpha\eta\,/(1+g\,)]^{t^*}\Big) / (1- [\alpha\eta\,/(1+g\,)]) \Big\}$$

$$+ \eta^{-T}\,OC_{fut}\,(\,C\,)(1/m_1) \bullet$$

$$\Big\{\Big([\alpha\eta\,]^{t^*+1} - [\alpha\eta]^{\tau+1}\Big) / \Big(1- \alpha\eta\,)$$

$$-(L_o + C/2)(1+g\,)^T / L_{avg}(T)\,)\Big([\alpha\eta\,/(1+g\,)]^{t^*+1} - [\alpha\eta\,/(1+g\,)]^{\tau+1}\Big) / (1-[\alpha\eta\,/(1+g\,)]\,)\Big\}$$

Cost-to-go Associated With Load Growth

Some investment trajectories require greater capacity than others. Load growth beyond the finite horizon T may exceed the capacity of the investments made over that finite period T. To account for both these effects, in order to compare fairly any alternative strategies, two additional costs, capital and operating, must be added to the terms specified above.

Let C(T) = total installed capacity at time T, the capacity at the end of the finite time horizon, along any investment trajectory, including original wires capacity.

Let max {C(T)} = largest total installed capacity at time T.

Let t_o' = the time when C(T) will be exhausted by load growth, $t_o' \geq T$.

Capacity must be added beyond t_o', to serve the load above C(T). All investment trajectories with actual C(T) less than the maximum can be brought to this maximum value by future investments in capacity as load grows beyond C(T). If load does not grow beyond this value, no future investments need to be made along that trajectory.

Let g = the user-supplied growth rate that is supposed to occur beyond T. Then $C(T)(1+g)^{k-to'}$

is the load at time $k > t_o'$.

Let k**= the time at which the load would grow beyond the installed capacity C(T) to the maximum value of installed capacity, max {C(T)}; i.e., k** satisfies

$C(T)(1+g)^{k^{**}-to'} = \max\{C(T)\}$.

Then the total capacity cost from t_o' to τ is

$$\sum_{k=to'}^{\tau} \alpha^k \gamma^{k-T} x \left[\min\{C(T)(1+g)^{k-to'}, \max\{C(T)\}\} - C(T)\right]$$

$$= \sum_{k=to'}^{k^{**}} \alpha^k \gamma^{k-T} x \left[C(T)(1+g)^{k-to'} - C(T)\right]$$

$$+ \sum_{k=k^{**}+1}^{\tau} \alpha^k \gamma^{k-T} x \left[\max\{C(T)\} - C(T)\right]$$

$$= x\, C(T)(\alpha\gamma)^{to'} \gamma^{-T} \Big[[(1-[\alpha\gamma(1+g)]^{k^{**}+1-to'})/(1-\alpha\gamma(1+g))] - [(1-[\alpha\gamma]^{k^{**}+1-to'})/(1-\alpha\gamma)] \Big]$$

$$+ x\,[\max\{C(T)\} - C(T)]\, \gamma^{-T} \Big[([\alpha\gamma]^{k^{**}+1} - [\alpha\gamma]^{\tau+1})/(1-\alpha\gamma) \Big]$$

The total operating cost is based on the area above C(T) and below max {C(T)}in the load duration curve. Let m_1 represent the slope of the top segment of the user-supplied load duration curve. Then the operating cost in each year k such that the peak load L(k) is greater than C(T) but less than max {C(T)} is

$$(1/2m_1(k))\ (L(k) - C(T))^2\ OC_{fut}\ .$$

with

$$L(k) = C(T)\ (1+g\ )^{k\ -\ to'}$$

and

$$m_1(k) = m_1\ C(T)\ (1+g\ )^{k\ -\ to'}$$

When the peak load is greater than max {C(T)} the operating cost is

$$OC_{fut}\ (\max\{C(T)\}- C(T))\ [L(k) - C(T)](1/m_1(k))\ \Big[1 - (\max\{C(T)\}- C(T))/\ [L(k) - C(T)]\ \Big]$$

This operating cost is based on a trapezoidal area under the load duration curve.

Therefore, the total operating cost is

$$\sum^{k^{**}}_{k=\ to'}\ OC_{fut}\ \alpha^k\ \eta^{k-T}\ (1/2m_1(k))\ (C(T)\ (1+g\ )^{\ k\ -\ to}\ - C(T))^2$$

$$+ \sum^{\tau}_{k=\ k^{**}+1} OC_{fut}\,\alpha^k\ \eta^{k-T}\ (\max\{C(T)\}- C(T))[L(k) - C(T)](1/m_1(k))\Big[1 - (\max\{C(T)\}- C(T))/\ [L(k) - C(T)]\Big]$$

$$= OC_{fut}\ (1/2m_1)\ (C(T))\ \eta^{-T}\ \ (\alpha\,\eta)^{\ to'}$$

$$\Big\{\ (1 -\ [\ \alpha\eta(1+ g\ )]^{k^{**}+1-\ to'}\ )/(1-[\alpha\eta(1+ g\ )])$$

$$-2\ \Big(\ (1-[\alpha\eta]^{\ k^{**}\ +1-\ to'}\ )/\ (1-\alpha\eta)\ \Big)$$

$$(1 -\ [\ \alpha\eta/(1+ g\ )]^{k^{**}\ +1-\ to'}\ )/(1-[\alpha\eta/(1+ g\ )])\Big\}$$

$$+ OC_{fut}\ (1/m_1)\ [\max\{C(T)\}\ -C(T)]\ \ \eta^{-T}\ \Big\{([\ \alpha\eta]^{k^{**}+1} - [\alpha\eta]^{\tau+1})/\ (1-\alpha\eta)$$

$$- (1+ g)^{to'} (1+ [\max(C(T)\}-C(T)]/2C(T)) ([\alpha\eta/(1+ g)]^{k^{**}+1} - [\alpha\eta/(1+ g)]^{\tau+1})/ (1- \alpha\eta/(1+ g)) \Big\}$$

That completes the specification of the cost-to-go computation.

6.2.4 Modifications in the Presence of Area Peak Load Saturation

When saturation occurs, load never increases beyond the user-supplied saturation level, S, for all future time. This causes some of the results stated above for the Terminal Value Model to change somewhat.

Capital Costs

Capital costs associated with investments made before T do not change in any way because of saturation. The assumption that capacity investments are replaced over the infinite horizon even if the total installed capacity is somewhat greater than saturation remains in force. The capital costs that are incurred because load growth beyond T causes load to exceed the total installed capacity at the end of the finite time horizon change somewhat under saturation.

Operating Costs

As before, there are three regimes of operating cost to consider. Saturation adds the further complication that there are, in addition, three cases of time to achieve saturation (TTS) that must be analyzed.

Let t = the time that an asset of capacity C was installed.

Let L = the lifetime of that asset.

Let T = the finite horizon.

Let TTS(t) = the time to saturation, measured from time t, TTS < ∞ if saturation is present and if saturation occurs along an arc emanating from t. The argument t in the definition of TTS will be suppressed in what follows.

Case 1: $t + TTS \leq T$

Case 2: $T < t + TTS \leq t + L$

Planning Period End-Effects

Case 3: $t + L < t + TTS$

Case 1

In this case, the area saturates before the end of the finite time horizon. The load grows to S and remains there forever.

<u>Regime 1: from t to T</u>

In this regime, we find first the effective growth rate g_{eff} from t to TTS by solving:

$$\log (1 + g_{eff}) = [\log (S/L(t))] / TTS$$

Then we set the load in this regime to $L(k) = \min \{ L(t) (1 + g_{eff})^{k-t} , S \}$ for $k > t$.

Once the load is known, logic identical to that which we have been using for operating costs along any arc applies.

<u>Regime 2: from T+1 to t+L</u>

In this regime, under case 1, the load is constant $L(k) = S$. The logic described above, for the non-saturation case, can be used with no changes. Notice that the first segment of the slope of the saturated load duration curve is $m_{1,S} = m_1 S$, where m_1 is the slope of the first segment of the user-supplied load duration curve expressed as percent of peak load per hour.

<u>Regime 3: from t+L+1 to τ</u>

In this regime, we assume that the capacity is provided by the fictitious future asset described by the user. The hours of operation of the fictitious asset used to replace the capacity C will depend on whether there are other assets installed when saturation occurs.

$$\text{Operating Cost} = \eta^{-T} (S - L_o) (\min\{C, S-L_o\}) (1/[2m_1 S]) (OC_{fut}) (([\alpha\eta]^{t+L+1} - [\alpha\eta]^{\tau+1}) / (1- \alpha\eta))$$

Since the arc saturates, the three operating costs, one for each regime, may be added to comprise the cost along the arc.

Case 2

In this case, the area saturates after the finite time horizon but before the end of the useful life of the last investment. This causes no essential difference compared with case 1.

Regime 1: from t to T

As in case 1, we find first the effective growth rate g_{eff} from t to TTS by solving:

$$\log (1 + g_{eff}) = [\log (S/L(t))] /TTS$$

Then we set the load in this regime to $L(k) = L(t) (1 + g_{eff})^{k-t}$ for $t < k \leq T$.

The identical logic applies along the arc until T.

Regime 2: from T+1 to t+L

In this regime, the load is $L(k) = \min \{ L(t) (1 + g_{eff})^{k-t} , S \}$ for $T+1 < k \leq t + L$. The load saturates at t + TTS and remains at S. Again, the logic stated above for the non-saturated case can be used with no changes. Notice that the first segment of the slope of the saturated load duration curve is $m_{1,S} = m_1 L(k)$, where m_1 is the slope of the first segment of the user-supplied load duration curve expressed as percent of peak load per hour.

Regime 3: from t+L+1 to τ

Since the load is identically S in this regime, the results for case 1 apply here as well.

Case 3

In this case, the area saturates after the lifetime of the last investment. Therefore, the cost in the third regime is somewhat more complex than in the previous cases.

Regime 1: from t to T

As in the above cases, we first find the effective growth rate g_{eff} from t to TTS by solving:

$$\log (1 + g_{eff}) = [\log (S/L(t))] /TTS$$

Then set the load in this regime to $L(k) =L(t) (1 + g_{eff})^{k-t}$ for $k > t$.

Once the load is known, logic identical to that which we have been using for operating costs along any arc applies.

Regime 2: from T+1 to t+L

In this regime, the load is also $L(k) = L(t) (1 + g_{eff})^{k-t}$ for $T+1 < k \leq t + L$. Again, the logic stated above for the non-saturated case can be used with no changes. Notice that the first segment of the slope of the saturated load duration curve is $m_{1,S} = m_1 L(k)$, where m_1 is the slope of the first segment of the user-supplied load duration curve expressed as percent of peak load per hour.

Regime 3: from t+L+1 to τ

Now the load is varying in this regime: $L(k) = \min \{ L(t) (1 + g_{eff})^{k-t} , S \}$ for $T+L+1 \leq k < \infty$. To find a closed form expression for the operating costs, define $\tau' = \mathrm{int}(t + TTS - 1)$, the greatest integer in t + TTS. Then, load grows until τ' and is saturated from $\tau' + 1$ to ∞. However, the costs are only computed until τ. There is no necessary relationship between τ and τ'.

The total operating cost is given by:

$$TC = \eta^{-T} C \, (OC_{fut}) \, (2/m_1) \bullet$$

$$\Big\{ L(t) \left([\alpha\eta(1+g_{eff})]^{t+L+1} - [\alpha\eta(1+g_{eff})]^{\min\{\tau, \tau'\}+1} \right) / (1- \alpha\eta(1+g_{eff}))$$

$$- 2 L_o \left([\alpha\eta]^{t+L+1} - [\alpha\eta]^{\min\{\tau, \tau'\}+1} \right) / (1- \alpha\eta)$$

$$+ [L_o^2/L(t)] \left([\alpha\eta /(1+g_{eff})]^{t+L+1} - [\alpha\eta/ (1+g_{eff})]^{\min\{\tau, \tau'\}+1} \right) / (1- [\alpha\eta / (1+g_{eff})]) \Big\}$$

$$+ \eta^{-T} (S-L_o) (1/2m_1S) OC_{fut} \, (\min\{C, S-L_o\}) \left([\alpha\eta]^{\min\{\tau, \tau'\}+1} - [\alpha\eta]^{\tau+1} \right) / (1- \alpha\eta)$$

This completes the discussion of the effect of saturation on Terminal Value (Infinite Horizon Cost-to-go)

6.3 Modeling Investment Leadtimes

6.3.1 Overview

We define leadtime as the time that must elapse between committing to an investment and actually having the investment available. For the population of nuclear power plants operating in the early 1980s, the average leadtime, defined as the time between construction permit issue date and on-line date, was more than five years. Although the nuclear power plant situation was unique and perhaps somewhat extreme, leadtimes can influence investment strategy.

The simplest way to think about the leadtime is that it affects the timing of cash flows. An investment with nonzero leadtime requires cash flows to occur earlier than the time-to-the-next decision logic would suggest if leadtime were zero. Leadtime can be thought of as a requirement for advance payment, before the time that the asset is actually needed to satisfy demand. If the user bears that in mind, then supplying data to the Strategy Model should be straightforward. Note that leadtime is supplied for each investment alternative and that is the only data required with respect to analyzing the effect of leadtimes on investment strategy.

If the leadtimes for all investments were identical, the phenomenon would present no modeling difficulty. The time to the next decision is based on the load growth and the capacity currently installed. If all leadtimes were the same, the time to commit to the next decision would be, for all alternatives, equally in advance of the time at which the installed capacity is exhausted. The equal advance is given by the common leadtime. (If we let t_L denote the common leadtime, measured in years, and T_1 denote the [calendar] time at which the next decision would take place, due to capacity exhaustion, then the effect of the leadtime is to advance the investment to the [calendar] time $T_1 - t_L$.) Because the leadtimes can vary by alternative, this approach will not suffice. The difficulty is that there is no obvious way to adjust the time to the next decision for leadtime without knowing what that next decision will actually be.

Since the Strategy Model can easily solve the problem with no leadtimes, the modeling approach taken is to convert the actual problem with leadtimes into an equivalent problem with no leadtimes, and let the Strategy Model solve that latter problem. The method used is essentially a change of variables and is described in the following sections.

Users not wishing to read about the details of the change of variables may skip the remainder of this section without loss of ability to use leadtimes in the Strategy Model.

6.3.2 Modeling Approach

The essential idea of the modeling approach is to adjust the input parameters so that the algorithm appears to be solving a zero-leadtime problem while in fact (a) evaluating the alternatives as if actual leadtimes were included and (b) measuring the present value of all cash flows as if actual leadtimes were included.

The modeling approach succeeds subject to the following assumptions.

[A.1] Leadtime is short enough so that load growth during the leadtime, from commitment to the time that the investment goes on-line, can be modeled as certain. That certain load growth rate is assumed to be the average load growth rate from the initial time to the time at which capacity is needed.

[A.2] Leadtime for any alternative is smaller than the minimum time to the next decision for any initial conditions of load and current set of alternative investments.

We describe the approach in terms of the model structure as follows.

Let T_o = time at which capacity is needed. This is the time at which the peak load is equal to the total installed capacity provided by investments and the original infrastructure. This is calendar time. Let the installed capacity or, equivalently, the load level, be denoted by L_o. This is the load level achieved at time T_o.

Let $\{p^j_i : i=1,2,...,n\}$ = discrete probability distribution on the time to the next decision if alternative j were selected at time T_o. That is, for each alternative j, with fixed capacity C_j, the time to the next decision is the time it takes for load growth to exhaust the capacity C_j. This is modeled discretely with n possible outcomes. In the current version of the model, n, the number of branches in the discrete approximation of the probability distribution of the time to the next decision is three.

Let $\{t^j_i : i=1,2,...,n\}$ be the times for alternative j to be exhausted that occur with the corresponding probabilities.

Let $\{V^j_i : i=1,2,...,n\}$ = the set of values of cost-to-go at the time of the next decision, given that alternative j were selected at time T_o.

If there were zero leadtime, the model would evaluate the decision at time T_o with respect to the expected present value of the cost-to-go, which is given by (suppressing the superscripts in the equation)

$$\sum_i p_i V_i / (1+r)^{ti}$$

where r is the discount rate. That is, every cost-to-go is discounted back to time T_o using the times $\{t^j_i : i=1,2,...,n\}$, which are the first passage times corresponding to the probabilities $\{p^j_i\}$.

This is the solution to the problem if there were zero leadtime.

Now suppose that the leadtime for alternative j is Δ_j. The actual decision would then be made at (calendar) time $T' = T_o - \Delta$, suppressing the subscript j.

One effect of making the decision earlier is that there is more uncertainty about the time that the capacity of alternative j will be exhausted by load growth. The way that effect is treated is as follows.

Under [A.1], we can state that the load level at T′ is given by $L_o' = L_o / (1+ \alpha)^\Delta$, where α is the average load growth rate up to T_o, which is given by the load model. Thus, L_o' is just the scaled-down value of the load that would be observed prior to exhausting the capacity of all the installed resources.

Further, we can also state, under [A.1], that the time to the next decision <u>relative to T′</u> is the time it takes to exhaust the difference $L_o - L_o'$ plus capacity C_j. That is, the amount of capacity available to absorb future load growth at the earlier time T′ is the amount left in the current set of installed assets after meeting the scaled-donw peak load plus any capacity that is added at time T_o.

Thus, in effect the capacity $C_j' = C_j + [\, L_o - L_o']$ is being added at time T′.

Now, if we run the load model to find the time to next decision from the earlier time T′, using C_j' we determine a probability distribution $\{p^j_i{}': i=1,2,...,n\}$ with corresponding times $\{t^j_i{}': i=1,2,...,n\}$. The primed values are the parameters that describe the actual situation at the earlier time T′, defined by leadtime Δ_j.

All that remains is to force the model to use these parameters while retaining the zero leadtime logic that is currently used to compute expected cost-to-go. This can be done as follows.

The actual decision is made at time $T' = T_o - \Delta_j$. The immediate capital cost of investment j is K_j, which is incurred at the time of decision. The costs-to-go at the time of the next decision remain the values $\{V^j_i : i=1,2,...,n\}$, and those occur at times $\{T' + t^j_i{}': i=1,2,...,n\}$. Note that the present worth (<u>at time T_o</u>) of the decision made at time T′ is equal to

$$\left(K + \sum_i p_i' \, V_i / (1+r)^{t_i'}\right) \; (1+r)^{\Delta}$$

which can be expressed equivalently as

$$K(1+r)^{\Delta} + \sum_i p_i{}' V_i / (1+r)^{t_i{}'-\Delta}$$

Hence, the equivalent set of parameters for calculating the expected cost-to-go at time T_o is the actual capital cost multiplied by $(1+r)^{\Delta}$, the actual probabilities given by the load model applied to the capacity addition $C_j{}' = C_j + [\, L_o - L_o{}']$, which are the values $\{p^j_i{}': i=1,2,...,n\}$, and the corresponding times decreased uniformly by the leadtime, $\{t^j_i{}' - \Delta_j : i=1,2,...,n\}$. This set of parameters will give the correct value of the expected cost to go at the time T_o.

In summary, the procedure is to make decisions in the model at the time T_o, as would be the case if lead time were zero. However, at that time, the capital cost of each alternative must be inflated by the factor $(1+r)^{\Delta}$. In addition, the time to the next decision is found by running the load model and moment matching procedure as if the capacity added were $C_j{}' = C_j + [\, L_o - L_o{}']$ at the time when the load level was equal to $L_o{}' = L_o / (1+ \alpha)^{\Delta}$. The load model will generate a set of probabilities and corresponding times to the next decision. The probabilities are to be used to find the expected cost-to-go, but the future costs are discounted with respect to the times found by the load model decreased by the leadtime, the values $\{t^j_i{}' - \Delta_j : i=1,2,...,n\}$. Furthermore, the calendar time for each of the next decisions is given by $T_i = T' + t^j_i{}' = T_o + t^j_i{}' - \Delta_j$. The procedure then restarts at the times $\{T_i : i=1,2,...,n\}$.

6.3.3 Discussion

The procedure is in essence a change of variables that makes the model behave as if there were no leadtime. We know the leadtime and adjust the capacities and times to reflect the consequences of nonzero leadtime.

An example may help to illustrate the ideas.

Suppose the following parameters are given:

$C = 1000$,
$L_o = 10{,}000$,
$\alpha = .01$,
$r = .05$,
$K = \$1$ M,
$\Delta = 2$.

Suppose that the load model were run for the given values of L_o and C. Let the resulting probability distribution on the time to the next decision be given the by probabilities {.20, .50, .30} with corresponding times {12, 10, 8}.

Now, apply the procedure above. The initial load level is $L_o' = L_o / (1+ \alpha)^{\Delta} = 10{,}000/(1.01)^2 = 9803$ and the effective capacity added is $C' = C + [L_o - L_o'] = 1000 + [10{,}000-9803] = 1197$. Suppose that the load model run for these values yields the probability distribution on the time to the next decision given by probabilities {.25, .40, .35} with corresponding times {15, 12, 7}, for example.

The model would then use the probabilities {.25, .40, .35} as given by the load model, but modify the times so that they become {13, 10, 5}. These are the times to the next decision relative to the time at which the capacity of the prior set of installed investments was exhausted. Furthermore, the value of the capital cost would be transformed to $1.1025 M. The present value of the costs would then be given by $1.1025 plus the expected discounted cost to go, where the discounting would be based on the times {13, 10, 5}.

It is worth noting that the effect of this procedure is to spread the tails of the leadtime distribution compared to that gotten with no leadtime. This is exactly the phenomenon that we wish to model.

6.4 After-Tax Cash Flows

6.4.1 Overview

This section describes how revenue requirements, taxes, and depreciation are treated in the strategy model. We shall first list the new inputs required of the user, with some suggested defaults that could be used as prompts. We shall then develop the equations required to define the appropriate cash flows. This development is based on Section 11 in the EPRI Capital Budgeting book, TR-104369.

The user may find the discussion of the inputs of interest. The subsequent sections provide the mathematical details of the various computations. These sections may be omitted without loss of ability to apply the Strategy Model.

6.4.2 Inputs to the Model

1. Method of deferred tax computation: the user must specify either the *flow-through method* or the *normalization method.* It is a matter of complete indifference to the model which is selected. Equations are given below for each method. Suggested default: flow-through method.

2. τ = tax rate. This is a percentage. The default is $\tau = 0.34$, the federal tax rate for large corporations.

3. i = cost of debt. This is a percentage. The default value is $i = 0.08$, a current representative corporate borrowing rate.

4. k = cost of equity. This is a percentage. The default value is $k = 0.10$, a current representative rate of return on utility stocks.

5. d = percent debt financing. This a percentage. The default value is $d = 0.50$, which is a typical value for utilities.

6. D_t = Federal Tax depreciation schedule. The default is the 15-year property MACRS schedule, which is the following sixteen percentages: {.05, .095, .0855, .0770, .0693, .0623, .059, .059, .0591, .059, .0591, .059, .0591, .059, .0591, .0295}. Each number is the percent of initial investment that can be depreciated in each year, starting in the taxable year in which the investment was made.

7. f_t = straight-line depreciation. Using the half-year convention, this is the sequence of L+1 numbers, {1/2L, 1/L, ..., 1/L, 1/2L}, which clearly sum to one. These numbers define the straight-line depreciation allowance apparently used by most utilities. L is the useful life of the asset.

6.4.3 Annual Revenue Requirements.

To find the annual revenue requirements, it is necessary to compute the following values.

Let RR_t = revenue requirements in year t.

Then

$$RR_t = E_t + DEP_t + ROI_t + TAX_t$$

where

E_t = operating costs in year t, which comes from the model;

$DEP_t = f_t\ K_o$, straight-line depreciation, where K_o = initial capital investment;

$ROI_t = INT_t + ROE_t$, the return on investment,

where

$INT_t = id\ BASE_{t-1}$, interest expense

and

$ROE_t = k(1-d)\ BASE_{t-1}$, return on equity

and

$BASE_{t-1} = NET_{t-1} = K_o - \sum_{j=1}^{j=t-1} DEP_j$, t=2, 3, ..., the unrecovered capital investment;

$$TAX_t = [\tau\, ROE_t - \tau\, (D_t\, K_o - DEP_t)\,]/\,(1-\tau) \quad \text{if flow-through}$$

and

$$TAX_t = \tau\, ROE_t \,/\, (1-\tau) \quad \text{if normalization.}$$

6.4.4 Finding The Present Worth of the Revenue Requirements.

Given k, i, d, as above, the discount rate is determined: $r = di + (1-d)k$.

This is the *before-tax discount rate.*

The after-tax discount rate is given by $r^* = r - \tau\, d\, i$.

The present worth of the revenue requirement stream is found by discounting $\{RR_t\}$ by the after-tax discount rate, r^*.

6.4.5 After-Tax Cash Flow

The after-tax cash flows are based on the initial investment (a real, lumpy cash flow), expenses, taxes, and interest payments. The cash flow is thus a cost. The MACRS depreciation schedule is used to find the decrease in taxes due to capital investment.

Let $ATCF_t$ = after-tax cash flow in year t.

Then

$$ATCF_t = (1-\tau)\, E_t + K_t - T_t + (1-\tau)\, INT_t$$

where

E_t = operating costs in year t, which comes from the model;

K_t = capital investments made in year t;

T_t = decrease in taxes in year t. This will be based on the depreciation allowance for capital investments made in prior years and is given by

$$T_t = \tau \sum_{m=0}^{t-1} K_m D_{t-m}$$

INT_t = interest expense in year t, as above.

The present worth of $\{ATCF_t\}$ is found by discounting by r*.

6.4.6 Net Cash Flow (Profitability)

The net cash flow is the difference between after-tax revenues collected (based on actual sales) and the after tax cash flow (essentially the after-tax net costs of operations and investment, as described above).

Let NCF_t = after-tax cash flow in year t.

Then

$$NCF_t = R_t - ATCF_t$$

where

$$R_t = p_t S_t$$

where

p_t = price of energy sold, \$/kWh

and

S_t = energy sold.

The present worth of $\{NCF_t\}$ is found by discounting by r*.

6.4.7 Discussion

There is a certain arbitrariness associated with revenue requirements. Revenue requirements comprise a way of recovering costs that is not unique. In fact, there are an infinite number of cost recovery streams that produce the same present value of revenue requirements. However, for a given problem there is a unique set of actual cash flows that give rise to the need for revenue requirements. That is, there is a relationship between the present value of actual cash flows and the revenue requirements. Given the arbitrary nature of annual revenue requirements, it may be better to report only the present value of revenue requirements, rather than reporting a set of annual numbers given by $\{RR_t\}$. Moreover, we are trying to find a method to select capital investments. It seems that the net cash flow sequence, if revenues are important, or the after-tax cash flow sequence provide a better measure of the net worth or costs of alternative investments.

Indeed, following the capital budgeting book, a simple *gross-up* rule (see pp. 11-22 ff.) permits one to find the present value of the revenue requirements given the present value of the after-tax cash flows. That is:

$$PV\{RR\} = PV\{ATCF\}/(1-\tau).$$

Therefore, we implemented the following.

(1) The model calculates the annual after-tax cash flows and the present value of the after-tax cash flows.

(2) The model calculates the net cash flows and the present value of the net cash flows.

(3) The model permits the user to select which of those two cash flows to use as the optimization criterion.

(4) The model calculates the present value of the revenue requirements based on the gross-up rule applied to the after-tax cash flows.

(5) If a stream of revenue requirements is requested by the user, this may be expressed as a stream as follows:

$$RR_t = PV\{RR\}\ (A_1\ |P\ r, g, n)$$

where

$$(A_1\ |P\ r, g, n) = (r - g) / [1-(1+g)^n (1+r)^{-n}]\ ,$$

the factor that converts a present worth into a geometric series, with growth rate g. Note that if g=0, this reduces to the more traditional uniform stream conversion factor.

6.5 Load Uncertainty

6.5.1 Overview

The dynamic behavior of the peak load is governed by the Load Dynamics Model. The model provides probabilistic information of future load conditions. There is a stand-alone version of the model that is available to users. As it is integrated into the strategy model, the Load Dynamics Model provides a probability distribution on the elapsed time from installation until the installed capacity is exhausted by load growth. The output of the model is expressed as a discrete approximation, using three discrete values of time, that matches the first five moments of the actual distribution.

The model responds to the observation that load is uncertain but correlated over time; i.e., trends occur. The question is how to use this observation when developing load projections. One approach would be to project the current trend into the future. For example, suppose that area load has followed a consistent growth trend for a few years, and we believe that the trend will persist. In this situation, simple extrapolation could be used to forecast the future (extrapolation is sometimes described as being like driving by looking in the rearview mirror). We could make the extrapolation more sophisticated. For example we could guess about (1) when growth is likely to transition to a new trend and (2) the level of the new trend.

There is a fundamental problem with such extrapolations. The actual path that future load will follow is uncertain; there are a very large number of possible trajectories. Thus, picking a load trajectory or two is not going to provide a complete basis for evaluating how different investment strategies behave under uncertainty. The purpose of the Load Dynamics Model is to provide more complete load uncertainty information that can then be used for developing and evaluating investment strategies.

The Load Model is a Markov dynamic probabilistic model. Each year, a user-specified growth rate governs load growth. At the end of the year, the growth rate makes a transition that determines the growth rate that will apply over the next year. That transition is governed by a probability distribution that is also supplied by the user. By building a sequence of load growth rates, the model generates all possible load trajectories. One of the probability distributions derived from these trajectories is of particular interest. That is the first passage time distribution. The first passage time is the time required for load to grow from a given initial level to a given final level. This time is equivalent to the time to the next decision when an investment of a given capacity is made at a given load level. Thus, the time to the next decision is a random variable governed by the Load Model dynamics.

The following sections discuss the input variables that the user must provide, the output of the Load Dynamics model, and a method for assessing the input parameters required by the model. These sections are not mathematically abstract and the user is advised to consult these sections before submitting input parameters to the Strategy Model.

6.5.2 Input Variables

The Load Dynamics Model requires a relatively small number of input variables. The user must specify (1) the collection of possible load growth rates that could occur in a given year, (2) the transition probabilities between these trend states, and (3) load saturation parameters.

Load Growth Trends

The model allows from two to five load growth rates. The user must specify these trends, expressed as percentage annual growth rates.

Transition Probabilities

The dynamic behavior of peak load is governed by transition probabilities. These probabilities measure the relative likelihood of the value of load growth rate in a given year if the load growth rate in the immediately previous year were known. It is most convenient to express this in terms of a square matrix the rows of which represent the growth rate at the beginning of a year (row 1 represents growth rate 1, row 2 represents rate 2, etc.) and the columns represent the growth rate in the next year. The elements in each row are the transition probabilities from one year to the next. The diagonal elements are the probabilities of remaining in the same trend. The off-diagonal elements are the probabilities of shifting to new trends.

A procedure exists for determining these probabilities and is described below. There is also an accompanying spreadsheet implementation for the assessment procedure. The main idea in the procedure is that there are constraints and relationships that these probabilities must satisfy. Among these are the following: (i) the sum of the probabilities in any row is one; (ii) the average time that the load growth remains in a trend state is equal to 1 / (1 - p(staying in same trend); (iii) the relative likelihood of making a transition from a given load state to a different load state is equal to the ratio of the corresponding off-diagonal elements in a particular row.

Saturation Parameters

When a local area can absorb no further growth, it is said to *saturate*. We say that saturation is the maximum load that can occur in an area. As load approaches saturation, one would expect that load growth would slow. For load levels below some *saturation lower bound* there is no slowing effect. For load levels above the saturation lower bound, load grows at progressively slower rates, and approaches saturation asymptotically. The saturation model contains a so-called *gravitational constant* that determines the nature of the attenuation effect. With gravitational constant values close to 0, say 0.01, the attenuation effect increases linearly between the saturation lower bound and saturation. This kind of increase provides a uniformly decreasing approach to saturation. As values increase beyond 1, the effect becomes progressively more nonlinear, as such that the slowing effect is slight when load is just above the saturation lower bound but becomes very pronounced just before reaching saturation.

The saturation level, the saturation lower bound are provided by the user. The gravitational constant is set at 0.01 and cannot be changed by the user. Some experimentation would be appropriate to determine reasonable settings for these parameters.

6.5.3 Outputs of the Load Dynamics Model

The Load Dynamics Model finds the probability distribution on the time required for load growth to exhaust installed capacity. The results of the model computations are displayed in the output report, the Optimal Tree, that presents the optimal solution to the investment policy. At each decision point, the alternative selected is listed and three paths follow each alternative. Each path represents the time required to exhaust the additional capacity provided by the alternative. The longest path corresponds to the slowest average growth rate. The probability that each path occurs is given on the path as well. The three paths are determined by a moment-matching procedure such that the first five moments of the actual distribution are matched by the three-branch discrete approximation presented in the output report.

6.5.4 Assessing the Parameters of the Load Dynamics Model

This assessment procedure allows the user to specify the load dynamics model parameters indirectly. In this approach, assessment logic has been implemented using software that determines the model parameters that best fit the user-supplied data. The user is alerted to major data discrepancies, but is allowed to overspecify the system without worrying about minor inconsistencies.

The load assessment process has been implemented as an add-in tool to the Strategy Model.

The Procedure

The procedure for assessing the matrix in the Load Dynamics Model is comprised of five steps.

Step 1. The user enters the starting peak load, the current load growth trend, and the time interval for the forecast.

Step 2. The user describes two to five load growth scenarios that describe potential load growth outcomes for the planning period. These scenarios must be mutually exclusive and collectively exhaustive. For each scenario, the user provides the average growth rate over the forecast interval, the minimum and maximum one-year growth rates that could be experienced over the forecast interval, and the probability of occurrence of the scenario. The sum of the scenario occurrence probabilities must be one.

Step 3. The user assesses the diagonal elements of the transition matrix indirectly. For three special growth rates, the maximum and minimum one year rates (over all possible scenarios), and the current rate, the user specifies the expected holding times. The holding time is the time the growth rate is expected to persist once it occurs. The diagonal terms of the transition matrix (p_{ii}) are related to the expected holding times (t_i) by the equation $p_{ii} = 1 - 1/t_i$. The user also specifies the a measurement of the degree of volatility of the load growth process by stating whether the load growth rate is more likely to jump between adjacent growth rates or to jump to a growth rate that is farther away.

Step 4. The user is ask to estimate the maximum potential area load (saturation) and the point at which saturation starts to slow the growth rate (saturation onset).

Step 5. The assessment software determines the growth rates and transition matrix that is most consistent with all the user inputs. The transition matrix contains the one-step transition probabilities that describe the dynamics of the annual growth rates. The transition matrix that best fits the user-supplied data is found by searching a set of transition matrices. The objective of the search is to minimize a weighted sum of the squared deviations between each of the first three moments of two different end-of-period load distributions. The first distribution is calculated from the user supplied load growth scenarios. The second distribution is calculated using the load growth trends and the transition matrix

6.6 Salvageable Bandaids

6.6.1 Overview

We have defined a separate class of assets, the so-called bandaid class, as distinguished from the strategic class. Bandaids are smaller capacity alternatives, such as capacitors or other modular capacity-increasing hardware, DSM programs, and distributed resources. Bandaids are constrained to be installed in the sequence they appear in the input form. (See II.2-8.) It is natural to permit bandaids to be removed from service as larger capacity alternatives are added. This suggests that one use of bandaids is to delay implementation of the larger capacity alternatives until the need for such capacity additions becomes less uncertain. This chapter describes the approach taken in the model to implement salvaging bandaids along a trajectory.

6.6.2 Assumptions

There are several assumptions that characterize the approach. We list them below.

[A.1] All bandaids vanish at the salvage time (q.v. [A.2]). We do not permit only some of the bandaids to be salvaged. This means that all bandaids are assumed to be salvageable.

[A.2] The salvage time is given by the time at which sufficiently much additional capacity is installed. That additional capacity must be great enough to exceed the aggregate capacity of the installed bandaids. If additional capacity is installed that is less than the aggregate capacity of the installed bandaids, then, following [A.1], no bandaids are salvaged. The additional capacity required to salvage bandaids must be greater than the installed bandaid capacity by some fixed increment. That fixed increment is a planning assumption, with a default value of 10 capacity units.

[A.3] After salvaging bandaids, the time to the next decision is based on the incremental capacity added.

[A.4] The lead time effect on time to the next decision is also based on the incremental capacity added. The capital cost treatment is based on the additional capital cost.

[A.5] The economic consequence of salvaging bandaids is based on the fact that we have already included in the path cost the following: the full capital cost plus the replacement capital cost from the time the bandaids' lives end to the end of the economic horizon, $\tau = T + L_{max}$, plus the operating costs from T to τ. The unused capital cost plus the sum of the last two terms must be subtracted from the path cost to measure the effect of salvaging the bandaids.

6.6.3 Procedure

Step 6. Suppose that we are at some decision node at some time t along the planning period. A new alternative is installed with sufficient capacity to salvage the bandaids currently installed on the trajectory. The next node is the chance node that determines the time to the next decision.

Step 7. Following the assumptions above ([A.3] and [A.4]), the lead time effects and the time to the next decision are based on the net capacity added, which is the actual capacity of the new alternative less the total bandaid capacity that is salvaged.

Step 8. Since the capacity provided by the bandaids will vanish, the capital cost and operating costs induced by the new installation are given by the initial capital cost (adjusted for inflation and discounting) and the operating cost equations found in the manual, III.2-6 –III.2-18.

Step 9. The terms that must be subtracted from the path cost are the unused part of the initial capital cost of the bandaids plus the replacement capital cost and the operating costs given by the same equations noted in point 3, immediately above. Aspects of these computations are discussed below.

a. Unused Capital Cost.

Suppose that one of the bandaids that will be salvaged was installed at time $t_o < t$, and has capacity C, initial capital cost P_o, and useful life L. Thus, the bandaid would have provided capacity until calendar time $t_o + L$. The initial capital cost added to the path cost is $P_o \alpha^{to}$. The unused capital cost is the avoided future rents from t+1 through $t_o + L$. Hence the unused capital cost is

$$V_t = P_o [\alpha^t - \alpha^{L + to}]/[1- \alpha^L] \quad \text{if } \alpha \neq 1$$

$$= P_o \ (L + t_o - t) / L \qquad \text{if } \alpha = 1.$$

This number must be subtracted from the path cost.

b. Capital Cost after t_o + L.

This is the cost of replacing an investment of capacity C after t_o +L. This cost is given by

$$PW_{cap} = \sum_{k=0}^{\tau} \alpha^{to + L+1+k} \gamma^{(to + L-T)+1+k} Cx$$

$$= (\alpha\gamma)^{(to + L+1)} \gamma^{-T} Cx [1-(\alpha\gamma)^{\tau+1}] [1-\alpha\gamma]^{-1}$$

where γ is the real growth rate of the rental price, x, of future capacity. This equation is identical to that on page III.2-6 of the manual.

This number must be subtracted from the path cost.

c. Operating Costs over the interval [T+1, t_o+L].

This is part of the operating costs, based on the actual operating costs of the installed assets. The method of computation of these costs is identical to that in the manual, pp. III. 2-9 -- 2-10, Operating Costs, Regime 2: from T+1 to t+L. Since nothing is changed (other than the trivial notational change from t to t_o), the equation is not reproduced here.

The result must be subtracted from the path cost.

d. Operating Costs over the interval [t_o+L+1, τ].

These are the final parts of the operating costs. In this third regime, described in the manual on pp. III.2-11 -- 2-13, we assume that capacity is provided by the fictitious future asset. The procedure is identical to that in the manual, so it is not reproduced here.

The result must be subtracted from the path cost.

e. Modifications in the presence of saturation.

These modifications are discussed in the manual on pp. III.2-15ff. The capital costs are not affected by saturation, but the operating costs noted in parts c and d, above, may be changed. TTS(t_o) to be the time to saturation, measured from time t_o. This time will be finite if saturation is present and *if saturation occurs on an arc emanating from* t_o.

The above statement is in italics to highlight the fact that saturation has an effect on operating costs only if saturation occurs before the bandaid capacity is exhausted by load growth. If that occurs, the procedure described in the manual applies.

The result must be subtracted from the path cost.

There are no other consequences of salvaging bandaids. This concludes this section

6.7 Operating Costs

6.7.1 Overview

The purpose of this section is to discuss how operating costs are computed in the Strategy Model. Investment alternatives are used to meet peak load in a local area and are dispatched in order of increasing operating cost to meet the load at any time. Since the model is not an operations model, it seems appropriate to treat costs on an annual basis. The load served by each asset is determined by the dispatch order and the load duration curve.

The details of the dynamics of the load duration curve, the dispatching and costing algorithm, and the varying treatment of costs associated with different asset types—traditional T&D investments, load-following generation, non-load-following generation, and DSM—are discussed below.

Users not interested in these details may skip the remainder of the section without loss of ability to use the Strategy Model.

6.7.2 Load Duration Curve Dynamics

The load duration curve is provided by the user and is approximated by a broken line. The user provides information on the break points of the load duration curve in terms of annual hours for which the load is above a reported percentage of the peak load. A representation of a load duration curve is given in **Error! Reference source not found.**. The peak load in each year is the maximum height of the annual load duration curve. Each year the peak load changes, as determined by the Load Dynamics Model. Therefore, in each year, the load duration curve changes. We assume that the percentages of the break points in the load duration curve remain constant. Thus, in each year, the load duration curve is known. This load must be served by the assets present along a particular trajectory.

6.7.3 Dispatching and Costing Algorithm

The investments present along a trajectory at a particular time provide capacity above the original wires capacity in the local area. Denote the aggregate capacity at time t by C(t) and the original wires capacity by L_o. The peak load at time t is denoted L(t), and we have the relationship $C(t) > L(t)$, except when the time to the next decision occurs and the load has exhausted the capacity.

We identify three classes of investments, as follows.

If the (avoided) system energy cost is zero, then the asset is said to be *load-following*. Load following assets can provide variable amounts of capacity, from zero to the nameplate value.

If the asset is not load-following, and if both variable O&M and fixed O&M are zero, the asset is identified as *non-load-following DSM*. Any other setting of the O&M costs for a non-load-following asset identifies the asset as *non-load-following other*.

Wires and substations are considered to be load-following. Engines are typically non-load-following, although it is possible for the user to define load-following engines. The model will treat such an asset in a manner identical to a wire asset.

Note that one may also define *load-following DSM* using this data structure, by setting (avoided) system energy costs to zero; the operating costs for such a DSM program need not be zero. An example of such a program is some form of load control. It is reasonable to assume that when a customer is interrupted, some cost is incurred in order to perform the interruption.

The table below summarizes these classes of assets.

Operating Costs

		Zero	**Positive**
Avoided System Costs	**Zero**	Load Following DSM	Load Following DSM, Wires & Engines
	Positive	Non-Load-Following DSM	Non-Load-Following Engines & Other

In the following discussion the dispatch rules and cost procedures are illustrated with respect to a load duration curve. We show such curves as straight lines. We let L_o represent the original (typically wires) capacity of the local planning area. At any time $t > 0$, the peak load on the system is $L(t) > L_o$, assuming that zero load growth will not occur.

Operating Costs

We also use the following notation.

Let C_{lf} = the capacity of load-following assets installed at any time, including wires, DSM, and engines, as appropriate.

Let C_{nd} = the capacity of non-load-following DSM installed at any time.

Let C_e = the capacity of non-load-following engines installed at any time.

Let t_L = the number of hours that the load is greater than some value L.

In particular,

let t_o = the number of hours that the load is greater than the original wires capacity L_o.

Portfolios with a Single Asset

Load-Following Assets.

An example of a load-following asset is a circuit or a substation (called *wires*, loosely). Necessarily, $C_{lf} > L(t) - L_o$.

Then the asset operates for t_o and only provides capacity sufficient to meet the load between the load duration curve and L_o. The triangular area thus defined is priced at the variable O&M cost specified for the asset. SeeFigure 6.9.

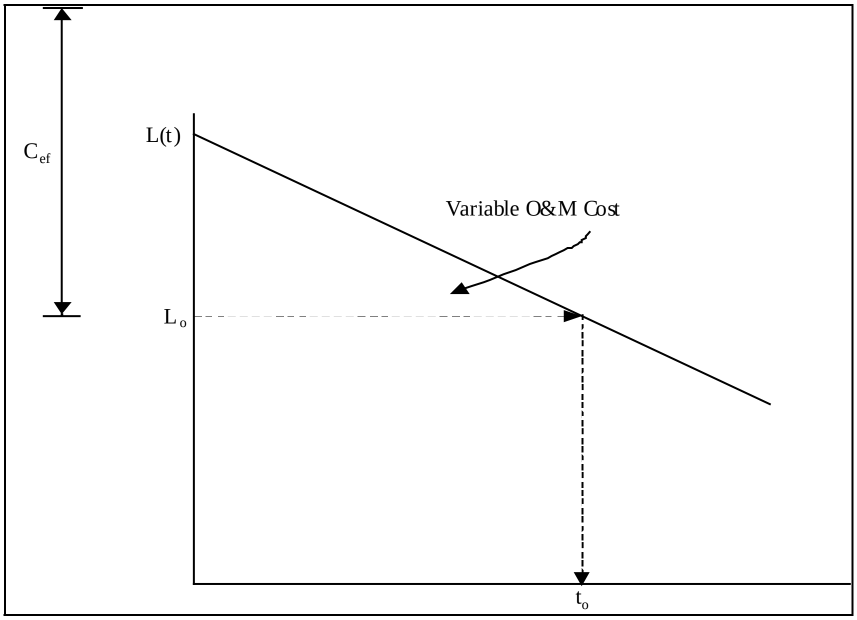

Figure 6.9 Load Following Dispatch"

All purely load-following technologies are treated in the same way. In the current version of the model, load-following technologies are defined as technologies that have zero avoided system energy costs. The operating costs can be positive or zero (see the table above). Note that this category includes load-following DSM.

Non-Load-Following Assets: DSM.

Non-load-following-DSM has zero operating costs and provides system benefits, as noted in the table, above. Such DSM programs are not dispatchable by the utility; the customer has control over their "dispatch." Again, necessarily $C_{nd} > L(t) - L_o$.

The non-load-following DSM asset operates for $t_{L(t)\text{-}Cnd}$, which is the number of hours that the load is above $L(t)\text{-}C_{nd}$. One way to think about this is that the non-load-following DSM capacity is fixed and can provide relief for loads that are below the original wires capacity. To justify this notion, consider what happens if a non-load-following DSM program were put in place but load did not grow beyond L_o. The benefits would still occur, and would be expressed as avoided energy costs for loads below L_o. Since the DSM does not follow load, the energy provided is given by a rectangle of width $t_{L(t)\text{-}Cnd}$ and height C_{nd}. The cost of providing this energy is zero. The unshaded triangular area is the load above the original wires capacity that is satisfied by the DSM. The shaded areas below and above the load duration curve are energy demands that otherwise would be met by system resources, hence they are avoided costs due to the non-load-following DSM program. Notice that as L(t) approaches L_o +C_{nd}, the time that the DSM is dispatched approaches t_o, and the area under the original wires capacity goes to zero. SeeFigure 6.10

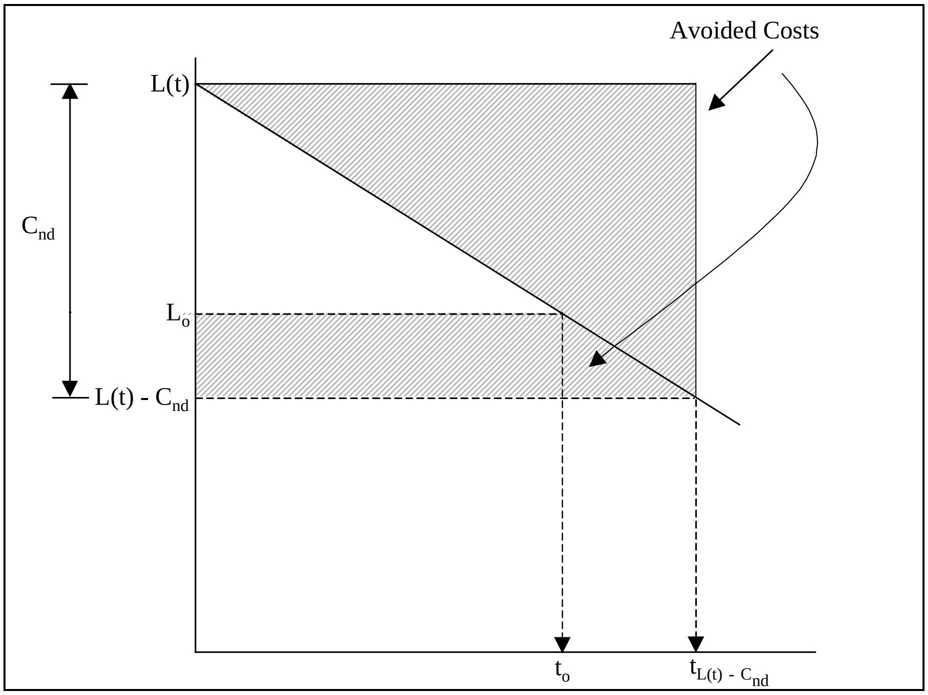

Figure 6.10 Non-Load-Following DSM Dispatch

Non-Load-Following Assets: Engines.

Non-load-following engines have operating costs that ought to be greater than the system avoided cost and provide system benefits, as indicated by the table above. That the engine is non-load-following means that when such an engine is on the full capacity is provided. Again, necessarily $C_e > L(t) - L_o$.

Non-load-following engines operate for t_o, the number of hours the load is above the original wires capacity. In each hour, the engines provide C_e, so that the energy provided is given by a rectangle of width t_o and height C_e. The cost of providing this energy is given by the operating cost of the engines. Therefore, a trapezoidal area is defined above the load duration curve that determines the energy provided by the engines that would otherwise be provided by the system. This shaded area determines the avoided costs due to the engines. Note that as L(t) approaches L_o + C_e the shaded area approaches the upper right-hand half of a rectangle. See Figure 6.11.

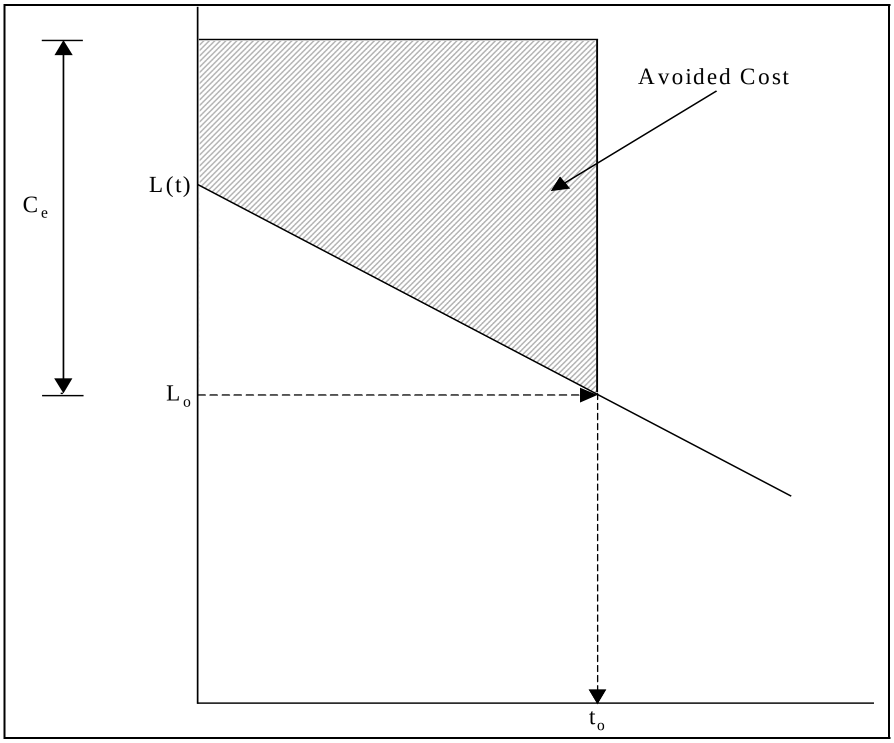

Figure 6.11 Non-Load-Following Engine Dispatch

Portfolios with Several Assets

Based on the above considerations, it is straightforward to describe the dispatch and costing of a portfolio with multiple assets. Now, necessarily, $C_e + C_{lf} + C_{nd} > L(t) - L_o$.

The concept is the following. Dispatch non-load-following DSM such that the energy benefit is maximized, subject to the asset portfolio's ability to handle the peak load. The number of hours over which non-load-following DSM is dispatched is no more than $t_{L(t) - (Ce + Clf + Cnd)}$. (See below for a more precise treatment of operating times for each asset type.) Load-following assets are dispatched next, for at most $t_{L(t) - (Ce + Clf)}$. (If there are multiple load-following assets, the model dispatches them in merit order.) Then, engines are dispatched as little as possible to meet the peak. The number of hours that engines run may be as few as zero or as many as $t_{L(t) - Ce}$. See Figure 6.12.

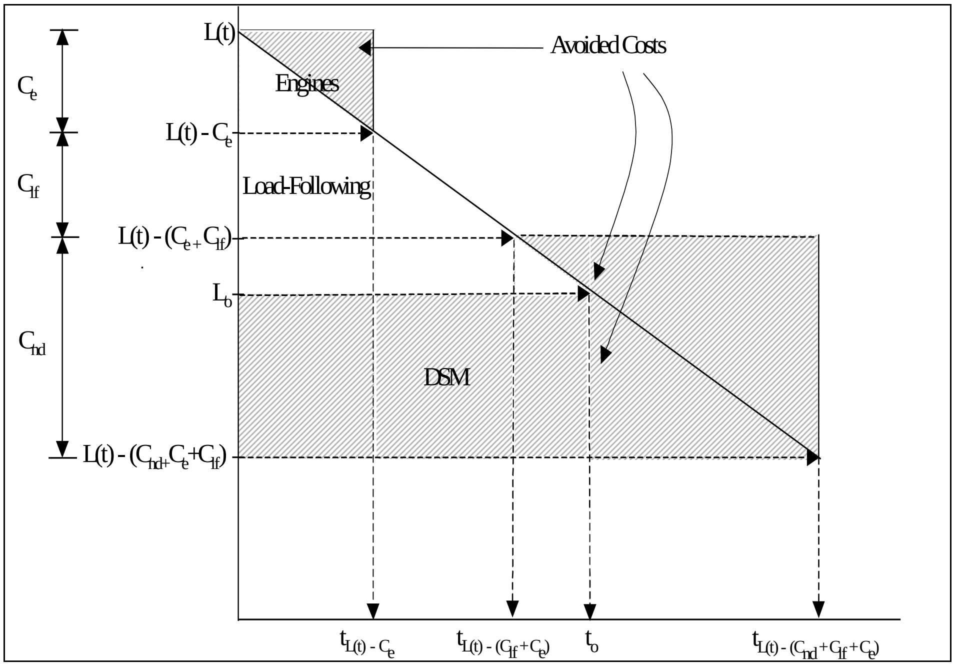

Figure 6.12 Combined Dispatch

There is one constraint that is important to recognize. We impose a limit on the number of hours that the non-load-following DSM can, in effect, operate. That number can be no more than $t_{Lo - Cnd}$. A graphically-oriented way to express this is that the DSM cannot slide down below L_o beyond its original capacity. In fact, if $L(t) - (C_e + C_{lf} + C_{nd}) < L_o - C_{nd}$, the latter bound applies and some of the engine capacity is not used to meet the peak load L(t) but is instead energy relieving, as in Figure 6.11.

We may proceed a bit more formally and specify the hours over which each asset type is dispatched. It is most convenient to suppress the time argument and focus on the load level which determines the operating time by projection through the load duration curve. For non-load-following DSM, the load at which the DSM program begins to operate is given by

$$L_{nd}^* = \max\{ L_o - C_{nd} , L(t) - (C_e + C_{lf} + C_{nd}) \}.$$

If the maximum is attained at $L_o - C_{nd}$, the constraint on non-load-following DSM dispatch is operative and there is excess capacity provided by all the assets in the current portfolio. Else, there will be exactly enough capacity to meet the peak load and the DSM has been used as much as possible.

The load level at which load-following assets are dispatched is greater by amount C_{lf} than the load at which DSM is dispatched. This is

$$L_{lf}^* = L_{nd}^* + C_{nd} = \max \{ L_o, L(t) - (C_e + C_{lf}) \}.$$

The interpretation of this result is exactly analogous to the previous result. If L(t) is sufficiently small relative to the capacities in the portfolio, then non-load-following DSM can be used as much as the constraint permits, so that the load-following technologies begin to meet load above L_o. Else, there will be exactly enough capacity to meet the peak load and both non-load-following DSM and load-following assets have been used as much as possible.

The load level at which engines are dispatched is greater by amount C_{lf} than the load at which load-following assets are dispatched, or some value less in order to meet the peak load using load-following technologies only. This is

$$L_e^* = \min \{L_{lf}^* + C_{lf}, L(t)\} = \min \{ \max \{ L_o + C_{lf}, L(t) - C_e\}, L(t) \}.$$

Consider the second expression for L_{lf}^*. If the interior maximum is achieved by the first term, then there are two possibilities. Either some engines are needed to meet the load

($L_o + C_{lf} < L(t)$), or the load can be satisfied with just non-load-following DSM and load-following technologies ($L_o + C_{lf} > L(t)$).

(Note that the condition $L_o + C_{nd} + C_{lf} > L(t)$, when $C_e > 0$, does *not* indicate that engines are never used. The dispatch rules stated here are designed to maximize the number of hours that non-load-following DSM is dispatched while using engines as little as possible to meet load. Therefore, although there would be enough capacity provided by non-load-following DSM and load-following assets to meet load L(t), the number of hours non-load-following DSM is used would be increased if some engines were used. Thus, the only way that total load can be satisfied with just non-load-following DSM and load-following assets is if $L_o + C_{lf} > L(t)$. In that case, non-load following DSM is effective for the maximum hours possible, given by $t_{L_o - C_{nd}}$, and load-following capacity C_{lf} satisfies the demand from L_o to L(t).)

If the interior maximum is achieved by the second term, then the entire engine capacity must be used to satisfy the remaining load. This is the case illustrated in Figure 6.12. The maximum hours that non-load-following DSM could be dispatched can only be achieved by using the entire capacity of the engines. Again, it may be possible to satisfy the load without any engines at all, but that would reduce the number of hours that non-load-following DSM is dispatched.

See Figure 6.13 for an illustration of these expressions.

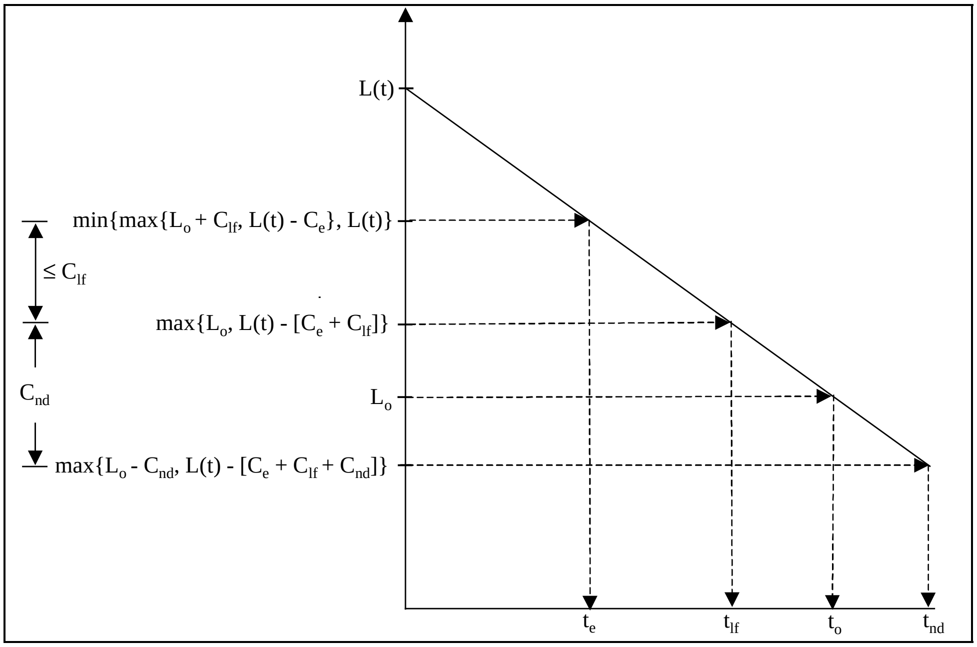

Figure 6.13 Combined Dispatch: Load Levels that Define Dispatch Times

We offer some further comments on this dispatch procedure. Note that there are avoided system costs due to both the engines and the non-load-following DSM. This dispatch is in merit order, with respect to operating cost. It is not necessarily optimal, but it appears to be consistent with the single-asset dispatches, it is simple, and it has an internal logic that captures some notion of the way non-load-following DSM and engines would actually operate. The essential feature of this dispatch is that there is a tradeoff between the energy benefit provided by the DSM and the operating cost of the engines. There are surely other ways to determine the dispatch but they have not been selected for use in the Strategy Model. In particular, we do not attempt to find the optimal dispatch.

If there are multiple non-load-following DSM programs, multiple load-following assets, and multiple engines comprising the capacities C_{nd}, C_{lf}, and C_e, the individual programs are added until (a) all non-load-following DSM is exhausted, followed by (b) load-following capacities until either load is met or load-following assets are exhausted, and then (c) adding engines until L(t) is achieved or first exceeded.

6.7.4 Examples

Consider the following example:

$L(t) = 2500$, the current peak load,

$L_o = 2000$, the original wires capacity,

$C_{nd} = 300$, comprised of two programs of sizes 100 and 200,

$C_{lf} = 150$, comprised of a single wires asset,

$C_e = 175$, comprised of two engines of sizes 50 and 125.

Apply the inequalities above.

$L_{nd}^* = \max\{ L_o - C_{nd} , L(t) - (C_e + C_{lf} + C_{nd}) \} = \max \{1700, 1875\}$, so the non-load-following DSM is dispatched over t_{1875}.

To maximize the energy savings, dispatch the larger program first, then the smaller. The dispatch hours of the smaller program is given by t_{2075}. The load level is now 1875+300=2175.

$L_{lf}^* = \max \{ L_o, L(t) - (C_e + C_{lf}) \} = \max \{2000, 2175\} = 2175$, so the wires asset is

dispatched over t_{2175}.

$L_e^* = \min \{ \max \{ L_o + C_{lf}, L(t) - C_e\}, L(t) \} = \min \{\max \{2150, 2325\}, 2500\} = 2325$,

so the engines are used to meet the peak load. A simple rule is to dispatch the engines in decreasing capacity order until the peak load is met or first exceeded. Thus, the larger engine is dispatched over t_{2325} hours and the smaller engine is dispatched over t_{2450} hours.

This completes the dispatch for this example. See Figure 6.14*a*.

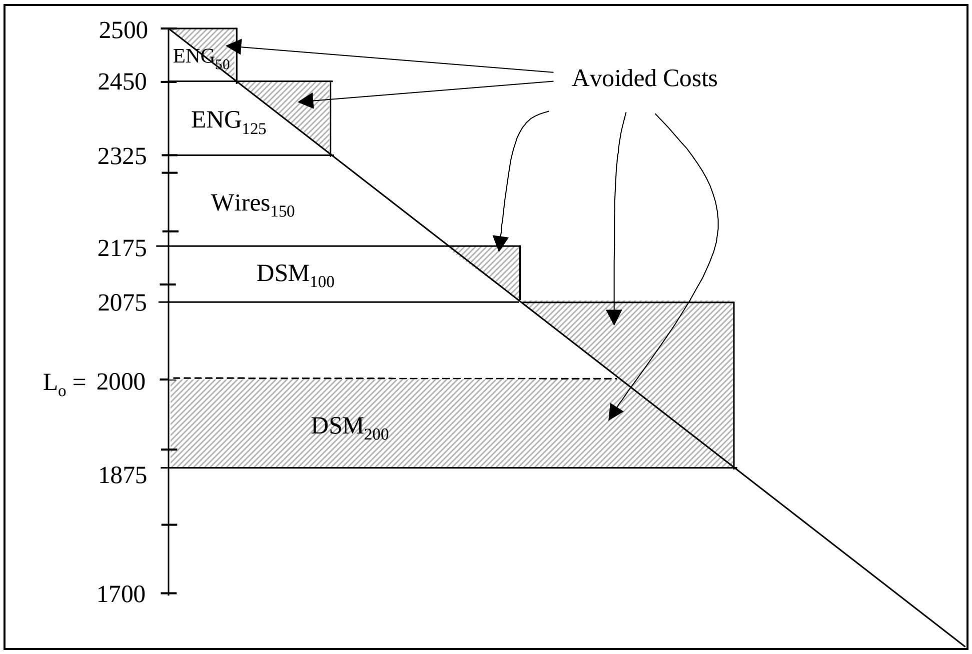

Figure 6.14a Example Dispatch

Now modify the example by adding an additional engine at size 250.

$L_{nd}^* = \max\{ L_o - C_{nd} , L(t)- (C_e + C_{lf} + C_{nd}) \} = \max \{1700, 1625\}$, so the non-load-following DSM is dispatched over t_{1700}. Dispatching the larger program first, the smaller is dispatched over t_{1900} hours, which takes the load to 2000 (L_o).

$L_{lf}^* = \max \{ L_o, L(t)- (C_e + C_{lf}) \} = \max \{2000, 1925\} = 2000$, so the wires asset is dispatched over t_{2000}.

$L_e^* = \min \{ \max \{ L_o + C_{lf}, L(t)- C_e\}, L(t) \} = \min \{\max \{2150, 2075\}, 2500\} = 2150$,

so the engines are dispatched in the sequence {250, 125, 50}, until the peak is satisfied. Thus, the largest engine (250 capacity) is dispatched over t_{2150}. This satisfies load until 2400. The next engine (125 capacity) is dispatched over t_{2400}. This exceeds the peak, so the last engine need not be dispatched.

This completes the dispatch for this example. See Figure 6.15*b*.

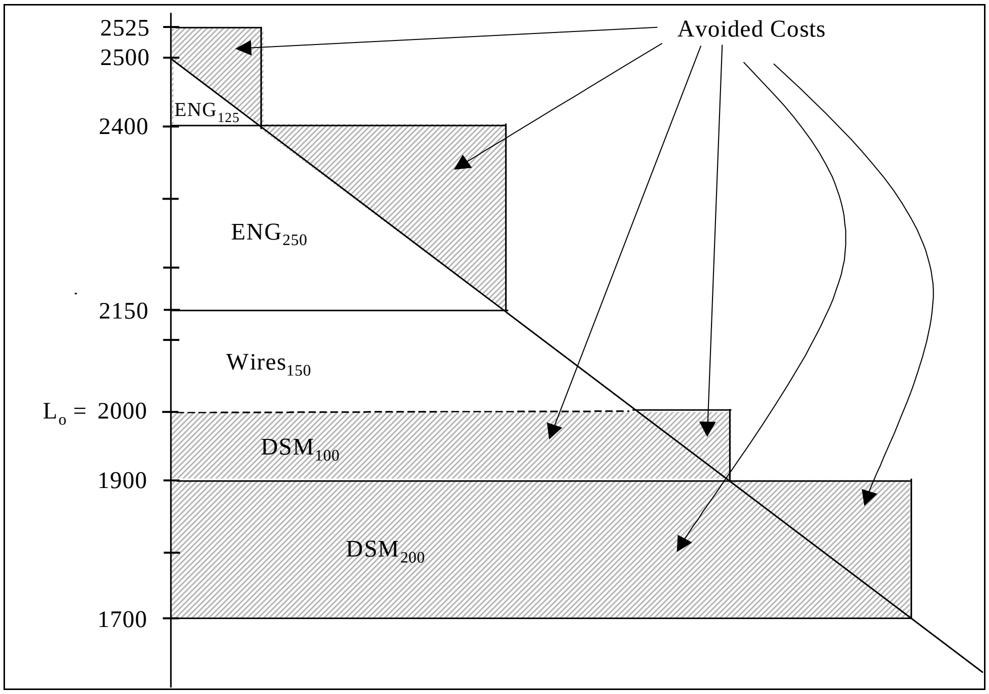

Figure 6.15b Example Dispatch

Now modify the example by letting the load grow to fill the available capacity. Thus load is at 2000 + 300 + 150 + 425 = 2875. Let the new addition, which is required, be a load-following DSM program of size 850. Now suppose the load grows to L(t)=2900.

$L_{nd}* = \max\{ L_o - C_{nd} , L(t)- (C_e + C_{lf} + C_{nd}) \} = \max \{1700, 1175\}$, so the non-load-

following DSM is dispatched over t_{1700}. Dispatching the larger program first, the smaller is dispatched over t_{1900} hours, which takes the load to 2000 (L_o).

$L_{lf}* = \max \{ L_o, L(t)- (C_e + C_{lf}) \} = \max \{2000, 1475\} = 2000$, so the load-following

assets are dispatched over t_{2000}.

$L_e* = \min \{ \max \{ L_o + C_{lf}, L(t)- C_e\}, L(t) \} = \min \{\max \{3000, 2475\}, 2900\} =$ 2900,

so the engines are never dispatched. The load is satisfied by non-load-following DSM, dispatched for the maximum possible time, and the load-following technologies. Those are dispatched in merit order. If the added load-following DSM costs less to operate than the wires asset, the former is dispatched first, over t_{2000} hours. This satisfies load up to 2850. The remaining load is met by the wires asset, but there is no avoided cost due to either the wires or the load-following DSM operation.

This completes the dispatch for this example. See Figure 6.16*c*.

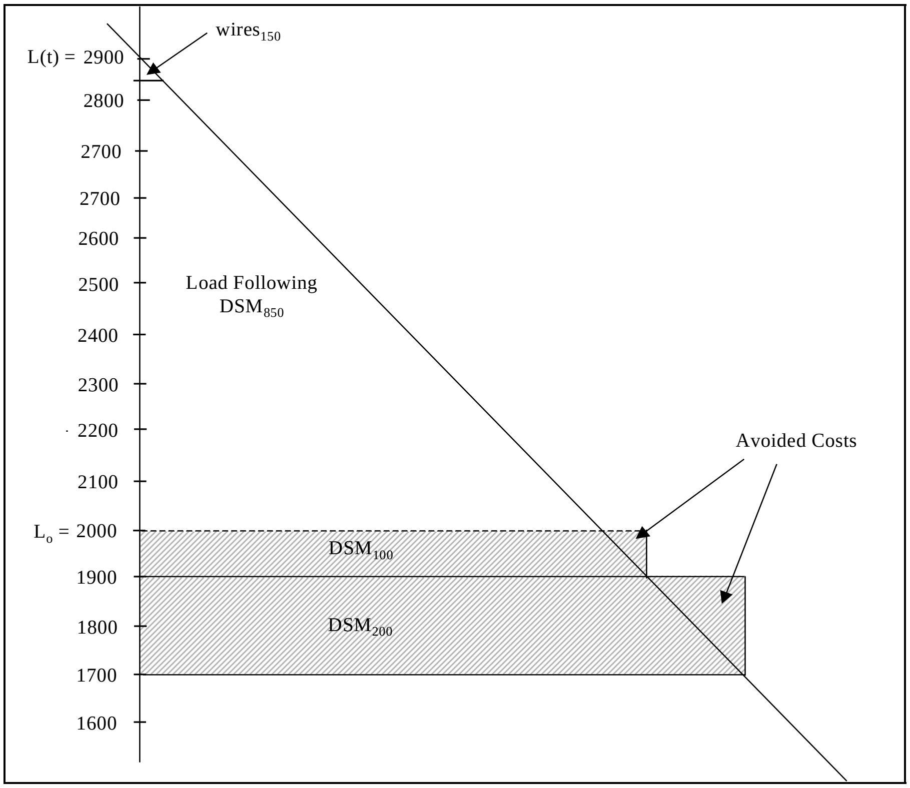

Figure 6.16c Example Dispatch

6.8 Losses and Unserved Energy Costs

6.8.1 Overview

The purpose of this section is to discuss the procedure used to measure the effects of losses and unserved energy. As load increases, additional capacity investment is required in order for demand to be satisfied. Nevertheless, even if there is sufficient capacity to meet the demand, events may occur that prevent the system as it is configured to satisfy all customers. *Unserved energy* is the term used to describe that demand that is not met for any reason at any time. It is natural to consider unserved energy as an annual value. Further, the physical nature of the apparatus in the distribution system entails *losses*, which are expressible as an annual cost. These losses are a function of the impedance of circuit elements and the load on the system. The Strategy Model includes a procedure for measuring the costs associated with unserved energy and losses. The procedure can be thought of as a sequence of steps, as described below.

The basis of the procedure is a collection of data the user supplies. The data exhibit the losses and unserved energy that are associated with each investment alternative as load on the system grows. Once these data are provided, the Strategy Model can compute losses and unserved energy costs for any trajectory. The description of the procedure is not mathematically abstract and the user is advised to consult these sections before submitting input parameters to the Strategy Model.

6.8.2 Strategic Alternatives.

As described in CHAPTER 1 of this Guide, the user defines a set of *main strategic alternatives*. For example, this set could consist of the elements {feeder, substation}, with one feeder and one substation as possible expansion alternatives. It would be possible to have more than one feeder or substation or both. The elements of the set and their multiplicities are arbitrary. The assumption made is that it is possible to characterize a collection of pure strategies. *Bandaids* comprise an additional class of expansion alternatives and are treated differently than the strategic expansion alternatives. We discuss bandaids in more detail below.

6.8.3 Loss and Unserved Energy Curves.

The user provides a set of curves that describe Losses and Unserved Energy costs as a function of load for each of the main strategic alternatives. That is, if the only investments made were a single strategic alternative, repeated sufficiently to respond to load growth, then the losses and unserved energy costs as a function of load are specified. These functions characterize the pure strategies only. We suppose that these curves or tables can be gotten by users in a straightforward manner, perhaps by using their favorite load flow software or similar analysis tools. It is also possible to consider the strategy of doing nothing. This pure strategy also entails unserved energy and losses. The user is also asked to provide a set of curves for this “strategy.” These curves are used to describe the effects of installing bandaids only.

6.8.4 Bandaids.

When bandaids are present, the user describes them with respect to effective capacity, capital cost and operating cost. The user may specify a single value for effective capacity of each bandaid, or a sequence of such capacities. The sequence reflects the order in which the bandaids are actually installed. For example, the user might specify that all bandaids are 300 kW. Alternatively, the user might specify the sequence {300, 500, 450, 1200,...}. Any specification is acceptable.

If a sequence is specified, the model assumes that the installation rules are such that if a bandaid decision is made, the earliest feasible bandaid in the sequence is the one that is installed. Feasibility is defined by prior installations and constraints. Thus, bandaids are a separate class of alternatives, distinct from the strategic expansion alternatives. The constraints that apply to strategic expansion alternatives must be augmented by this precedence or ordering constraint for bandaids.

6.8.5 Cumulative Loss Reduction (%) and Unserved Energy Reduction (%).

Relative to the curves or tables supplied for the pure strategies above, each incremental bandaid installation can reduce the cost of losses and unserved energy. The user may specify a percentage reduction provided by each bandaid. The cumulative is perhaps the easiest to provide. The model permits the user to supply a single percentage for all bandaids, regardless of what else might be installed, or a single percentage for all bandaids that depends on what else is installed, or variable percentages depending on what else is installed. The sample table below illustrates the effect of bandaids. The "Pre" case occurs before either a feeder or substation is installed, where the strategic alternatives are {feeder, substation}. Hence, that case corresponds to doing nothing. The reduction percentages indicate the effect of installing a sequence of bandaids only to respond to load growth. The "Feeder" case occurs after a feeder is installed. The substation (Substn) case occurs after a substation is installed. Thus, the effect of bandaids is given with respect to pure strategies. The interaction rules for mixed strategies are given below.

	Cumulative Loss Reduction %			Unserved Energy Reduction %		
	Pre	Feeder	Substn	Pre	Feeder	Substn
Bandaid 1	.03	.02	0	.04	.03	0
Bandaid 2	.05	.04	0	.06	.04	0
Bandaid 3	.07	.05	0	.08	.05	0
Bandaid 4	.08	.05	0	.09	.05	0
Bandaid 5	.08	.05	0	.10	.05	0

Table 6.1 Effects of Bandaids on Losses and Unserved Energy

6.8.6 Interaction Among Major Strategic Alternatives.

Trajectories need not be restricted to pure strategies, and it is necessary to characterize the joint effect of several strategic alternatives. The Strategy Model computes the interaction effects using some simplifying assumptions.

Losses

We assume that losses are determined by both the total load and the assets installed. We assume that the actual amount of losses can be best approximated by the minimum value of the losses given by the curves provided by the user, where the minimum is taken over the strategic expansion alternatives actually installed.

Unserved Energy

We assume that the unserved energy is determined by the incremental load on the assets installed. We also assume that the total load can be allocated to the strategic expansion alternatives actually installed based on the capacity of those alternatives. Total unserved energy is additive and the total cost is found by superposition.

6.8.7 Procedure: Sample Model Computations

The purpose of these sample descriptions is to indicate how the model computes the losses and unserved energy along a trajectory based on the inputs provided by the user. Both pure strategies and mixed strategies can be very easily evaluated.

Pure Strategies.

1. Install Bandaid 1 at L_o. As load grows in each year, the losses and unserved energy incurred are given by the curve for the Do-Nothing case, modified by the appropriate percentages in the “Pre” columns of the table above. As more bandaids are added, the model adjusts the reduction percentages accordingly.

2. Install Feeder or Substation at L_o. As load grows, the losses and unserved energy costs are given by the curve for the appropriate strategic expansion alternative. If there are multiple feeders or multiple substations, the procedures for losses and unserved energy are different.

   a. Losses. As load grows, the model uses the value of losses given in the table, regardless of the number of feeders or substations. If the load grows beyond the table domain, the model extrapolates the value of losses linearly.

   b. Unserved Energy. As load grows, the model adds the unserved energy that corresponds to the capacity of the strategic expansion alternative, per unit. When load is less than total capacity, the unserved energy for the last unit installed is given by the table for the incremental value of load on that last unit.

Mixed Strategies.

1. Let an installation be made at some time t, at some load L(t), t<T, augmenting some collection of previously installed assets.

2. The losses are given by the minimum value of the losses corresponding to the load level, where the minimum is taken over the installed strategic expansion alternatives. The model applies the minimum bandaid reduction percentage as appropriate to the values in the curves.

3. The unserved energy is given by a procedure analogous to 1.2.b., above. The model adds the unserved energy incrementally, as load is allocated to the installed strategic expansion alternative units, so that the minimum value of unserved energy is added. If there are bandaids present, the total load may be more than the aggregate capacity of the installed strategic expansion alternatives. If so, the model continues along the unserved energy curve beyond the unit capacity for the unit that would provide the minimum unserved energy. The model applies the bandaid reduction percentage as appropriate to the values in the curves.