




analytica[®]
Beyond the Spreadsheet

Optimizer Guide

Analytica 4.1

January 15, 2008



Lumina Decision Systems, Inc.
26010 Highland Way
Los Gatos, CA 95033
Phone: (650) 212-1212
Fax: (650) 240-2230
www.lumina.com



Copyright Notice

Information in this document is subject to change without notice and does not represent a commitment on the part of Lumina Decision Systems, Inc. The software program described in this document is provided under a license agreement. The software may be used or copied, and registration numbers transferred, only in accordance with the terms of the agreement. It is against the law to copy the software on any medium except as specifically allowed in the license agreement. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the licensee's personal use, without the express written consent of Lumina Decision Systems, Inc.

This document is © 1993-2008 Lumina Decision Systems, Inc. All rights reserved.

The software program described in this document, Analytica, is copyrighted:

© 1982-1991 Carnegie Mellon University

© 1992-2008 Lumina Decision Systems, Inc., all rights reserved.

Analytica was written using MacApp®: © 1985-1996 Apple Computer, Inc.

Analytica incorporates Mac2Win technology, © 1997 Altura Software, Inc.

The Analytica® software contains software technology licensed from Carnegie Mellon University exclusively to Lumina Decision Systems, Inc., and includes software proprietary to Lumina Decision Systems, Inc. The MacApp software is proprietary to Apple Computer, Inc. The Mac2Win technology is technology to Altura, Inc. Both MacApp and Mac2Win are licensed to Lumina Decision Systems only for use in combination with the Analytica program. Neither Lumina nor its Licensors, Carnegie Mellon University, Apple Computer, Inc., and Altura Software, Inc., make any warranties whatsoever, either express or implied, regarding the Analytica product, including warranties with respect to its merchantability or its fitness for any particular purpose.

Lumina Decision Systems is a trademark and Analytica is a registered trademark of Lumina Decision Systems, Inc.

Analytica Optimizer incorporates SolverSDK.dll from Frontline systems, Inc.: Software copyright © 1991-1999 by Frontline Systems, Inc.

Portions copyright © 1989 by Optimal Methods, Inc.

Portions copyright © 1994 by Software Engines.

About Analytica Optimizer	1
Introducing the Analytica Optimizer	2
What do I need to know?	2
What is the Analytica Optimizer?	2
How do I obtain the Analytica Optimizer?	3
Activating the Optimizer for Analytica	3
Activating Analytica Optimizer for ADE	4
Installing Optimizer add-on engines	5
What's new in Analytica Optimizer 4.0	5
Chapter 1: Quick Start	7
Quick start	8
Browsing Analytica Optimizer functions	8
Linear programs	8
A nonlinear program	14
Chapter 2: Formulating an Optimization Problem	19
Parts of an optimization problem	20
Continuous, integer, and mixed-integer programs	20
Choosing the type of optimization	21
Solving simultaneous equations	21
Chapter 3: Linear Optimization	23
Defining a linear optimization problem	24
Optional parameters	25
Obtaining the solution	25
Secondary aspects to a solution	27
Examples	30
Integer, binary, and grouped decision variables	30
Controlling the search	31
Linear programming settings	33
Array abstraction	34
Chapter 4: Quadratic Optimization	35
Defining a quadratic program	36
Solution properties	37
Common quadratic situations	38
Obtaining the solution	39
Search control settings	39
Examples	39
Chapter 5: Nonlinear Optimization	41
Nonlinear programs	42
Problem formulation	42
Obtaining the solution	44
Optional parameters for NLP	44
Array abstraction	45

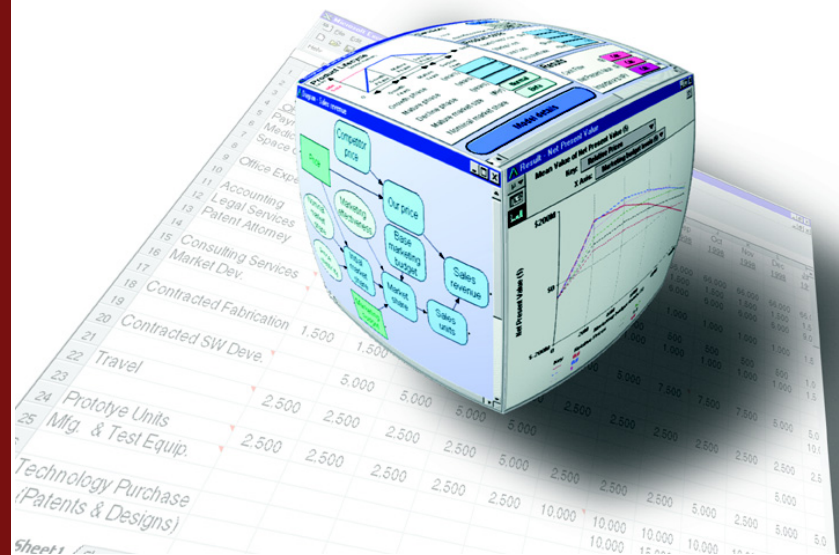
Integer, binary, and mixed-integer programs	46
Airline example for NLP	47
Intelligent arrays, array abstraction, and NLP	51
Summary of array abstraction for NLP	59
Solving systems of equations	60
Other examples	60
Giving hints to help the Optimizer	61
Dealing with local optima	62
Chapter 6: Control Settings	65
Controlling the search	66
Selecting the optimization engine	66
Examining engine capabilities	67
Specifying settings	68
Examining available settings	69
Termination controls	69
Algorithm selection	71
Numeric tolerance and precision	81
Chapter 7: Debugging Optimization Problems	83
Writing and reading from a file	84
Diagnosing conflicting constraints	84
Debugging a nonlinear optimization	85
Chapter 8: Optimizer Function Reference	87
Problem definition functions	88
Other functions	91
Chapter 9: Logistic Regression Function Reference	95
Logistic regression functions	96
Index	99

Introduction

About *Analytica Optimizer*

This introduction explains:

- What the Analytica Optimizer is
- How to obtain the Analytica Optimizer
- How to activate the Analytica Optimizer
- How to activate the Analytica Optimizer for ADE
- How to activate add-on engines



Introducing the Analytica Optimizer

This *Optimizer Guide* explains how to use the Analytica Optimizer. The Chapter 1, “Quick Start,” is a tutorial that takes you through the key steps to create some simple example Analytica models that use linear and nonlinear optimization. Chapter 2, “Formulating an Optimization Problem,” helps you to formulate your model for optimizing, and to choose whether it requires linear programming (LP), quadratic programming (QP), or nonlinear programming (NLP). The other chapters provide more details on these three types of optimization and their options and settings, and provide an overview of creating and debugging optimization problems. The last two chapters provide a concise reference for all the optimization functions.

What do I need to know?

This guide, including the *Quick Start* chapter, assumes you have basic knowledge of building models and writing expressions in Analytica. If you do not, you should first work through the *Analytica Tutorial* and scan through the *Analytica User Guide*.

This guide provides an introduction to the basic concepts of optimization, including linear, quadratic, and nonlinear programming. It is not, however, a complete textbook on optimization. You might find it useful, especially for more challenging applications, to consult one of the many good textbooks on optimization.

What is the Analytica Optimizer?

The Analytica Optimizer enhances Analytica with powerful functions to find optimal decisions and to solve equations. An optimal decision strategy can maximize value or minimize costs of any quantified objective. The optimization might be subject to a set of constraints. Analytica Optimizer offers linear programming (LP), quadratic programming (QP), and nonlinear programming (NLP). LP requires linear objective functions and linear constraints. QP requires quadratic objective functions and linear or convex quadratic constraints. NLP handles general nonlinear objective and constraint functions. All three methods handle decision variables that are continuous, discrete (integer, Boolean, grouped), or mixed.

The Analytica Optimizer uses the Premium Solver Platform licensed from Frontline Systems, Inc. Frontline developed the Optimizer/solvers in Microsoft Excel and other spreadsheets, and is the world leader in spreadsheet optimization. Their Premium Solver is the leading add-on software for spreadsheet optimization, and incorporates state-of-the-art technologies. The LP and QP methods handle up to 8000 variables and 8000 constraints in addition to variable bounds (up to 2000 variables can be integers, and the limit is 2000 variables when quadratic constraints are present). The NLP methods offer hybrid methods using classical gradient-search and evolutionary (genetic) algorithms for smooth and discontinuous objective functions, with up to 500 decision variables and 250 constraints.

The Analytica Optimizer performs optimization under uncertainty to maximize expected values and minimize loss percentiles, as well as other statistical functions of objectives and constraints. The LP and QP methods fully support Analytica's Intelligent Arrays. Thus, you can easily create arrays of optimizations conditioned on samples from uncertain variables, for parametric analysis of effects of key assumptions, and for each time period in a dynamic model. The NLP functions do not fully support Intelligent Arrays. But, you can optimize nonlinear objectives that aggregate over dimensions, e.g.,

expected net present value to aggregate over uncertainty and time, and you can manually configure an optimization problem to abstract over explicitly named dimensions.

The Analytica Optimizer is an edition of Analytica that includes all the functionality of the Analytica Enterprise edition. After developing Optimizer-based models with Enterprise, you can deliver them to end users on the desktop using Analytica Power Player with Optimizer, or via a web browser on a server using the Analytica Decision Engine (ADE) with an Optimizer license.

How do I obtain the Analytica Optimizer?

You can purchase a license for the Analytica Optimizer or the Analytica Power Player with Optimizer from Lumina Decision Systems. Or you can purchase an upgrade to Optimizer if you already have a license for Enterprise or Professional editions.

If your copy of Analytica is for release 3.1 or earlier, you need to upgrade to release 4.1 to obtain the newest Optimizer features as described in this manual. Substantial discounts are available if you have a maintenance agreement for Analytica 3.1 (included free for 12 months from purchase).

Use of Optimizer from ADE requires a special ADE license. ADE Optimizer licenses include limits on the number of concurrent ADE process instances, with licensing pricing based on the maximum number of concurrent instances allowed on a machine.

For more information:

- Visit the Lumina web site: <http://www.lumina.com>
- Call Lumina at 650-212-1212

Activating the Optimizer for Analytica

If you already have any edition of Analytica 4.1 installed, your installation includes the Optimizer files and there is no need to download new software. To activate the Optimizer software, you need to enter a new license code with the Optimizer option. Follow these steps:

1. Start Analytica in the usual way, e.g., via the Windows Start menu, or by double-clicking an Analytica model file.
2. From Analytica's **Help** menu, select the **Update license** option, to show the **Analytica Licensing Information** dialog box.
3. Replace the existing license code at the bottom of the dialog box with a new code that activates the Optimizer. If you have received the new license code in an e-mail, you can copy and paste it directly into the dialog box.
4. Click **OK**.
5. Exit and restart Analytica.

You can verify successful activation of Analytica Optimizer by examining the splash screen when Analytica starts up, or by going to **Help > About Analytica**. The splash screen should display "Analytica Optimizer," as shown below.



Activating Analytica Optimizer for ADE

The Analytica Decision Engine (ADE) Optimizer edition is also available in release 4.1. This kit includes a license code for Analytica Enterprise for developing models, as well as a code ADE for the production server. Similarly, ADE Optimizer includes a license code for Analytica Optimizer as well as a code for ADE Optimizer. See “Activating the Optimizer for Analytica” on page 3 for instructions on activating Analytica Enterprise with Optimizer. ADE Optimizer licenses include a maximum limit on the maximum number of concurrent ADE process instances that can be running concurrently on the same computer.

If you currently use ADE release 3.1 or earlier, you first need to upgrade it to release 4.1.

To upgrade an existing non-Optimizer installation of ADE 4.1 to activate the Optimizer, follow these steps to enter a new license code for the ADE Optimizer:

1. Open a command prompt. From the Start menu, select **Run**, type `cmd` and press **OK**.
2. Type `cd ADE_Dir`, where `ADE_Dir` is the path to the directory for ADE 4.1. On most computers this is `cd c:\Program Files\Lumina\ADE 4.1`
3. Type `ade.exe /RegServer`
A dialog appears with a text box to enter your new license code.
4. Enter the new license code and press **OK**.

Installing Optimizer add-on engines

With Analytica Optimizer 4.0, it is possible to add on other engines for solving optimization problems. Some engines provide superior performance on particular classes of optimization problems, and some engines handle larger numbers of variables or constraints. Among the available add-on engines are MOSEK, large-scale SQP, XPRESS, KNITRO, OptQuest, large-scale GRG, and large-scale LP/QP. These add-on engines are available at additional cost and require special installation steps.

To install an add-on engine in Analytica Optimizer, you must first receive the special license code and the DLL file containing the engine. With Analytica Optimizer already installed, place the DLL in your Analytica install directory (or, alternatively, remember the full file path to the DLL). Use **Regedit** to add the following registry key/folder (if it does not already exist):

```
HKEY_LOCAL_MACHINE/Software/Lumina Decision Systems/Analytica/4.0/  
SolverEngines
```

In this folder, create a string value using the name of the add-on engine, and set its value to the full file path of your DLL. For example:

```
KNITRO : C:\Program Files\Lumina\Analytica 4.0\Knitro.dll
```

Next, find the `solver.lic` file in your Analytica install directory and open it in a text editor such as Notepad. Add the license key provided for the add-on engine to this file.

To test for proper installation, open Analytica, and create a variable defined as follows:

```
Variable Engines := SolverInfo("AvailEngines")
```

Show the result for this variable and verify that your engine name appears in the resulting list.

What's new in Analytica Optimizer 4.0

Analytica Optimizer has several new features and has become easier to use since Analytica 3.1. Enhancements include:

- It has been upgraded to Frontline's SDK version 7.1 (from 4.5).
- There is the capability to add in other external solvers engines (at an additional cost), including XPRESS, KNITRO, MOSEK, OptQuest, and Frontline large-scale engines.
- All Optimizer algorithm/engine settings are specified through two parameters, named **parameter** and **setting**, to **LpDefine()**, **QpDefine()**, and **NlpDefine()**. This scheme generalizes to the use of different engines, including new add-on engines.
- A new type of integer constraint, group integers, is supported. In this integer type, decision variables belonging to the same group are prohibited from having the same value.
- Quadratic constraints are allowed in quadratic programs defined using **QpDefine()**. Special algorithms are highly effective in solving convex quadratic constraints and second-order cone programs.
- A new function, **SolverInfo()**, provides access to the list of installed Optimizer engines, the components of an optimization problem definition, and the attributes of an Optimizer engine.


- When specifying nonlinear optimization problems, almost all parameters to **NlpDefine()** are now optional, so it's easy to specify simple optimization problems. For example, the **Vars** and **Constraints** indexes can be omitted when there is only a single scalar decision variable, or zero or one constraints. For example, a simple unconstrained scalar optimization requires only two parameters.
- New **SetContext** and **Over** parameters to **NlpDefine()** provide more flexibility for structuring your model so that your optimization array abstracts properly. Also, **NlpDefine()** can be specified within a dynamic loop, so that the definition of an optimization problem at time t is based on the optimal solution of a previous optimization problem at time t-1.
- The new **traceFile** parameter to **NlpDefine()** makes it easy to log the points visited during an optimization search to a file for debugging.
- The **ObjNL** and **LhsNI** parameters to **NlpDefine()** allow quadratic dependence to be specified as a further hint.
- The **MultiStart** setting can be used with the GRG Nonlinear engine in **NlpDefine()**, which is often quite effective when local minima are present.
- The nonlinear engines, such as GRG Nonlinear or Evolutionary, can optionally be used to solve problems defined using **LpDefine()** or **QpDefine()**, as an alternative to the default LP/Quadratic and SOCP Barrier engines. With **QpDefine()**, this might be necessary if the quadratic constraints are non-convex.
- **LpWrite()**, **LpRead()**, and **LpWriteIIS()** now support the three file formats LP, MPS, and LPFML. These formats are standards used for exchange of linear programs between other Optimizer software products.
- The set of status codes returned from **LpStatusNum()** and the set of status messages returned from **LpStatusText()** have changed. Legacy models that tested against specific status numbers might need to be adjusted.

Chapter 1

Quick Start

This chapter shows you:

- How to browse Analytica functions
- How to obtain the Analytica Optimizer
- How to optimize a linear program
- How to optimize a nonlinear program



Quick start

Who this is for

This section leads you through a series of steps to create Analytica models that solve some simple linear and nonlinear optimization problems. The reader should follow along by performing the steps in Analytica.

If you're familiar with LP, QP, and NLP

If you are already familiar with concepts of linear, quadratic, and nonlinear programming, this provides a fast way to get started creating Analytica optimization models. Since this chapter does not cover all the functions and features of Analytica Optimizer, and their use in complex situations, you should review the rest of the manual as well, especially “Problem definition functions” on page 88.

If you're not an expert

If you do not have a background in linear and nonlinear programming, performing the step-by-step instructions in this section might be the best place to start, even if you don't yet understand why each step is being done. Afterwards, read the rest of this manual, returning to the examples in this section as you learn more about Analytica Optimizer. Also, be sure to explore the Optimizer example models included with Analytica Optimizer.

Analytica prerequisites

This manual assumes a basic knowledge of modeling and writing expressions in Analytica. If you do not have this background, you should go through the *Analytica Tutorial* and *Analytica User Guide* prior to continuing with this manual.

Browsing Analytica Optimizer functions

To begin, follow these steps:

1. Start Analytica using the menu commands **Start > Programs > Analytica 4.1 > Analytica 4.1**.
2. In the main application menu, select **Definition**.
3. Move your cursor down to the **Optimizer** submenu.

On the submenu that pops up, take a minute to scan the Analytica Optimizer function names. If you do not have an **Optimizer** option on your **Definitions** menu, it means that you do not have a license code to activate Analytica Optimizer. You need to contact Lumina at sales@lumina.com.

4. Select the diagram window and press *Control+2* to create a new variable, and *Control+e* to edit its definition.
5. Select **Paste Identifier** on the **Definition** menu.
6. Using the library pull-down menu, select **Optimizer**.

From here you can review the Optimizer functions along with parameters and function descriptions. The two main functions to study initially are **LpDefine()** (to define a linear program) and **LpSolution()** (to solve a linear program).

Linear programs

This section shows you the process of encoding a linear program in Analytica Optimizer. The model you create here is included in the **Example Models/Optimizer**

Examples directory installed with Analytica under the name **Two Mines.ana**. This is the problem you will encode:

The Two Mines Company owns two different mines that produce an ore that, after being crushed, is graded into three classes: high, medium, and low-grade. The company has contracted to provide a smelting plant with 12 tons of high-grade, 8 tons of medium-grade, and 24 tons of low-grade ore per week. The two mines have different operating characteristics as detailed below.

How many days per week should each mine be operated to fulfill the smelting plant contract?¹

Mine	Cost per Day (\$1000)	Production (tons/day)		
		High Grade	Medium Grade	Low Grade
X	180	6	3	4
Y	160	1	1	6

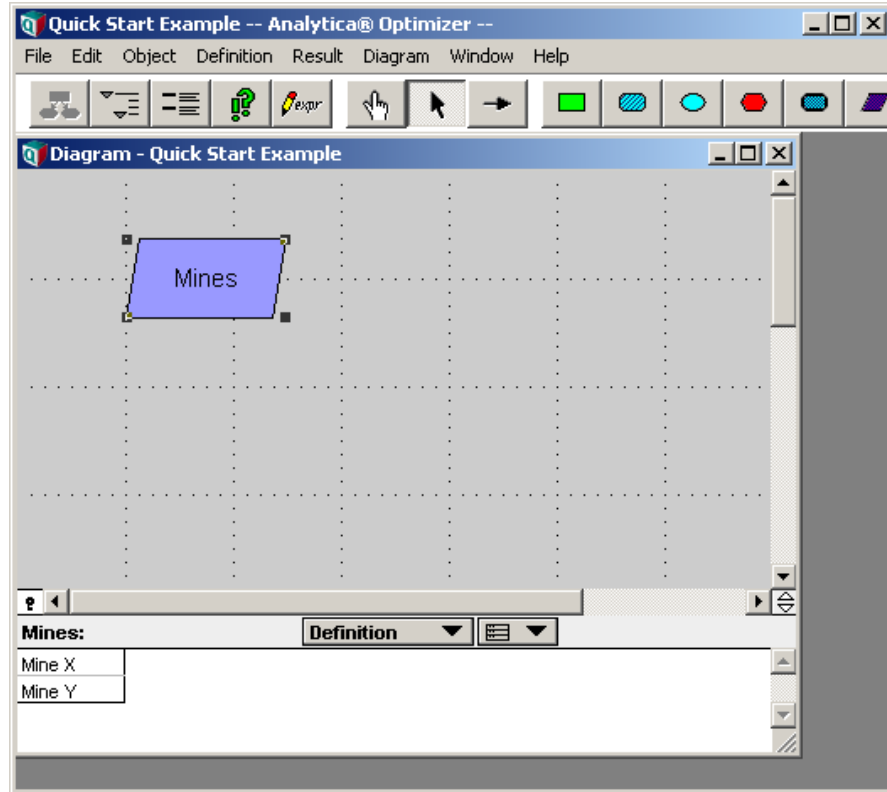
The first step is to identify the decision variables, in this case the number of days per week to operate each mine, and then create an index variable naming each decision variable:

1. Create an index name and name it **Mines**.

We will use this as the index for the objective variables, i.e., the number of days per week to operate each mine.

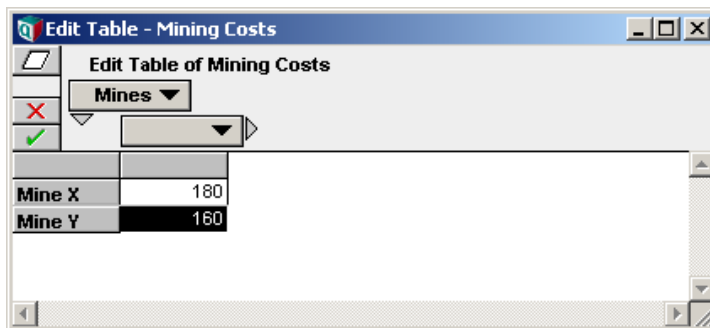
2. Edit its definition attribute and set **List Of Labels** as its definition from the pull-down menu
3. Enter the labels *Mine X* and *Mine Y*.

1. This example was created by J.E. Beasley.
Cf. <http://www.brunel.ac.uk/depts/ma/research/jeb/or/contents.html>




Next, enter the mining costs, which will become the objective coefficients that define the objective as a linear function of the decision variables:

4. Create a variable and name it **Mining_costs**. Set its units attribute to **\$/day**.
5. Edit its definition attribute and set the definition type to **Table**. In the index chooser, select **Mines** and press **OK**. Populate the table with the operating costs as follows.

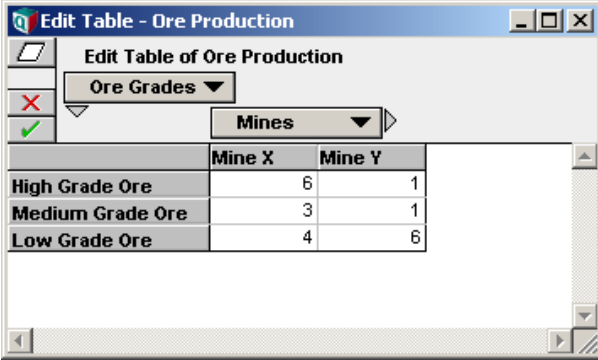


In this problem, there is one production constraint for each grade of ore. Thus, an index for ore grade can serve as the constraint index.

6. Create an index variable and name it **ore_grades**. Set its definition to a list of labels.

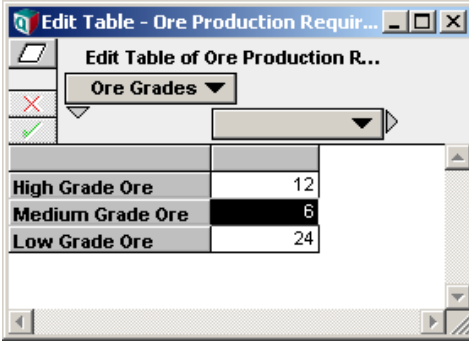


7. Create a decision variable and name it `ore_Production`. Set its units to `tons/day`. Set its definition type to **Table** and in the index chooser select both `Mines` and `Ore_Grades`. Fill in the table.



	Mine X	Mine Y
High Grade Ore	6	1
Medium Grade Ore	3	1
Low Grade Ore	4	6

8. Create a variable and title it "Ore Production Requirements." For convenience, set its identifier to `ore_prod_req`. Set its units to `tons/week`.
9. Edit the definition attribute for **Ore Production Requirements** and set the definition type to **Table**, selecting `ore_grades` as its index. Fill in the edit table.



High Grade Ore	12
Medium Grade Ore	6
Low Grade Ore	24

Note that the constraints for the problem for each ore grade are

$$\text{Sum}(\text{Ore_production} * \mathbf{x}, \text{Mines}) \geq \text{Ore_prod_req}$$

where \mathbf{x} is the objective, i.e., the number of days per week to operate each mine.

We now have all the inputs required to define the linear program.

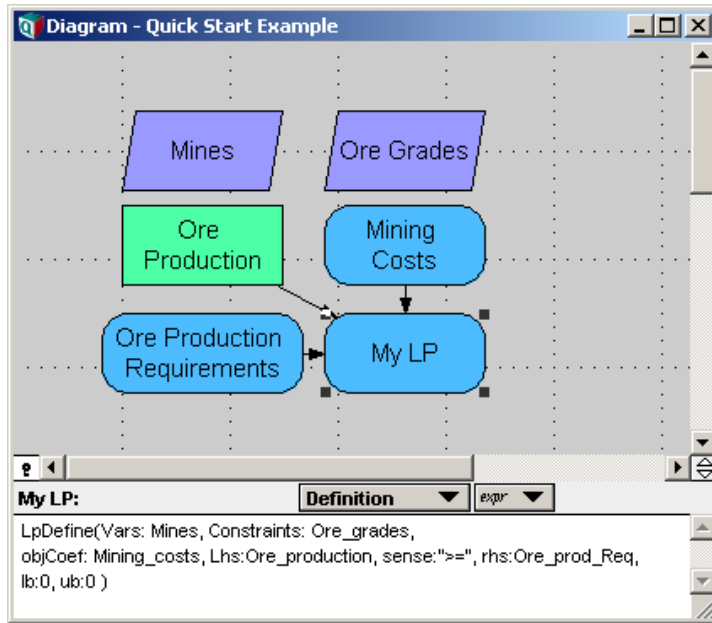
To create the linear program to solve this problem:

10. Create a variable and name it `my_LP`. Enter the following definition.

```
LpDefine(Vars: Mines,
constraints: Ore_grades,
objCoef: Mining_costs,
lhs: Ore_production,
sense: ">",
rhs: Ore_prod_req,
lb: 0,
ub: 5)
```


The parameters **lhs**, **sense**, and **rhs** refer to the left-hand side of the constraint equations, the constraint equation comparator (greater than, equals, less than), and the right-hand side of the constraint equations, respectively. The last two parameters, **lb**

and **ub** (the lower and upper bounds), specify the limits on the number of days per work week that a mine can operate.



Note that the example above uses **name-based calling syntax** for the function **LpDefine()**. In this syntax, you give each parameter by name, colon, and the expression to be passed, e.g., **Vars: Mines**. You can also use more conventional position-based syntax, but that is less comprehensible for functions like **LpDefine()** with many parameters and options (see “Name-based calling” in Chapter 20 of the *Analytica User Guide*).

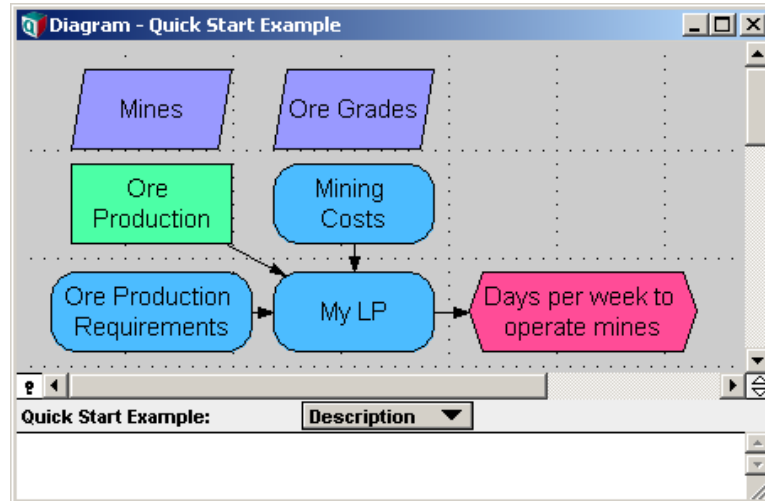
11. Select the **My_LP** node and press **Control+r** to evaluate it.



The **LpDefine()** function defines the linear program and returns a special object which displays as **<<LP>>**; however, it does not solve for the optimal solution. Follow these steps:

12. Create an objective node and title it “Days per Week to Operate Mines.” Set its units attribute to **Days per week**, and set its definition to **LpSolution(My_LP)**.

Your model should now look something like this.



- Press *Control+r* to evaluate the linear program. The result view shows the optimal number of days per week to operate each mine.

Mid Value of Days per week to operate mines (days/week)	
Mines	Totals
Mine X	1.5
Mine Y	3

It is always a good idea to check the status of the optimization as well. To check on the status of the optimization:

- Create an objective variable and name it **status**. Enter the definition **LpStatusText(My_LP)**.
- Evaluate the variable **status**.

In this case, **status** should be “Optimal solution has been found,” indicating that the solution viewed earlier was indeed the optimum. If the search had terminated early, or it could not find a feasible solution, **status** would show you the situation. See “Obtaining the solution” on page 25 for the full list of possible status values.

The example produced a non-integer solution. If you could operate each mine only for an integral number of days, and partial days are not possible, you need an integer solution. You can easily modify the problem to achieve this:

- Click **My_LP** and change its definition by adding a **Ctype** (continuity type) parameter to indicate that you want an integer solution.

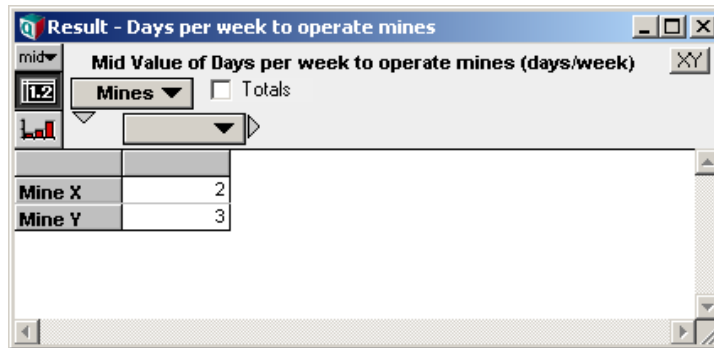
```

LpDefine(Vars: Mines,
constraints: Ore_grades,
objCoef: Mining_costs,
lhs: Ore_production,
sense: ">",
rhs: Ore_prod_req,
Ctype: "I",

```

lb: 0,
ub: 5)

17. Click **Days per Week to Operate Mine** and press *Control+r* to view the result.



A nonlinear program

You will now define and solve a nonlinear optimization. Nonlinear optimizations are treated differently from linear and quadratic optimizations. In the previous linear programming example, the coefficient matrices completely describe the problem, and the optimum solution is simply computed. A nonlinear optimization, by comparison, repeatedly re-evaluates expressions or portions of your model during a search. You will indicate the portion of your model to re-evaluate to the **NlpDefine()** function.

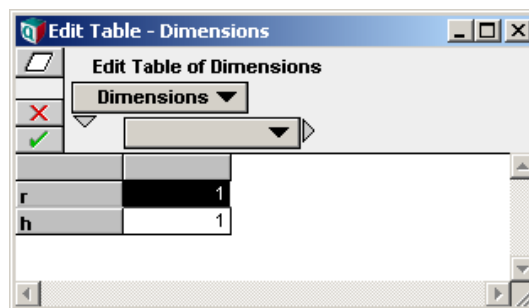
We will formulate the following optimization problem:

Find the dimensions of a cylinder with minimum surface area with a volume of at least 500 cm³.

This example can be found in the **optimal can dimensions.ana** example model in the **Example Models/Optimizer Examples** directory installed with Analytica.

To model this, we first create a self-indexed table, **Dimensions**, to index the decision variables and to hold candidate solutions.

1. Start Analytica, or select **File > New** to start a new model.
2. Create a decision variable and name it **Dimensions**.
3. Set the definition type to **Table**, select **Dimensions (Self)** for the index, and fill in the edit table.



Since it is self-indexed, the **Dimensions** variable serves both as the optimization vector and as the **Vars** index. During the optimization search, the cell values will be set to candidate solutions and other portions of the model evaluated.

For convenience, we can break out the decision variables as Analytica variables. Follow these steps:

4. Create a variable node named `radius`. Give it the definition
`Dimensions[Dimensions="r"]`

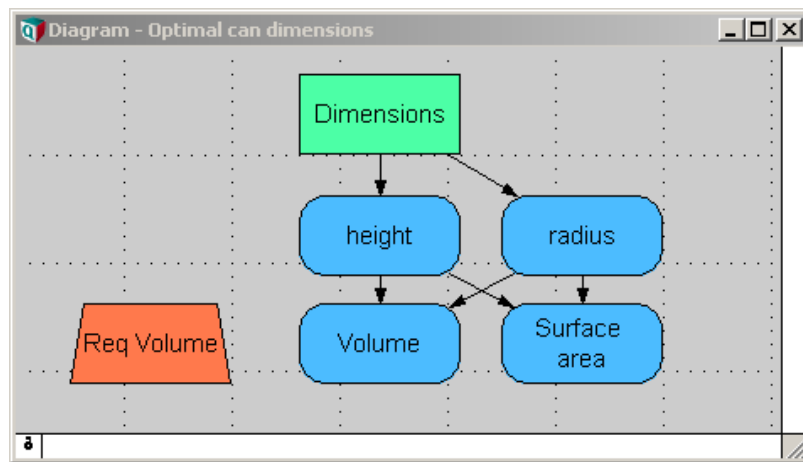
5. Create a variable, named `height`. Give it the definition
`Dimensions[Dimensions="h"]`

Next compute the surface area and volume. `Surface_area` will become the objective function. Volume will become a constraint.

6. Create a variable named `volume`. Give it the definition
`height * Pi * radius^2`

7. Create a variable named `Surface_area`. Give it the definition
`2 * Pi * radius^2 + 2 * Pi * radius * height`

8. Create a constant named `Req_Volume` (title: `Required Volume`). Set its value to 500.



Next, set up the constraints; in this case there is only one. For nonlinear problems, this involves setting up a constraint index, a left-hand side (which will be a computed expression) and a right-hand side. Sometimes it is convenient to do this as follows:

9. Create an index named `cp` with the title `Constraint Parts`. Define it as a list of labels `["lhs", "sense", "rhs"]`.

10. Create a variable named `Constraints`. Set its definition to a table and select `Constraints (Self)` and `Constraint Parts` as the indexes. Set up the edit table so that `Constraint Parts` is on the horizontal dimension and `constraints` is on the vertical dimension. Fill in the edit table as shown here.

	lhs	sense	rhs
req volume	Volume	'>	Req_volume

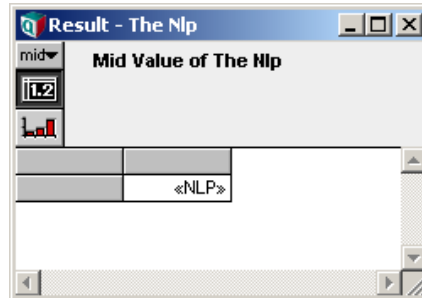
Now, define the nonlinear optimization problem:

11. Create a variable named `The_NLP`. Give it the following definition.
`NlpDefine(Dimensions, Constraints,`

```
x: Dimensions,
obj: Surface_area,
lhs: Constraints[cp="lhs"],
sense: Constraints[cp="sense"],
rhs: Constraints[cp="rhs"])
```

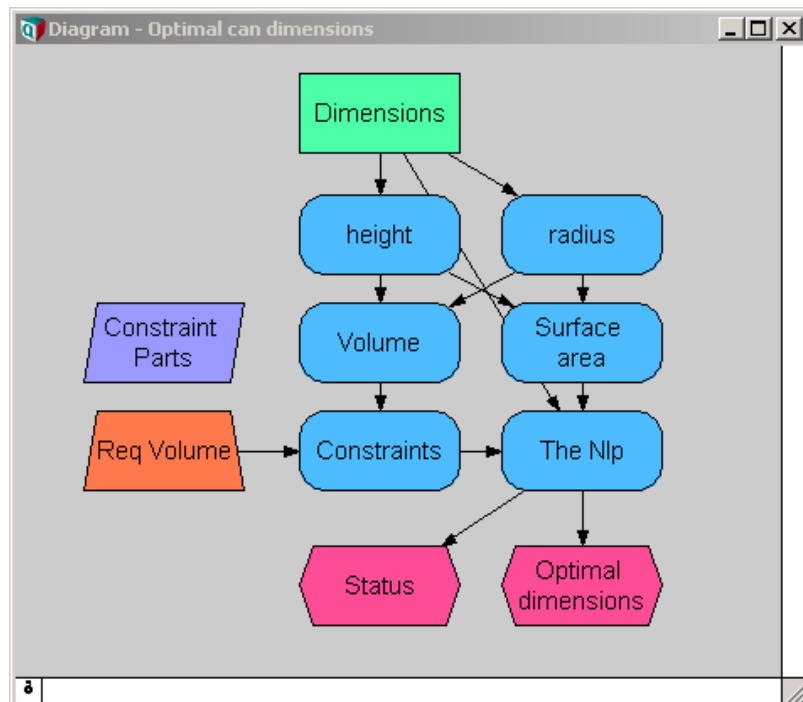
This defines the nonlinear optimization problem. The objective function is **Surface_area**, which is computed from the values in the **Dimensions** node. The left-hand side of the constraint is also computed from **Dimensions**.

When **The_NLP** is evaluated (by selecting the node and entering *Control+r*), an object is created that displays as **<<NLP>>**.



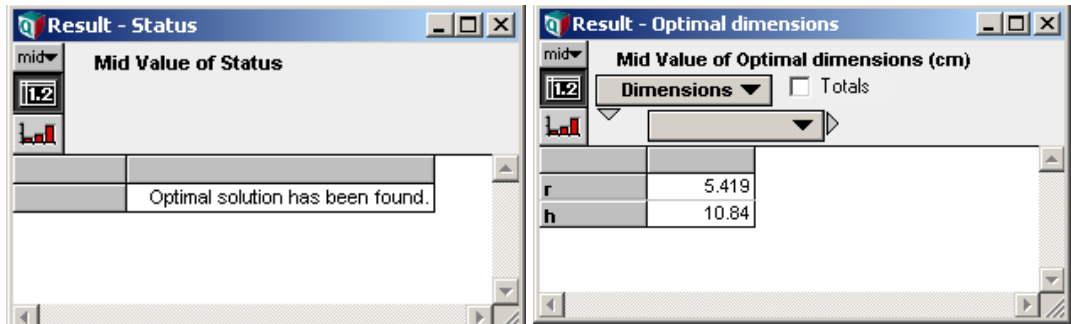
At that point, the NLP is not solved, it is only defined. It is solved when a function such as **LpStatusText()** or **LpSolution()** is evaluated. To get the solution:

12. Create an objective node named **status**, and set its definition to **LpStatusText(The_NLP)**
13. Create an objective node named **Optimal_Dimensions** and set its definition to **LpSolution(The_NLP)**



When either of these objective nodes is evaluated, the optimization engine searches for and reports the optimal solution. View the **status** node's result to make sure the optimi-

zation was successful, and view the `Optimal_dimensions` node to view the solution and its status.




Chapter 2

Formulating an Optimization Problem

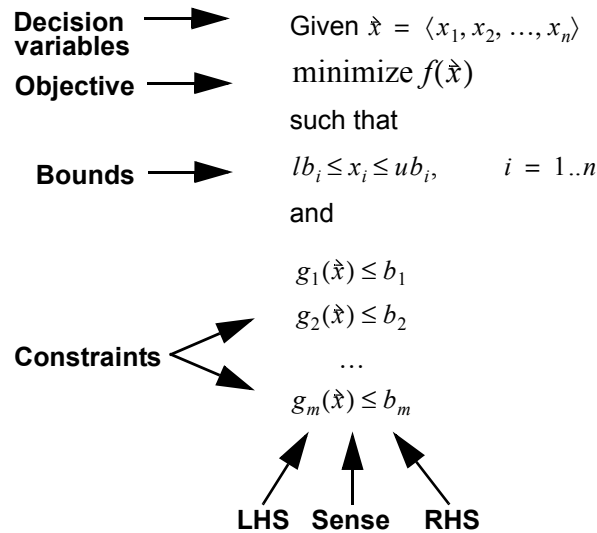
This chapter shows you:

- The different types of optimization problems
- How to choose the proper type of optimization
- How to optimize when solving simultaneous equations



Parts of an optimization problem

The first step in performing an optimization is to formulate the problem appropriately. An optimization problem is defined by four parts: a set of decision variables, an objective function, bounds on the decision variables, and constraints. The formulation looks like this.



Decision variables A vector (one-dimensional array) $\hat{x} = \langle x_1, x_2, \dots, x_n \rangle$ of the variables whose values we can change to find an optimal solution. A **solution** is a set of values assigned to these decision variables.

Objective A function $f(\hat{x})$ of the decision variables that gives a single number evaluating a solution. By default, the Optimizer tries to find the value of the decision variables that minimizes the value of objective. If you set the optional parameter **Maximize** to True, it instead tries to maximize the objective. For a linear program (LP), the objective is defined by a set of coefficients or weights that apply to the decision variables. For a nonlinear program (NLP), the objective can be any expression or variable that depends on the decision variables.

Bounds A range $lb_i \leq x_i \leq ub_i, \quad i = 1..n$ on the decision variables, defining what values are allowed. These bounds define the set of possible solutions, called the **search space**. Each decision variable can have a lower bound and/or an upper bound. If not specified, the lower and upper bounds are **-INF** and **+INF** — that is, there are no bounds.

Constraints The constraints, e.g., $g_1(\hat{x}) \leq b_1$, are bounds on functions of the decision variables. They define which solutions are feasible.

Each constraint consists of a **left-hand side (LHS)** $g_1(\hat{x})$, which is a function of the decision variables, \hat{x} , a **sense**, (<, =, or >) defining the direction of the constraint, and a **constant**, e.g., b_1 .

Continuous, integer, and mixed-integer programs

Each decision variable can be specified as **continuous**, meaning it is a real number (between bounds if specified), as **integer**, meaning a whole number, as **binary** or **Boolean**, meaning its values can be True (1) or False (0), or as a member of an integer **group**, where each member of the group must have a different integer value. Optimiza-

tion problems are classified as **continuous**, meaning the decision variables are all continuous; **integer**, meaning they are all integer, binary, or group variables; or **mixed-integer** if they are a mixture of continuous and integer, binary, or group variables. In this naming convention, binary or Boolean variables are treated as integer variables. The Optimizer engine uses these distinctions to select which algorithms to use.

Choosing the type of optimization

A critical issue in formulating an optimization problem is determining whether it is linear, quadratic, or nonlinear. For a **linear program (LP)**, the objective must be a linear function of the decision variables. For a **quadratic program (QP)**, the objective and constraints must be linear or quadratic functions of the decision variables. The problem is a **nonlinear program (NLP)** if the objective or any of the constraints are nonlinear in any of the decision variables.

You define the type of a problem by using the function **LpDefine()**, **QpDefine()**, or **NlpDefine()**, respectively. You provide the decision variables, objective, bounds, and constraints as parameters to the selected function, along with other optional parameters.

Linear and convex quadratic optimization problems are often relatively fast to compute. But general nonlinear optimization is a computationally difficult problem. Many of the most famous and notoriously difficult computation problems can be cast as optimization programs, from the traveling salesman to the solution (or non-solution) of Fermat's last "theorem." It is, therefore, unreasonable to expect the Optimizer engine to succeed on any possible nonlinear problem you can formulate. While the Frontline Solver engine used in the Analytica Optimizer is among the best of the general-purpose optimization engines available, success with hard optimization problems depends on your ability to formulate the problem effectively, provide appropriate hints for the Optimizer, and adjust the search control settings.

Linear and quadratic optimization in Analytica fully support Intelligent Arrays™ — that is, any of their parameters can be arrays with additional dimensions, and Analytica performs an array of optimizations to compute an array of optimal values. For example, any parameter can be uncertain, defined as a random sample, and the optimization can be carried out within a dynamic loop, for each time step. In contrast, NLP is subject to restrictions on array abstraction, particularly in models with uncertain factors in the objective or constraints, or when used in dynamic loops. There are ways around these limitations, which we describe in "Array abstraction" on page 45. However, it is easier to manage array abstraction, particularly in dynamic simulation, with linear or quadratic optimization problems.

There are often several ways to formulate the same problem. Linear and quadratic formulations are faster and more flexible, so it is worth careful thought to see if it is possible to reformulate a nonlinear optimization into a linear or quadratic optimization. Often a simple transformation, combination, or disaggregation of the decision variables can turn an apparently nonlinear problem into a linear or quadratic problem.

Solving simultaneous equations


The Optimizer first attempts to find a feasible solution. If found, it then attempts to optimize within the set of feasible solutions. Thus, solving a set of simultaneous equations is a special case of the optimization problem, where each constraint has a sense of =, the objective is irrelevant (unless you want to express a preference among feasible solutions), and any feasible solution is a solution to the system of equations.

Chapter 3

Linear Optimization


This chapter shows you how to:

- Define a linear optimization problem
- Obtain a solution
- Deal with integer and binary decision variables
- Control the search



Defining a linear optimization problem

A linear optimization problem has the following standard formulation.



In this standard form, all decision variables, x_i , are real-valued and unconstrained, ranging from **-INF** to **+INF** ($-\infty$ to ∞).

To encode this in Analytica, use the function **LpDefine()**.

```
LpDefine(Vars, Constraints: Index;
  ObjCoef: Number[Vars];
  LHS: Number[Vars, Constraints];
  RHS: Number[Constraints] )
```

These are the required parameters of **LpDefine()**.

Vars An index over the **n decision variables**, $[x_1, x_2, \dots, x_n]$, for which we wish to find the optimal **solution** — that is, the values that minimize (or maximize) the objective. The index has one element for each decision variable. You can define it as a list of numbers, **1..n**, or a list of labels to give meaningful names to each decision variable.

Constraints An index over the set of **m constraints**, with one element for each constraint. Again, you can define it as a list of numbers, **1..m**, or a list of labels to give meaningful names to each decision variable.

ObjCoef The **objective coefficients**, an array of **n coefficients**, $[c_1, c_2, \dots, c_n]$, indexed by **Vars**. The objective we are trying to minimize (or maximize) is the dot product of these objective coefficients and the decision variables — that is, $c_1 x_1 + c_2 x_2 + \dots + c_n x_n$.

LHS The **left-hand side (LHS)** of the constraints is an **n by m array of coefficients**, indexed by **Vars** and **Constraints**, $a_{11}, a_{12}, \dots, a_{ij} \dots a_{mn}$.

RHS The **right-hand side (RHS)** of the constraints, being an array of **m constants**, (b_1, b_2, \dots, b_m) indexed by **Constraints**. The constraints, e.g., $g_1(\hat{x}) \leq b_1$, are bounds on functions of the decision variables. They define which solutions are acceptable.

Each constraint consists of a **left-hand side (LHS)** $g_1(\hat{x})$, which is a function of the decision variables, \hat{x} : a **sense**, (**<**, **=**, or **>**) defining the direction of the constraint, and a **constant**, e.g., b_1 .

When **LpDefine()** is evaluated, the result is a special linear program object, which displays as **<<LP>>**. This defines the linear program, but does not compute the optimum. That information is obtained through a series of functions described in “Obtaining the solution” on page 25.

Optional parameters

You can specify a wide set of optional parameters to **LpDefine()** for variations on the basic formulation shown above. These options include lower and/or upper bounds on the decision variables, maximizing instead of minimizing the objective, and changing the direction (**sense**) of the constraints from \leq to \geq or $=$.

You can specify these optional parameters to **LpDefine()** in any order by listing each parameter name, followed by a colon, followed by the value. Here is an example.

```
LpDefine(Vars: VarIndex, Lb: 0,
         Constraints: ConIndex, Lhs: lhs, Sense: ">=", Rhs: rhs,
         ObjCoef: ObjCoef, Maximize: True)
```

In this case, **Vars**, **Constraints**, **Lhs**, **Rhs**, and **ObjCoef** are the required indexes and coefficients described in the previous section, and the other parameters are optional parameters, specifying that we want to maximize the objective, each decision variable (x_1, \dots, x_n) has a lower bound (**Lb**) of zero, and all constraints have sense \geq , instead of the default \leq .

Lower and upper bounds on decision variables

You can specify lower and upper bounds on decision variables using these optional parameters.

```
Lb, Ub: Optional Number[Vars]
```

By default, **Lb** = **-INF** and **Ub** = **+INF**. If you give a single number to one of these parameters, it specifies the same bound for all decision variables. To specify a different bound for each decision variable, give it an array of values indexed by **Vars**.

Lower and upper bounds for binary and grouped-integer variables are ignored, even if specified, since these are fixed. Binary variable bounds are 0-to-1, and grouped-integer bounds are $1..N$, where **N** is the number of decision variables in the same group.

Maximizing the objective

The optional parameter **Maximize** should be either True or False, specifying whether Analytica Optimizer should attempt to maximize or minimize the objective function. If not specified, it defaults to False, and minimizes the objective function.

Sense of constraints

The **sense** of a constraint refers to whether the left-hand side is \leq , \geq , or $=$ to the right-hand side. The **Sense** parameter is used to specify the sense for each constraint.

```
Sense: Optional Text[Constraints]
```

When omitted, it assumes \leq by default. The following text values are recognized.

- **<**, **<=**, or **L** : LHS is less-than or equal to RHS
- **>**, **>=**, or **G** : LHS is greater-than or equal to RHS
- **=** or **E** : LHS is equal to RHS

If a single value is passed to the **Sense** parameter, that sense applies to all constraints. If each constraint has a different sense, then the **Sense** parameter should be an array indexed by constraints.

Obtaining the solution

The optimal values for the decision variables, x_1, \dots, x_n , are obtained using the **LpSolution()** function, which takes as a single parameter the **<<LP>>** object created by **LpDefine()**, and which returns an array indexed by the **Vars** index. The value of the objective function at the optimum is obtained using the **LpOpt()** function.

LpSolution(lp: LpType)

Returns the optimal solution to the programming problem **lp** defined by **LpDefine()**. The result is an array of decision variables indexed by **Vars**. If the Optimizer cannot find an optimal solution, it returns the best values found during the search so far.

LpOpt(lp: LpType)

Returns the value of the objective function for linear program **lp** at the optimum. For a linear problem, the value it returns is equal to this.

```
Sum(LpSolution(lp) * ObjCoef, Vars)
```

LpStatusNum(lp: LpType) and LpStatusText(lp: LpType)

These two functions return, respectively, the status number and the corresponding text describing the status of the solution, or why the optimization search terminated, for the programming problem **lp**. If the optimization is successful, these results will be 0 and “Optimal solution has been found” respectively. If not successful **LpStatusNum()** returns another number and **LpStatusText()** returns different text explaining why it has not found an optimal solution.

Note: *The numeric codes from **LpStatusNum()** and the corresponding text from **LpStatusText()** have changed in Analytica Optimizer 4.0. These reflect changes in the status numbers and text returned by the Frontline System's new, restructured Optimizer.*

Tip Starting with version 4.0, you can use a different engine than the Frontline Optimizer that comes with Analytica Optimizer. If a different Optimizer engine is used, different engine-specific codes might be returned. In these cases, **LpStatusText()** returns “unknown Frontline solver status code,” and the result returned by **LpStatusNum()** depends on the status number returned by the Optimizer engine. Consult the documentation for the engine you are using.

Possible outcomes to an optimization include:

1. It found a global optimum.
2. There is no feasible solution, because the constraints are contradictory.
3. The optimal solution is unbounded, because the constraints (if any) do not prevent the objective function from approaching $-\infty$ (for a minimization problem).
4. The search terminates with a feasible solution, but before an optimal solution is found. This happens when the computation time or number of pivots exceeds the termination criteria before a feasible solution has been located (see “Controlling the search” on page 31).
5. The search terminates before finding a feasible solution.

These different cases can be detected using the **LpStatusNum()** or **LpStatusText()** functions, both of which take **lp** as a single parameter, and which return the values shown in the table below for a continuous linear program.

Tip **LpSolution()** often returns the best solution “point” so far, even in cases in which the global optimum was not located, so it is important to check the status.

Status Number	Status Text
-3	Invalid status.
-2	Ignore status. Used when dummy result code needs to be overridden.
-1	Invalid license status. (License expired, missing, invalid, etc.)
0	Optimal solution has been found.
1	The Solver has converged to the current solution.
2	"No remedies" status. (All remedies failed to find better point.)
3	Iterates limit reached. Indicates an early exit of the algorithm.
4	Optimizing an unbounded objective function.
5	Feasible solution could not be found.
6	Optimization aborted by user. Indicates an early exit of the algorithm.
7	Invalid linear model. Returned when a linearity assumption renders incorrect.
8	Bad data set status. Returned when a problem data set renders inconsistent.
9	Float error status. (Internal float error.)
10	Time out status. Returned when the maximum allowed time has been exceeded. Indicates an early exit of the algorithm.
11	Memory dearth status. Returned when the system cannot allocate enough memory to perform the optimization.
12	Interpretation error. (Parser, Diagnostics, or Executor error.)
13	Fatal API error. (API not responding.)
14	The Solver has found an integer solution within integer tolerance.
15	Branching and bounding node limit reached. Indicates an early exit of the algorithm.
16	Branching and bounding maximum number of incumbent points reached. Indicates an early exit of the algorithm.
17	Probable global optimum reached. Returned when MSL (Bayesian) global optimality test has been satisfied.
18	Missing bounds status. Returned for EV/MSL Require Bounds when bounds are missing.
19	Bounds conflict status. Indicates \leq $=$ $>$ bounds conflict with existing binary or all different constraints.
20	Bounds inconsistency status. Returned when the lower bound value of a variable is greater than the upper bound value, i.e., $lb[i] > ub[i]$ for some variable bound i .
21	Derivative error. Returned when API_Jacobian has not been able to compute gradients.
22	Cone overlap status. Returned when a variable appears in more than one cone.
999	Exception occurred status. Returned when an exception has been caught by try/catch top-level.
1000	Custom base status. (Base for Solver engine custom results.)
1102	The quadratic constraints are non-convex, the SOCP engine cannot solve this problem.

Secondary aspects to a solution

The solution to a linear program contains more information than just the optimal solution, (x_1, \dots, x_n) . Often these secondary elements of the solution are of more value than the solution itself for decision making purposes, since they indicate how changes (e.g., different decisions) impact the optimum. These secondary aspects of the solution are accessed using the functions `LpSlack()`, `LpObjSa()`, `LpRHSSa()`, `LpShadow()`, and `LpReducedCost()`.

Slack or surplus: `LpSlack(lp: LpType)`

When you have a constraint

$$a_{i1} x_1 + a_{i2} x_2 + \dots + a_{in} x_n \leq b_i$$

the slack (or surplus) for that constraint is the positive value that, when added to the LHS, makes both sides equal, that is

$$a_{i1} x_1 + a_{i2} x_2 + \dots + a_{in} x_n + \text{slack}_i = b_i$$

The constraints that have zero slack are of particular interest, since they are instrumental in constraining the optimum. If these constraints are relaxed (e.g., by increasing b_i), a larger maximum value can be obtained. However, as critical constraints are relaxed, other constraints might become relevant. For the constraints, the non-zero slack gives an indication of how close they are to becoming critical.

The slack for each constraint is obtained from this function.

```
LpSlack(Lp)
```

It takes as input the object returned from **LpDefine()** and returns an array indexed by **Constraints**, containing the slack at the optimum for each constraint.

Coefficient sensitivity: **LpObjSa()** and **LpRhsSa()**

If we change a coefficient in the objective function, the solution (x_1, \dots, x_n) continues to be the optimal solution as long as the coefficient remains within a certain range. Note that the solution point is the same, but the value of the objective function at the optimum is effected. This range can be computed with this function.

```
LpObjSa(Lp: LpType; Var: optional)
```

The first parameter, **Lp**, is a linear program defined using **LpDefine()**. When called with only a single parameter, the range is computed for all decision variables, and the result is indexed by the linear program variable array **Vars**. If the range for only a single decision variable (or a small subset) is required, the second parameter, **Var**, is used to indicate the decision variable for which the sensitivity is to be computed. The second parameter should be an element (or a subset) of the **Vars** index.

The result returned from **LpObjSa()** is dimensioned by a local index, **.range:= ['lower', 'upper']**. Thus, to get the smallest value for each coefficient in the objective that would continue to produce the same solution, you would use an expression like this.

```
Var sa:= LpObjSa(myLp) DO  
sa[.range='lower']
```

Note: The **LpObjSa()** function can only be used with a linear program. It is not meaningful for quadratic or nonlinear programs.

The sensitivity of the right-hand side coefficients can be computed using this function.

```
LpRHSSa(Lp: LpType; constraint: Optional)
```

This computes the range over which the coefficient in the RHS can vary without changing the basis of the solution. In other words, over the returned range, the set of constraints with zero slack remains the set of constraints with zero slack (i.e., the critical constraints).

The result is indexed by a local index, **.range:= ['lower', 'upper']**, containing the smallest and largest values for the corresponding RHS coefficient. If the optional second parameter is not specified, the range is computed for all variables and the result is indexed by **Vars**. If the range is needed for only a single coefficient, the second parameter specifies an element of the **Constraints** index, and only the range for that constraint is computed.

When a coefficient can be changed an arbitrary amount without changing the solution basis, the corresponding entry in the result returned by **LpRhsSa()** or **LpObjSa()** is **-INF** for the lower value or **+INF** for the upper value.

Dual values: shadow prices and reduced costs

If a constraint is relaxed, i.e., by increasing the right-hand side, b_i , by one unit, how does this impact the objection function? This is referred to as the **shadow price**, or **dual value**, of the constraint. A shadow price is valid only for small changes in b_i (the actual range for which it is valid can be obtained from the **LpRhsSa()** function), and is computed by the function

LpShadow(lp: LpType)

where **lp** is a linear program object returned by **LpDefine()**. The result is indexed by **Constraints**. Mathematically, the shadow price is given by this equation.

$$\text{Shadow}_i = \frac{\partial \text{Obj}}{\partial b_i}$$

This is the partial derivative of the objective function relative to the constraint RHS coefficient.

Warning: *Not all linear programming packages use the same convention for the sign of shadow prices. If you have used the LINDO package, note that the convention used by Analytica Optimizer differs from the sign produced by the LINDO package.*

How far can a coefficient in the objective function be increased (in a minimization program) or decreased (in a maximization program) before the objective function changes? When a decision variable has a non-zero value in the optimal solution, any change in the objective function coefficient changes the objective value, so for those decision variables the answer would be zero. But for decision variables that are zero, the coefficient can change until that variable eventually enters the basis. This amount is known as the **reduced cost** (or dual value) of the variables and is returned by the function

LpReducedCost(lp: LpType)

The result is indexed by **Vars**.

The shadow price and reduced cost are known as **dual values**, the shadow price being a dual to the solution in the original (or “primal”) problem, and the reduced cost being a dual to the slack price in the original problem. To each problem in the standard form (see “Defining a linear optimization problem” on page 24) there corresponds a dual linear program given by this.

$$\text{maximize } b_1 y_1 + b_2 y_2 + \dots + b_m y_m$$

such that

$$a_{11} y_1 + a_{21} y_2 + \dots + a_{m1} y_m \geq c_1$$

...

$$a_{1n} y_1 + a_{2n} y_2 + \dots + a_{mn} y_m \geq c_n$$

The new variables in this program, y_1, y_2, \dots, y_m , are the shadow prices, and the slack value for each constraint is the reduced costs in the primal problem. Note that the variables in the primal problem correspond to constraints in the dual problem, and constraints in the primal problem correspond to decision variables in the dual problem.

Examples

Several example linear-programming optimization models are included in the **Example Models/Optimization Examples** folder installed with Analytica. The linear program examples include:

- **Automobile production.ana**: Taking differences in unit production cost, and labor and material availability into consideration, figure out how many cars to produce at each factory to meet a production goal. This example demonstrates the use of linear program-related sensitivity functions.
- **Big Mac Attack.ana**: Optimize your McDonald's-based diet to fit your budget and nutritional needs, and minimize your calorie or carbohydrate consumption.
- **Capital Investment.ana**: Simple case of selecting which projects to pursue given a fixed budget.
- **Optimal production planning.ana**: A classic textbook linear program of selecting how much of each product to produce given resource limitations.
- **Production Planning LP.ana**: Another take on the same problem, but demonstrating the interpretation of the secondary solution aspects.
- **Two Mines Model.ana**: Schedule production at multiple mines to meet production goals given capacity constraints. (This is the example used in Chapter 1, "Quick Start.")
- **Sudoku with Optimizer.ana**: Solves Sudoku puzzles. Demonstrates use of grouped integer variable types with many groups.

Integer, binary, and grouped decision variables

In a standard linear program the decision variables are assumed to be continuous (real-valued) numbers. However, you can also use Analytica Optimizer to define and optimize a linear program with some or all of the decision variables constrained to be integers, Boolean or binary, grouped-integer, or a mixture of continuous and integer, binary and grouped-integer variables (a ***mixed-integer program***).

You can specify the type of each decision variable as continuous, integer, binary, or grouped using the optional parameter

```
Ctype: Optional Text[Vars]
```

This parameter takes one of the following values:

- **c**: Continuous
- **i**: Integer
- **B**: Binary or Boolean value, i.e., 0 or 1
- **G**: Grouped integer, i.e., one of **1..N**, where **N** is the number of decision variables in the same group, and such that no variables in a group have the same value.

If you give the **Ctype** parameter a single text character, it specifies the same type for all decision variables. Here is an example.

```
LpDefine(..., Ctype: "B")
```

This specifies that all decision variables are binary. To specify a ***mixed-integer program***, you supply an array of characters, indexed by **Vars**, specifying the type of each decision variable.

When you have grouped-integer variables partitioned into two or more groups, the **group** parameter specifies which group each of the grouped-integer variables belongs to

```
group : Optional Number[Vars]
```

For example, if your grouped-integer decision variables are partitioned into three groups, the elements of the array passed to **group** would be 1, 2, or 3 to indicate the group that each variable belongs to, and 0 for each of the non-grouped-integer decision variables. Each decision variable can belong to at most one group, and each group should have at least two grouped-integer decision variables. The example model **Sudoku with Optimizer.ana** demonstrates the use of multiple grouped-integer groups.

In general, integer and mixed-integer linear programs are harder to solve than linear programs with exclusively continuous variables. The Optimizer uses a combination of a simplex algorithm with a memory-efficient branch-and-bound algorithm.

In some cases, the Optimizer might fail to find a solution to a large integer or mixed-integer linear program. Use the **LpStatusNum()** and **LpStatusText()** functions to see whether it has been successful, and if not, why not. For a complete list of the possible values returned by **LpStatusNum()**, see “LpStatusNum(lp: LpType) LpStatusText(lp: LpType)” on page 92.

Controlling the search

A linear program with all continuous decision variables is solved using a simplex algorithm. The space of feasible solutions is called a **simplex** and is a convex polyhedron in N-dimensional space, where N is the number of decision variables. A simplex algorithm traverses the simplex from corner to corner, moving to an adjacent corner with an improved objective value at each iteration (**pivot**). The objective is improved with each pivot until the global optimum is reached. The same algorithm is used on an augmented simplex initially to find an initial feasible solution.

An integer, binary, or mixed-integer program uses the simplex algorithm in combination with a branch-and-bound algorithm. It first uses the simplex to solve the continuous version of the problem. This bounds the optimal objective from one side and provides a starting point for a search. Whenever it finds a feasible integer solution, this provides a bound on the optimal objective on the other side and allows the branch-and-bound search to prune alternative integer solutions that would be provably inferior to the ones already found. As the algorithm explores solutions having one integer decision variable set to a particular integer value, the continuous LP subproblem is solved again using repeated invocations of the simplex algorithm on increasingly constrained problems. It terminates the search when the search space has been exhausted (i.e., the global optimum located), when the termination criteria have been exceeded, or when the best solution found is within the solution (gap) tolerance. In addition, logical implications of integer constraints can often be deduced before simplex is even run. Various algorithms for finding these constraints are referred to as **cuts**, and various algorithms for recognizing and using cuts of different types can be switched on or off.

Many settings controlling the precise behavior of the Optimizer can be altered using the two parameters to **LpDefine()** named **parameter** and **setting**. The list of all possible parameters is the topic of Chapter 6, “Control Settings.”

Viewing and specifying control settings

When you have defined a linear program using **LpDefine()**, the following function returns the set of control settings used by the engine

```
SolverInfo("Setting", Lp: myLp )
```

Just replace *myLp* with the name of the variable holding the result from **LpDefine()**.

You can also access the range of allowed values for each setting, as well as the default value, using **SolverInfo()**. For this, you need to know the name of the Optimizer engine used on your problem. For linear programs, this is always **LP/Quadratic** unless you have installed an add-on engine. To obtain the name of the engine used in the general case, use this.

```
SolverInfo("Engine", Lp: myLp )
```

Using the name of the engine, the range (min/max) of possible values for each setting and the default value can be obtained using this format.

```
SolverInfo( ["MinSetting", "MaxSetting", "Defaults"],
           Engine: "LP/Quadratic" )
```

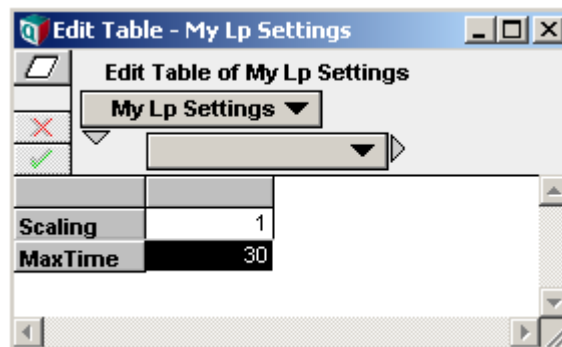
If you want to change the value for a single control setting, you can specify values for two optional parameters, **parameter** and **setting**, to **LpDefine()**, providing the name of the setting to **parameter**, and the value to **setting**. For example, if you want to set the **Scaling** parameter to 1, modify your call to **LpDefine()** as follows

```
LpDefine( .., Parameter: "Scaling", Setting: 1 )
```

To alter more than one control setting, you need to supply arrays to these parameters. The arrays passed to **parameter** and **setting** should have a single common index. If the index of the array passed to **setting** is a list of labels, where the index labels contain the name of each control setting, then you only need to include the **setting** parameter.

It is often convenient to specify control settings in a self-indexed edit table. Follow these steps:

1. Drag a variable node to your diagram, and title it **My Lp Settings**.
2. In the definition pane, set the definition type to **Table**.
3. In the **Index Chooser** dialog, select **My Lp Settings (Self)** as the table index.
4. Click the row heading cell, and change **Item 1** to **Scaling**.
5. With the row header still selected, press *down-arrow* to add a row.
6. Change the second row header cell to **MaxTime**.
7. Enter **1** into the first table body cell.
8. Enter **30** into the second body table cell.



9. In your call to **LpDefine()**, insert a setting parameter as follows

```
LpDefine( ..., setting: My_lp_settings )
```

The Optimizer scales parameters and terminates after 30 seconds if the optimum has not been found. A self-indexed table set up in this fashion makes it easy to adjust multiple control settings if the need arises.

Note: In Analytica Optimizer 3.1, control settings were specified as optional parameters to **LpDefine()**. These legacy parameters are still supported for backward compatibility; however, the use of the **setting** parameter is recommended. This change reflects a change in Frontline's architecture, and more readily generalizes to other add-on engines and future Optimizer engine extensions.

Linear programming settings

The list and descriptions of all settings to Optimizer engines is covered in Chapter 6, "Control Settings." Continuous linear problems usually solve very quickly and reliably, and seldom require much tuning. Integer programs, on the other hand, can be very complex to solve in some cases, so that experimentation with engine settings can be justified with hard problems. Here is a brief list of some of the settings that might be of interest, but see Chapter 6 for detailed descriptions.

- Termination Control**
- Iterations:** Maximum number of pivots by the simplex algorithm.
 - MaxTime:** Maximum number of seconds the Optimizer spends on the problem.
 - MaxTimeNoImp:** The maximum number of seconds with no substantial improvement.
 - Tolerance:** The amount of improvement required within **MaxTimeNoImp** in order to continue.
 - IntTolerance:** If branch-and-bound can prove its current solution is within this percentage of the true optimal, it stops.
- Preprocessing**
- Scaling:** Whether to scale decision variables and constraints for simplex. If coefficients vary by many orders of magnitude, numeric instabilities might result without scaling.
 - Presolve:** When on, the engine performs a presolve step removing singleton rows or columns, fixed variables, and redundant constraints, and tightening bounds.
 - StrongBranching:** Performs experiments prior to solving to estimate the impact of branching on each integer variable. The up-front cost might pay off later in more efficient branch-and-bound searches.
- Branch-and-bound hints**
- IntCutoff:** A bound you provide in advance on the objective function value for the optimal solution. This is an upper bound for a minimization or a lower bound for a maximization. If you can provide such a bound, the branch-and-bound algorithm might be able to prune huge portions of its search space.
 - UseDual:** Controls whether the dual or primal basis is used by simplex to solve subproblems generated by branch-and-bound.
 - ProbingFeasibility:** Engine attempts to deduce values of certain binary variables based on settings of others, prior to actually solving a subproblem.
 - PrimalHeuristic:** Uses a heuristic method to attempt to discover a solution early in the branch-and-bound process. This can be an effective way to locate an **IntCutoff** early, to dramatically prune the branch-and-bound search space later.
- Cut generation**
- MaxRootCutPasses, MaxTreeCutPasses:** Control how many cut passes are carried out at each step of the solution process.

GomoryCuts, KnapsackCuts, ProbingCuts, OddHoleCuts, MirCuts, TwoMirCuts, RedSplitCuts, SOScuts, FlowCoverCuts, CliqueCuts, RoundingCuts, LiftAnd-CoverCuts: These are all different *cut* methods that attempt to deduce constraints from existing integer commitments, thus reducing the space of feasible solutions. Each introduces a time overhead, that could be spent searching instead, but in problems where a method is effective, speed-up can be dramatic.

Tolerance and precision **ReducedTol, PivotTol:** Control which rows or columns are candidates for pivots during the simplex algorithm.

Precision, PrimalTolerance: Control the amount of numeric error by which a constraint can be violated and still be considered satisfied. For example, in an equality constraint, you don't want 1-bit of numeric error to result in an unsatisfied constraint.

Array abstraction

As with most Analytica functions, **LpDefine()** and all the functions used to retrieve the solutions to a linear program are fully array-abstractable. If, for example, you supply an array of coefficients to the **ObjCoef** parameter of **LpDefine()** that is indexed by index *In1* in addition to the *Variables* index, **LpDefine()** returns multiple <<LP>> objects, with the collection being indexed by *In1*. When such a result is solved, multiple optimization problems are run.

If any parameter that expects a particular dimension is supplied an object without that dimension, **LpDefine()** will treat it as if that dimension were specified with the value constant across that dimension. So, for example, specifying the parameter **RHS: 1** would treat the right-hand-side of every constraint as having the value 1.

Because these functions are fully array abstractable, any coefficient, bound, or other parameter can be uncertain, evaluated as a sample (indexed by **Run**), computed from probability distributions or chance variables. When evaluated in probabilistic mode, these models solve a separate optimization problem for each sample.


Linear programs involving time can also be embedded in dynamic loops (see Chapter 17, "Dynamic Simulation," in the *Analytica User Guide*). By specifying a parameter value that is a function of a previous time step, and using **LpDefine()** from within a dynamic loop, a separate optimization can be performed at each **Time** point.

Chapter 4

Quadratic Optimization

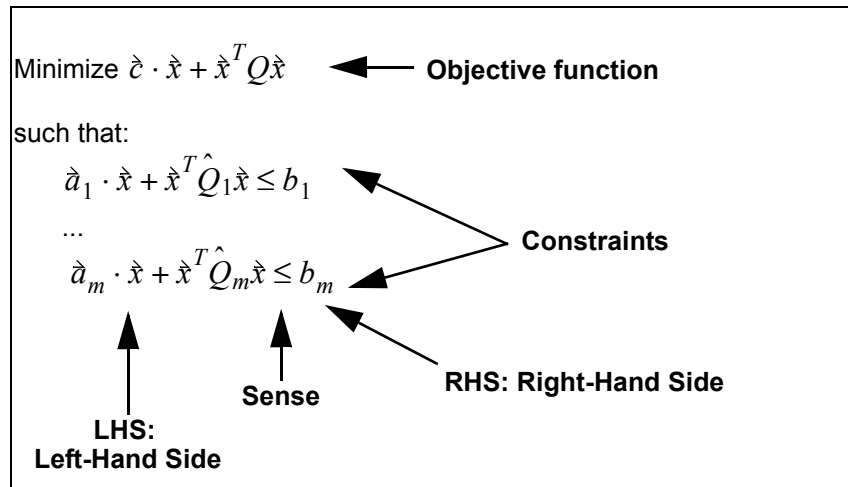
This chapter shows you how to:

- Define a quadratic optimization
- Solve a quadratic optimization
- Use sample quadratic optimizations as a starting point for your own optimizations



Defining a quadratic program

The general form for a quadratically constrained quadratic program accepted by Analytica Optimizer is:



The objective and constraint left-hand sides are written here in matrix notation. The term $\hat{c} \cdot \hat{x}$ is the linear part of the objective, and each constraint has a linear part $\hat{a}_i \cdot \hat{x}$. These linear parts are the same as the objective and constraint left-hand sides of a linear program, with coefficient vectors \hat{c} , \hat{a}_1 , \hat{a}_2 , ..., \hat{a}_m . This formulation augments the linear program by adding quadratic terms to the objective and to constraints. In a pure quadratic program, only a quadratic objective is used, and the constraints are all linear, so that the Q_i matrices are omitted (or 0). The more general case in which constraints can also be quadratic is referred to as a **quadratically constrained quadratic program**, but for conciseness, we use **quadratic program** to cover the general case.

The quadratic terms, i.e., $\hat{x}^T Q \hat{x}$ in the objective and $\hat{x}^T \hat{Q}_i \hat{x}$ in the constraint are specified by the $n \times n$ matrices Q , Q_1 , Q_2 , ..., Q_m , where n is the number of decision variables and m is the number of constraints, and \hat{x}^T denotes the vector transpose of the decision variables.

To ensure that the Q matrices are square, you need to specify a second index (**Vars2** in the example on the next page) with the same number of elements as the first index, **Vars**. (An array in Analytica can be indexed only once by the same index.)

A quadratic program is defined with these parameters:

Required parameters

```
QpDefine(Vars, Vars2, Constraints: Index:
  c: Optional Number[Vars];
  Q: Numeric[Vars, Vars2];
  Lhs: Numeric[Vars, Constraints];
  LhsQ: Number[Vars, Vars2, Constraints];
  Rhs: Numeric[constraints];
```

Optional parameters

```
sense: Optional Text[Constraints];
maximize: Optional Boolean;
lb, ub: Optional Number[Vars];
```

```

Ctype: Optional Text[Vars];
group: Optional Text[Vars];
guess: Optional Number[Vars];
Parameter: Optional Text;
Setting: Optional Number;
Engine: Optional Text
)

```

Note: The parameters to **QpDefine()** are not shown here in exactly the same order that the function expects. The **LhsQ** parameter has been moved up to its logical position in the text for clarity, and a several deprecated parameters are not shown. Because there are so many optional parameters to this function, you should always use the named-parameter calling convention when using **QpDefine()**. In the named calling convention, you precede each argument the parameter name, as in the following example.

```

QpDefine ( Vars: Part, Vars2: Part2, Constraints: ConstraintIndex,
c: coef, Q: Qcoef, Lhs: LhsCoef, LhsQ: LhsQcoef, Rhs: b
sense: '<=', maximize: True )

```

When evaluated, **QpDefine()** returns a quadratic program object, which displays as <<QP>>. The optimum solution is not solved until one of the routines to access the solution, such as **LpStatusNum()** or **LpSolution()**, is called.

Optional parameters The parameters **sense** (<=, =, or >=), **maximize** (true to maximize objective), **lb** and **ub** (upper and lower variable bounds), **Ctype** (continuous/integer type **c**, **v**, **i**, or **g**), **group** (integer-group), and **parameter** and **setting** (for specifying search control settings) are all optional parameters of **LpDefine()** and are described in “Linear Optimization” starting on page 23. As with linear programs, Analytica Optimizer supports integer, binary, and mixed-integer quadratic programs.

The optional **guess** parameter provides an initial guess for a solution which might be used by the optimization engine to disambiguate multiple extrema when a Q matrix is indefinite (see below).

Engine The optional **engine** parameter can be used to explicitly select which optimization engine to use to solve the problem. By default, **engine:LP/Quadratic** is used with a quadratic objective and linear constraints, and **engine:SOCP Barrier** is used when quadratic constraints exist. If the constraints are not convex, the SOCP Barrier engine can return a status 1102 (“The quadratic constraints are non-convex”). If this happens, you might need to use **engine:GRG Nonlinear** to obtain a solution.

Solution properties

The Hessian of the objective function is a second partial derivative of the objective relative to each pair of decision variables, and is given by $Q + Q^T$. Depending on the values in this Q matrix, the objective function can have a number of different shapes, and the objective can contain a single extreme (minimum or maximum), an infinite number of extrema, or no extreme values. The optimum value to a quadratic program can lie at the objective’s extrema, or it can exist on a constraint boundary.

Positive and negative-definiteness When the Q matrix of a minimization problem is *positive-definite*, meaning that for all non-zero \hat{x} : $\hat{x}^T Q \hat{x} > 0$: the objective function has a “bowl” shape with a single extrema. Similarly, for a maximization problem if it is *negative-definite* it has a bowl-shape with a single extrema. When the extrema is a feasible solution, it is the unique optimal solution

to the quadratic program. The quadratic programming algorithms are optimized for this case.

Semi-definiteness When the q matrix is *positive semi-definite* (or *negative semi-definite* for a maximization problem), the objective has a “trough” with an infinite number of extrema. In such a case, the Optimizer finds one of the feasible points in the trough.

Indefinite objective If the q matrix is *indefinite*, the objective has a **saddle point**. Like an extrema, a saddle point has a zero gradient, but is not an actual optimum. The true optima (one or many) lies on the constraint boundaries. In an indefinite case, the Optimizer converges either to the saddle point, or to one of the optimum solutions on the constraint boundaries. If it converges to the saddle point, which might not be optimal, **LpStatusNum()** returns 65 (“objective changing too slowly”). The final point reached by the optimization depends on the initial starting point for the search, which can optionally be specified using the parameter **guess** to **QpDefine()**.

```
guess: Optional Number[Vars];
```

The **guess** parameter is only relevant if q is indefinite, otherwise the same end result is reached regardless of the starting point.

Because **QpDefine()** is totally array-abstractable, you can provide multiple guesses by dimensioning the argument to this parameter by an index other than **Vars**, with different starting points. In that case, multiple quadratic optimizations are solved, each at different starting points.

Convexity The set of constraints are said to be **convex** when the set of feasible solutions is a convex subset of all potential solutions. A convex subset is one in which for every two points in the set, all points on the line segment connecting them are also in the subset.

A quadratic constraint is convex in these cases:

- The matrix \hat{Q}_i is positive semi-definite and the constraint sense is \leq .
- The matrix \hat{Q}_i is negative semi-definite and the constraint sense is \geq .
- The constraint is linear.

When all constraints are convex, then the set of feasible solutions is convex. Quadratic programs can be solved efficiently when the set of feasible solutions is convex. Positive (or negative) semi-definiteness of q can be tested for using Analytica’s **EigenDecomp()** function. The matrix is positive (negative) semi-definite if all Eigenvalues are non-negative (non-positive). Note: If your q is not symmetric, use **EigenDecomp()** on $Q + Q^T$.

Common quadratic situations

Quadratic programs arise in several applications, one of the most common being portfolio optimization.

Portfolio allocation Assume there are n investments, each with an uncertain outcome. The investments are not independent. For example, two investments in the same sector can be influenced by similar market forces and thus be highly correlated. Other pairs of investments can be negatively correlated. A symmetric covariance matrix, q : can be used capture the pair-wise covariances between investments, as well as the variances of the individual investments (the diagonal elements of q). Letting each element of the vector \hat{x} be the fraction of the total portfolio allocated to investment, the variance of the complete portfolio is $\hat{x}^T Q \hat{x}$. As a result, various objective functions used in portfolio optimizations depends on the net variance of the portfolio, and leads to quadratic programs of the form shown in “Choosing the type of optimization” on page 21. Two such examples are

demonstrated in the example model **Asset Allocation.ana** found in the **Example Models\Optimizer Examples** directory.

When sample covariances are computed from historical data, and the number of time periods used is greater than the number of dimensions (e.g., the number of investments), the resulting Q matrix is guaranteed to be positive-definite. As discussed in the previous section, then property lends itself well to solution by quadratic programming.

Obtaining the solution

The **QpDefine()** function defines the quadratic program, but does not solve it. The optimum is solved for when **LpSolution()**, **LpStatusNum()**, **LpStatusText()**, or any of the other functions that use the solution are called.

The functions **LpSlack()**, **LpShadow()**, and **LpRhsSa()** are all available for quadratic programs (see the discussion for each of these in Chapter 3, “Linear Optimization”)

The **LpReducedCost()** function can also be called on a quadratic program.

Search control settings

In addition to settings found on linear programs, several additional settings apply to quadratically constrained problems. These can be altered from their default values using the two parameters to **QpDefine()** named **parameter** and **setting**. See Chapter 6, “Control Settings”, for details on using search control settings, and for descriptions of available settings.

Examples

The **Example Models/Optimization Examples** directory, installed with Analytica, contains an example model demonstrating quadratic optimization.


- **Asset Allocation.ana**: Portfolio optimization is a classic quadratic programming application. This example demonstrates four formulations of an asset allocation problem, two of which are quadratic programs.

Chapter 5

Nonlinear Optimization

This chapter shows you how to:

- Formulate a nonlinear optimization problem
- Obtain the solution for a nonlinear optimization problem
- Give hints to help the Optimizer
- Control the search



Nonlinear programs

A nonlinear program (NLP) is the most general formulation for an optimization. The objective and the constraints can be arbitrary functions of the decision variables, continuous or discontinuous. This generality comes at the price of longer computation times, and less precision than linear and quadratic programs (LP and QP). There is also the possibility with smooth NLPs that the Optimizer will return a local optimum that is not the global optimum solution. In general, it is hard to prove whether a solution is globally optimal or not. For these reasons, it is better to reformulate nonlinear problems as linear or quadratic when possible.

Linear and quadratic problems define the objective function as arrays of linear or quadratic coefficients. They pass these arrays as parameters to the Optimizer, which operates on them directly to find a solution without further interaction with the rest of the Analytica model. For nonlinear problems, the objective function is defined as an Analytica expression or variable that depends on the decision variables. In this case, the Optimizer repeatedly evaluates the objective function as it assigns different values to the decision variables in its search for a solution. It does the same with expressions passed to **Lhs**, the left-hand side of the constraints.


This approach imposes certain restrictions on array abstraction (support for Intelligent Arrays) for NLPs — for example, requiring the objective function to return a single (scalar) number. We devote a section of this chapter to showing how to work with these restrictions so that you can apply NLP optimization to create arrays of optimizations for models with uncertainty (samples indexed by **Run**) for parametric analysis, dynamic models over time, and other indexes.

The Optimizer has a variety of methods, including gradient-based search, branch-and-bound, and genetic algorithms, from which it chooses to try to suit the problem. In many cases, you can give it information about the problem that can help it choose the most appropriate methods, and so work faster and more reliably. Such hints include:

- The type of dependence — i.e., whether the objective or constraint functions vary linearly, smoothly, or discontinuously with each decision variable.
- The **gradient** and **Jacobian** expressions to compute the needed partial derivatives for the objective much faster at each search point.
- Control parameters to influence how the search is performed.

Problem formulation

The basic formulation for a nonlinear optimization is



where \hat{x} is a vector denoting the n -dimensional candidate solution. A nonlinear optimization problem is defined using the function **NlpDefine()**, shown here without optional parameters.

```

NlpDefine(Vars, Constraints: Optional Index;
          X: Variable;
          Obj, LHS: Optional Expression;
          RHS: Optional Numeric[Constraints])

```

- Vars** An index for the decision variable \mathbf{x} , below. This can be omitted when there is only one scalar decision variable.
- Constraints** An index for the constraints **LHS** and **RHS**, below. This is optional when there are fewer than two constraints.
- X** The decision variables, indexed by **Vars**. The parameter passed to \mathbf{x} must be the name (identifier) of a global, decision, or local variable, not an expression. As the NLP Optimizer searches for better solutions, it assigns new values to the decision variable, and computes the corresponding value of the objective.
- Obj** The objective to maximize or minimize, according to the setting of optional parameter **Maximize**. It can be a variable or an expression. It must depend on the decision variable \mathbf{x} directly or indirectly via other variables. When no objective is specified, the Optimizer stops when a feasible solution, satisfying all the constraints, is found.
- LHS** The *left-hand side (LHS)* of the constraints, indexed by **Constraints**, if the **Constraints** index is specified. You can omit **LHS** if you have no constraints. If you have a single constraint (and **Constraints** is omitted), the expression should evaluate to a scalar. If you have more than one constraint, then **LHS** should be indexed by **Constraints**. Each element of **LHS** can be an expression or a variable. They must depend on the decision variable \mathbf{x} , directly or indirectly.
- RHS** The *right-hand side (RHS)* of the constraints, indexed by **Constraints**, if the **Constraints** index is specified. You can omit this if you have no constraints. If you have a single constraint, **RHS** should evaluate to a single number, and if you have multiple constraints, each element of the array passed to **RHS** must evaluate to a single number. It must not depend on the decision variables. By default, feasible solutions are those in which **LHS** is less than or equal to the corresponding value of **RHS**. You can change this with the optional **Sense** parameter, described below.

Optimization without constraints

To solve for the unconstrained minimum of $\mathbf{f}(\mathbf{x})$ where \mathbf{x} is a scalar, only the decision variable \mathbf{x} and the objective need to be specified. For example, the following defines an NLP to find the minimum of $f(x) = \cosh(1 + \sinh(x))$.

```

var x:=0;
NlpDefine( X:x, Obj: cosh(1+sinh(x)) )

```

In this example, the decision variable \mathbf{x} is a local variable, and the objective expressed explicitly. Alternatively, the decision variable and objective can be variables in your model, as in this example.

```

NlpDefine( X:annualMaintExpense, Obj:costOfOwnership )

```

To minimize over several decision variables, \mathbf{x} is a vector, and the index of that vector is specified using the parameter **Vars**. The objective must be a variable or expression that depends on \mathbf{x} , but must evaluate to a scalar. For example, to find a vector of coefficient **Coeff**, indexed by \mathbf{k} , that minimize **errorMeasure**, the problem is formulated as follows.

```

NlpDefine( Vars: K, X: Coeff, Obj: errorMeasure )

```

Representing constraints

A system of equations and inequalities can be represented as a set of constraints without an objective. Discrete constraint satisfaction problems can also be encoded using integer-valued variables. When there is no objective, the Optimizer generally terminates when any feasible solution is located. Systems of nonlinear equations can be extremely difficult to solve, and it might be unrealistic to expect an Optimizer to find a solution just because it can be expressed, so a certain degree of realism is essential.

GoalSeek A *goal-seek functionality* seeks a value for a scalar x that causes scalar y to be equal to g , where y depends on x . Since x is scalar, the **Vars** index does not need to be specified, and because there is a single constraint, no **Constraints** index is necessary. When we use **GoalSeek**, equality is usually desired, so the **Sense** parameter is specified as = (otherwise it defaults to <=). Since any feasible solution is sought, no object needs to be specified. Thus, simple **GoalSeek** requires parameters **X**, **Lhs**, and **Rhs**.

```
NlpDefine( X:GrossIncome, Lhs:NetIncome, Sense:'', Rhs: 1M )
```

or

```
NlpDefine( X:Price, Lhs: Supply(Price)-Demand(Price), Sense:'',
Rhs:0)
```

or using a local variable

```
Var x:=0;
NlpDefine( X:x, Lhs:cosh(1+sinh(x)), sense='', rhs:1.001 )
```

Multiple constraints When there is more than one constraint, an index parameter, **Constraints**, must be specified. When **LHS** is evaluated during the optimization search, the result must be indexed by the **Constraints** index, and only by the **Constraints** index. It is often convenient to set up a variable that computes **LHS**, defined as a table indexed by **Constraints**, with a separate expression in each cell (the $g_i(\hat{x})$ in the basic formulation).

RHS should also evaluate to an array indexed by **Constraints**. However, unlike **LHS**, **RHS** is constant during the optimization search, so **RHS** should not depend on **X**.

Obtaining the solution

The same functions used to obtain the solution to LP and QP optimizations also work for NLP. These include **LpStatusNum()**, **LpStatusText()**, **LpOpt()**, **LpSolution()**, **LpSlack()**, **LpShadow()**, and **LpReducedCost()**.

For more, see Chapter 8, “Optimizer Function Reference,” on page 88.

Optional parameters for NLP

Every parameter of **NlpDefine()** except **X** is optional; however, either an objective or at least one constraint is required. In addition to the core parameters used for specifying the objective and constraints, the following optional parameters to **NlpDefine()** can also be specified.

Maximize By default, **NlpDefine()** defines a minimization problem. You should set the optional parameter

```
Maximize: Optional Boolean
```

to True when you wish to maximize the objective.

Sense By default, each constraint specifies that the left-hand side is less than or equal to the right-hand side. Using the optional **Sense** parameter, you can change the relationship between left-hand and right-hand sides:

Sense: Optional Text[Constraints]

- **<**, **<=**, or **L**: **LHS** is less-than or equal to **RHS**
- **>**, **>=**, or **G**: **LHS** is greater-than or equal to **RHS**
- **=** or **E**: **LHS** is equal to **RHS**

If you pass a single text value, such as **=**, to the **Sense** parameter, that sense will apply to all the constraints. If you want a different sense for each constraint, pass an array indexed by **Constraints**, with each cell containing its own text value **<=**, **>=**, **E**, etc.

Bounds As with LP and QP problems, you can define upper and lower bounds on each decision variable for an NLP problem using these optional parameters.

Lb, Ub: Optional Number[Vars]

If not explicitly specified, the Optimizer assumes bounds of **-INF** and **+INF**, i.e., that the decision is unbounded. If you pass a single number to either of these parameters, that bound applies to all decision variables. So, for example

NlpDefine(..., Lb:0,Ub:1 ...)

This specifies that all decision variables are in the range 0 to 1. If lower or upper bounds are different for each decision variable, pass them arrays of numbers indexed by **Vars**.

Array abstraction

NlpDefine() does not automatically array abstract over extra dimensions that appear in the results of evaluating the objective or constraint left-hand side expressions. When the objective function **Obj** is evaluated, it should be a single number with no extraneous indexes. Similarly, when **Lhs** is evaluated, it should contain only the index passed to the **Constraints** parameter. If you have multiple objectives, you must combine them (for example, as a weighted average) if they are all the criteria that you wish to optimize simultaneously. If instead you desired multiple optimization problems, for example, for each combination of scenarios, then you must tell **NlpDefine()** about these dimensions explicitly, otherwise an error results when **Obj** or **Lhs** is discovered to have extra dimensions.

Two parameters of **NlpDefine()** can be used in some situations to specify dimensions that must be iterated, in which a separate optimization should be conducted for each combination of elements in these extra dimensions.

Over : ... optional atomic
SetContext : ... optional Variable

Each of these are optional repeated parameters, allowing several indexes (in the case of **Over**) or several parametric variables (in the case of **SetContext**) to be specified.

The **Over** parameter specifies a list of indexes that should be abstracted over. For example, if we wish to run a separate optimization problem for each possible **Discount_rate**, and each possible **Initial_Investment** option, where **Discount_rate** and **Initial_Investment** are variables that can contain lists, we would specify

NlpDefine(..., Over: Discount_rate, Initial_investment)

If `Discount_rate` contains three elements and `Initial_investment` contains five elements, then 15 separate optimization problems are defined. No error results if `Discount_rate` or `Initial_investment` is not a list, which means that values that might be set to a list during a parametric exploration, but normally are not, can be listed here. In addition, if there is a result in your model that contains extra dimensions, the identifier of the result array can be listed, and each of the dimensions in that result are used. This fact is useful for defining an NLP that abstracts across the `Run` dimension in sample mode, but not in Mid mode, using for example

```
NlpDefine(..., Over: Uniform(0,1) )
```

The `Over` parameter does two things. It instructs Analytica to define multiple NLPs, and instructs `NlpDefine()` to pay attention to one particular slice of `Obj` or `Lhs` in each NLP when there are extra dimensions. It is important to note, however, that `Over` does not restrict how the rest of your model is evaluated when computing `Obj` and `Lhs`. If you rely only on the `Over` parameter, your evaluation of `Obj` and `Lhs` might compute far more than necessary, most of which gets ignored, resulting in slow evaluations during optimization.

The `SetContext` parameter provides a mechanism to restrict your computation to only the slice that is used by a particular NLP. For example, instead of computing `Obj` for every possible `Discount_rate` and then throwing away all results except the one corresponding to the current NLP, it is more efficient to set `Discount_rate` to a single value before running the optimization search, thus limiting the computation of `Obj` at each cycle to the single `Discount_rate` that is used. `SetContext` specifies a list of variables that is restricted to a single value for the duration of one particular optimization.

To use `SetContext`, the variable listed must contain a list-based domain attribute, containing the set of possible values for that variable. Each instance of NLP corresponds to one of those possible values, and when that instance runs, the variable is set to that single value. For example

```
Discount_rate := Choice(Self,0)
Domain of Discount_rate := [8%,10%,12%]

NlpDefine(..., SetContext: Discount_rate)
```

If the `Over` parameter is also specified, for example with other indexes, the parameter specified in `SetContext` should also be specified in `Over`. If the context variable evaluates to a single number, it is not set (hence, it works fine with `Choice()`, abstracting only when `All` is selected in the `Choice` pull-down menu).

The use of `Over` and `SetContext` allow array abstraction in some situations, but they are not always applicable, and not always the best way to achieve an array-abstractable optimization. In the `The Airline Example for NLP` below, a variety of other techniques are demonstrated.

Integer, binary, and mixed-integer programs

Like the LP and QP optimizers, the NLP optimizer can handle discrete decision variables — that is, integer or binary (Boolean) — as well as continuous values. Use the parameter

```
Ctype: Optional Text[Vars]
```

to specify the continuity type of each decision variable by providing one of the following text values for each variable:

- `c`: Continuous

- **i**: Integer
- **b**: Binary or Boolean
- **g**: Grouped integer

The nonlinear optimizer uses a genetic (or evolutionary) algorithm when discontinuous variables are present.

When the grouped integer parameter is used, each grouped integer variable belongs to one group, where each group contains at least two grouped integer variables. Within a group, a feasible solution assigns the integers $1..n$ to each variable, where n is the number of variables in the group, so that all variables have different values. When you have more than two groups, you should also specify the **group** parameter.

Group : Optional Number[Vars]

This parameter specifies the group number for each grouped integer variable. If **group** is omitted, all variables are assumed to belong to the same group. The **Traveling salesman.ana** example model demonstrates the use of grouped-integer variables (belonging to a single group).

Hard integers vs. soft integer constraints

The **Ctype** parameter can be used to indicate that a decision variable in any feasible solution is integer-valued. However, during the course of the search, the nonlinear optimizer can explore non-integer values for these variables while exploring the rate of change within the search space. If your model produces well-defined results for non-integer values, then this type of **soft-integer constraint** does not present a problem, and the Optimizer can benefit from being able to explore non-integer solutions.

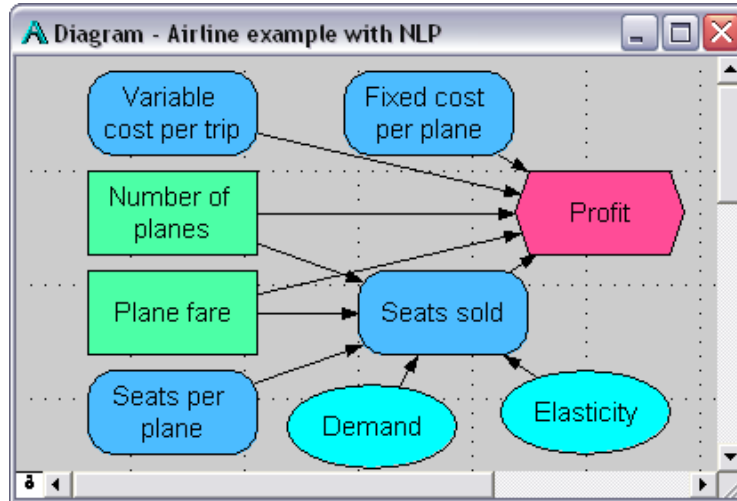
In some cases, however, you might need to enforce a **hard-integer constraint**. For example, a table lookup can encounter an error when a non-integer lookup is attempted. Or your model can encounter errors for other reasons when non-integer values are attempted. If you have hard-integers, you can either modify your model to return something sensible for non-integer values (for example, by linearly-extrapolating), or you should explicitly select the Evolutionary engine using this parameter.

Engine : "Evolutionary"

When using the Evolutionary engine, two control settings are relevant (these are set using the **parameter** and **setting** parameters to **NlpDefine()**). First, at each iteration, a new sample point is generated by mutating or combining two members of the population via cross-over. Then, that sample can be improved further using a local search before adding it to the population. For example, a gradient descent can be used to find the nearest local optima before adding it. **LocalSearch** controls whether any local searches are performed and by what method. By leaving **LocalSearch** at its default off position, non-integer solutions are not explored. If you do elect to use a local search, then the **FixNonSmooth** setting controls whether gradient information in local searches is limited to continuous variables. If you have hard-integer constraints with local search, then you should ensure that **FixNonSmooth** is indeed on. These parameter settings do not impact the GRG Nonlinear engine.

Airline example for NLP

Here we introduce the airline decision problem. We use this example in the rest of this chapter with eight cases that illustrate how to formulate problems for NLP, including situations in which parameters have extra indexes, for dealing with uncertainty, parametric analysis, and dynamic models over time. You can find this example in the **Example Models/Optimizer Examples/Airline NLP.ana**. It includes the eight different cases described below. Open the model in Analytica to see full details.



A small airline is trying to decide how many planes to lease and what fare to charge on a new route. It has two decision variables — `Num_planes`, the number of planes allocated for this route, and `Fare`, the price charged for trips on this route — and two chance variables — the `Base_demand` for seats (assuming the fare is \$200) and the `Elasticity1` of demand with respect to price.

```
Decision Num_planes := 2
Decision Fare := 200 ($/passenger trip)
Chance Base_demand :=
    Triangular(300K, 400K, 500K) (trips/year)
Chance Elasticity1 := Triangular(2, 3, 4)
```

Note: When you first load the model, `Num_planes` and `Fare` are not set to these single numbers. You can change them at this point to single numbers, or even to the following sequences to view `Profit` parametrically.

```
Decision Num_planes := 1..5
Decision Fare := 100..300
```

We assume that the demand is elastic with respect to changes in price, using a demand function that decreases as `Fare` increases at an exponential decay rate determined by `Elasticity1`. At a base fare of \$200, the demand is equal to the `Base_Demand`. We compute the actual `Seats_sold` as the lesser of the demand modified for price elasticity and the actual seats available (the product of the number of planes and annual `Seats_per_plane`).

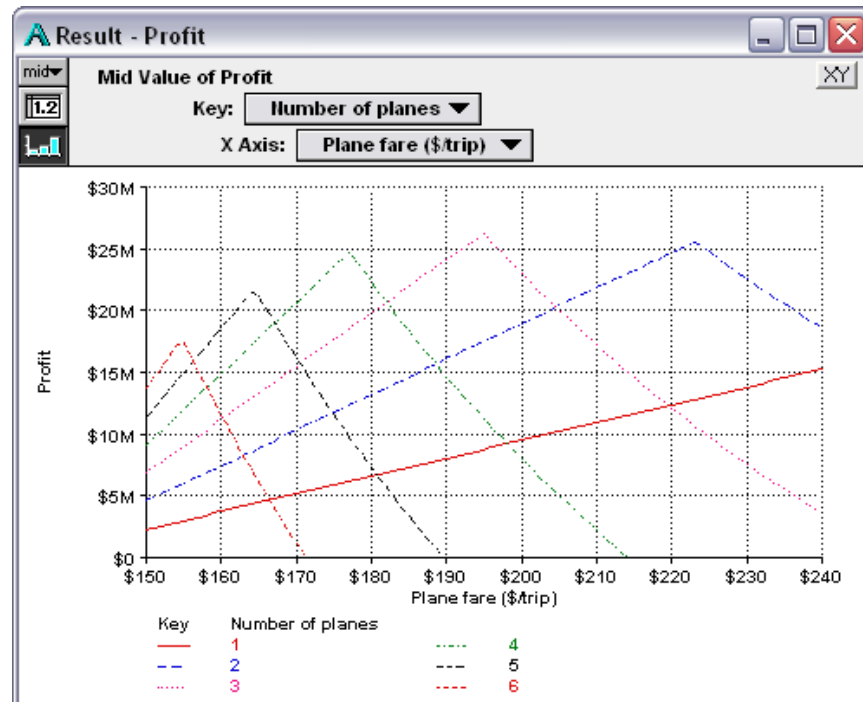
```
Variable Seats_per_plane := 200 * 360 * 2
Variable Seats_sold := Min([Base_demand *
    (Fare/200)^-Elasticity1,
    Num_planes * Seats_per_plane])
```

Finally, we model the objective variable `Profit` as the difference between revenues and costs, including `Fixed_cost`, the annualized fixed cost of leasing and operating each plane, and `var_cost`, the incremental cost for each new passenger.

```
Variable Fixed_cost := 12M ($/plane/year)
Variable Var_cost := 100 ($/passenger trip)
Objective Profit := Seats_sold*Fare
    - Seats_sold*Var_cost - Num_planes*Fixed_cost
```

This graph shows `Profit` as a function of the two decision variables, using the parametric approach to visualize the effects. Note that for each number of planes, 1 to 5, the

profit is a sharply peaked function of the fare. The optimum fare is at the highest peak, \$195 with 3 planes.



In this simple case, with only two decision variables, you can visualize the objective function and find the optimal values (or close) by parametric analysis. For more complex problems, the Optimizer is essential. Next we show how to apply that.

Reformulating the decision variables for NLP

We usually need to reformulate a decision problem, at least a little, to apply NLP. One reason is that `NlpDefine()` expects a single, array-valued decision variable for parameter **X**. So, if you want to apply NLP to optimize a model, like the airline example, whose decision variables are two or more separate Analytica variables, you need to combine these decisions into a single array-valued decision. If the model has n scalar decision variables, you should define a decision variable **Decisions** as a one-dimensional array with an index containing n elements. For the airline example, we define **Decisions** with two elements, corresponding to its two decisions, **Num_planes** and **Fare**.

```
Index Dvars := ['Number of planes', 'Plane fare']
Decision Decisions := Table(Dvars)(3, 200)
```

The values in the table are the initial values, prior to optimizing. We must now redefine the individual decision variables so that they obtain their values from the corresponding elements of **Decisions**.

```
Num_planes := Decisions[Dvars = 'Number of planes']
Fare := Decisions[Dvars = 'Plane fare']
```

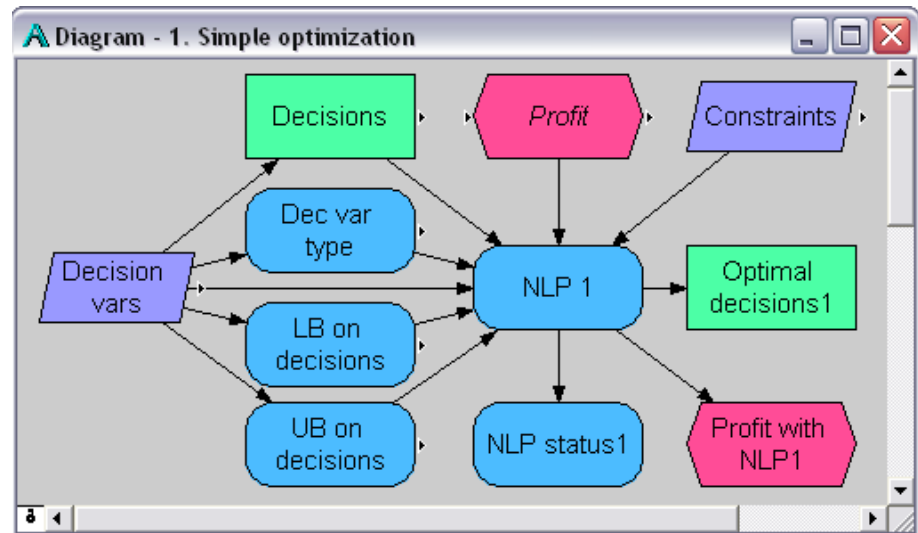
As the Optimizer searches for optimal values, it assigns successive new candidate solutions to **Decisions**, and gets the resulting value of **Profit**, which in turn gets its values from **Decisions**, via **Num_planes** and **Fare**.

If one or more of the original decision variables is an array, the new decision variable **Decisions** passed to **x** must still have only one dimension. Its size should be the sum of the sizes of all the original decision variables. Again, you should assign the current initial values of the original decision variables to the corresponding elements of **Decisions**. Then you redefine each original decision variable so that it gets each element

from the corresponding element of `Decisions`. See “Case 7. NLP with optimizations over time” for an example, where we add the `Time` dimension to `Num_planes` and `Fare`.

Case 1. Simple NLP optimization

We now complete the formulation of the NLP for the airline problem introduced above, creating a model that looks like this.



We need to specify the type of each decision — `I` (integer) for Number of planes, and `C` (continuous) for Fare — the lower and upper bounds for the two decisions, and the **Constraints** index.

```
Variable Dec_type := Table(Dvars)('I','C')
Variable Lb_decisions := Table(Dvars)(1, 100)
Variable Ub_decisions := Table(Dvars)(5, 300)
```

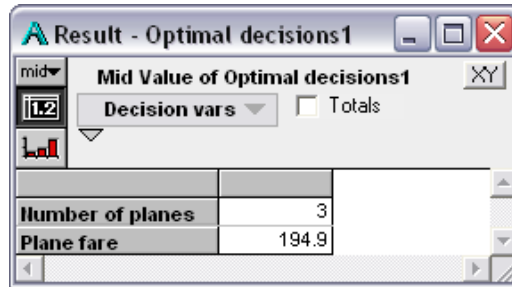
We can now define the NLP using these parameters.

```
Variable NLP_1 := NlpDefine(
  Vars: Dvars, X: Decisions, Ctype: Dec_type,
  LB: Lb_decisions, UB: Ub_decisions,
  Obj: Profit, Maximize: True )
```

Since we want the largest `Profit`, we set **Maximize** to `True`. Finally, we define the key results of the optimization: the optimal decisions, the profit with these decisions, and the status of the optimization.

```
Decision Optimal_decisions1 := LPSolution(Nlp_1)
Objective Profit_with_nlp1 := LPOpt(Nlp_1)
Variable Nlp_status1 := LPStatusText(Nlp_1)
```

When we display the result of any of these three variables, it performs the optimization. For example, `Optimal_decisions1`, gives this table (agreeing closely with the parametric analysis).



Intelligent arrays, array abstraction, and NLP

Unlike most other Analytica functions, including linear and quadratic optimization, nonlinear optimization does not fully support Intelligent Arrays — that is, it does not automatically generalize over extra dimensions for all parameters. Below we show how you can work around these restrictions to create and solve arrays of NLP problems, including handling uncertainty, parametric analysis, and dynamic optimization over time.

NLP's limitations are that the following required parameters must be dimensioned by the specified indexes *and no other indexes*:

- **X** must be indexed only by the index supplied to **Vars**.
- **Obj** must be scalar, a single number with no indexes.
- **Lhs** must be indexed by the index supplied to **Constraints**, or have no index.

Similarly, these optional parameters, if specified, must also be dimensioned by only the specified indexes:

- **gradient** must be indexed only by the index supplied to **Vars**.
- **Jacobian** must be indexed only by the indexes supplied to **Vars** and **Constraints**.

See page 62 for details on **gradient** and **Jacobian**.

Note that **NlpDefine()** *does* generalize fully over extra dimensions for all parameters other than the parameters listed above. But, for these parameters, it is up to you, the modeler, to make sure that they have only the required indexes. Otherwise it flags an error. Read on to see how to get around these limitations.

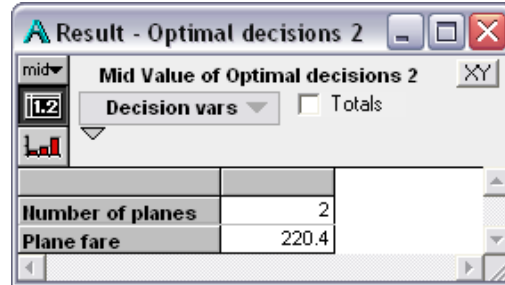
In many of the cases described below, there are two alternative approaches, one in which we encapsulate the NLP in a user-defined function, and a second using the **Set-Context** parameter. Both approaches usually entail certain changes to the underlying model, so there are tradeoffs between the two approaches, and in any particular problem you might find one more convenient than the other. The cases that follow demonstrate both alternatives when both are applicable.

Case 2. Maximize expected value: NLP with uncertainty

If you want to find the optimal decisions with an uncertain model, the most common approach is to define the objective as maximizing the *expected value* (i.e., mean) of the objective function — for example, maximizing the expected profit or the expected utility in a decision analysis formulation. For the Airline example, we define **NLP2**, which differs from **NLP1** only in that the objective takes the mean of the profit.

```
Variable NLP2 := NlpDefine(... ,
    X: Decisions, Obj: Mean(Profit), ...)
```

In this case, the objective is a single scalar number (i.e., the expected value). Although it is a function of an uncertain quantity, it is not itself uncertain. So you can apply **NlpDefine()** directly, and the restrictions on array abstraction mentioned above cause no problems. Note the results of doing the optimization using expected value are a bit different from the deterministic analysis, because the profit function is not symmetric.



Mid Value of Optimal decisions 2	
Number of planes	2
Plane fare	220.4

The same approach works if you want to maximize a statistic of the objective other than mean, such as to minimize the 1st percentile of an uncertain profit (loss), e.g., **Getfract(Profit, 1%)**. If there is uncertainty in the constraint functions, you can define the constraints using percentiles (using **Getfract()** or other statistical functions). For example, the constraint that the cumulative cashflow has a >95% chance of being nonnegative.

In these cases, you are trying to find the optimal decision now, *before* resolving the uncertainties that affect the objective or constraints. You can set the model to perform a single optimization and the result is a single optimal solution (set of decisions) and corresponding maximum expected value (or other statistic) of the objective. Given the optimal solution, you can then compute a probability distribution over the objective function to model the uncertainty over the value outcome.

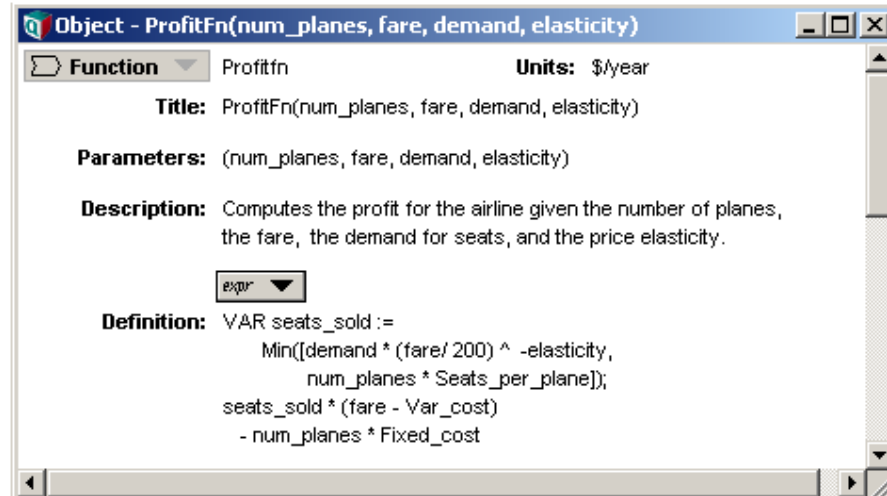
Case 3. NLP with uncertainty: probabilistic optimization

The second type of optimization under uncertainty is less common. In this case, the optimal decisions are made *after* resolving the uncertainty, and you want to compute probability distributions over what those optimal decisions are now while still uncertain. This is sometimes known as *preposterior* analysis because the optimization is performed *a posteriori* — after the uncertainty is resolved — but you are performing the analysis now, *before* the uncertainty is resolved. (Not to be confused with *preposterous* analysis, which we try to avoid.) This situation requires a sample of optimizations to be performed. It results in a random sample of optimal decisions, and a sample of corresponding values of the objective for each solution.

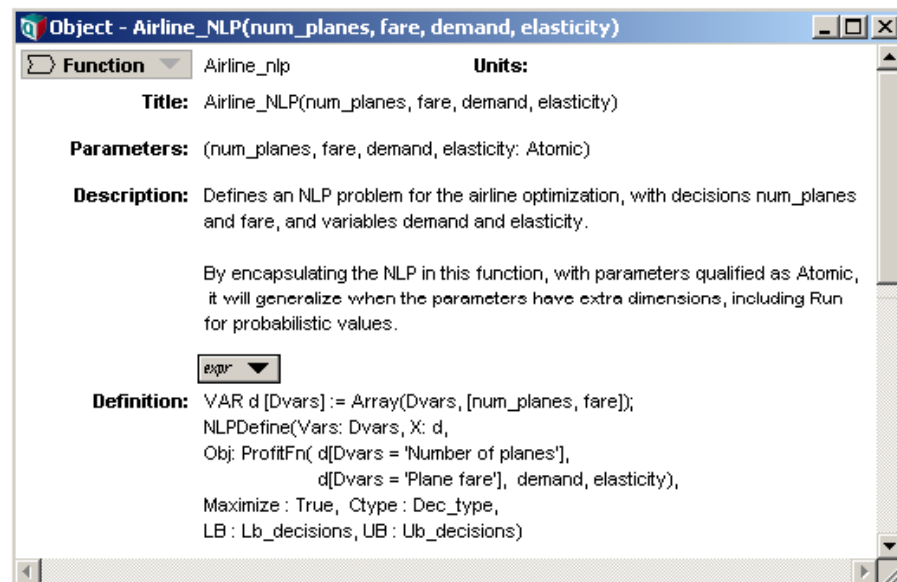
You might try simply to compute a probabilistic value of the optimal decision in case 1, from **NLP_1**, by selecting a uncertain view, e.g., **Sample**, in the result for **Optimal_decisions1**, shown previously. However, this generates the error “*The expression for the objective function in NlpDefine must evaluate to a single numeric value during optimization...*” This is because the objective, **Profit**, is no longer a single number at each iteration, but rather a Monte Carlo sample of numbers.

Instead, we need to create an NLP that abstracts over the chance variables, so that the **Run** index does not cause problems for **NlpDefine()**. Two techniques for accomplishing this are demonstrated here. The first encapsulates the NLP in a user-defined function, and the second uses the **SetContext** parameter.

UDF encapsulation For convenience, we define two functions, first **ProfitFn()** that encapsulates the objective **Profit** as a function of the decisions and chance variables as parameters. This function replicates **Profit** in the simple airline model (see below.)



Then we define a function **Airline_nlp()** that defines an NLP using **ProfitFn()** that we just defined for the objective.

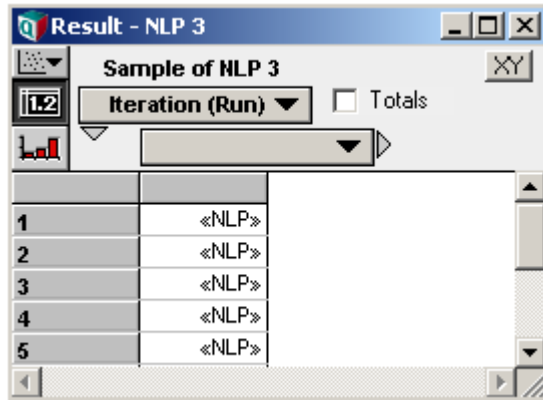


Airline_nlp() qualifies its parameters as *atomic*. This means that if the actual parameters are arrays, indexed by **Run** or anything else, it reduces them to scalar values and calls the function multiple times, once for each combination of scalar values. On each call, it passes scalar parameters to **ProfitFn()**, so that the objective passed to **Obj** in **NlpDefine()** is scalar, as required. In this way, it restores the Intelligent Array behavior that NLP otherwise lacks.

We now define a variable using this function.

```
Variable Nlp_3 := Airline_nlp( Num_planes, Fare,
                             Demand, Elasticity1)
```

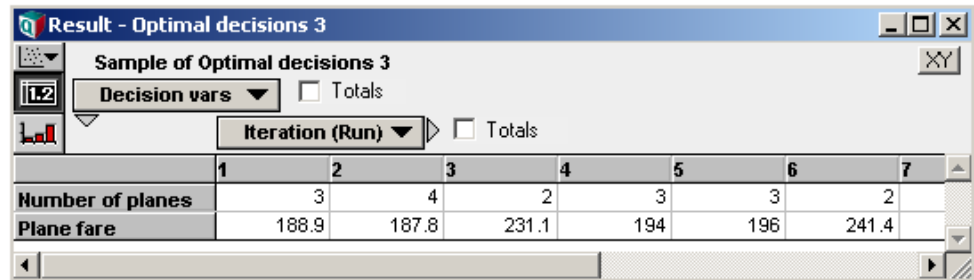
If you show the result of this variable in a sample view (with **SampleSize** set to 5 for rapid execution), it shows a sample of NLP problems.



When we show the resulting optimal decisions

```
Decision Optimal_decisions_3:=LPSolution(Nlp_3)
```

it evaluates each sample of the NLP and generate a corresponding sample of optimal decisions.



This computation involves doing **samplesize** optimizations. So, it could take a long time if the NLP problem is difficult and the sample size is large.

Use of SetContext

The uncertainty in the objective comes from the uncertainty in **Seats_per_plane**, which is uncertain as a result of its uncertain inputs, **Base_demand** and **Elasticity1**. When we use the value of **Seats_per_plane**, we can slice out only the sample corresponding to the current optimization. To do this, we set up a context variable.

```
Run_context := if IsSampleEvalMode then Run else 1
domain of Run_context := Index Run
```

Next we modify our **NlpDefine()** call, specifying **Run_context** as a context for the optimization.

```
Nlp_3b := NlpDefine(..., SetContext:Run_context )
```

When **Nlp_3b** is viewed in sample mode, a separate NLP appears for each element of the **Run** index. When the NLP corresponding to **Run=3** is evaluating, **Run_context** is set to **1** during that search. Finally, we must slice out a single element of the sample for **Seats_per_plane**. We do this by modifying the definition of **Profit** (in the example, a copy of **Profit** as been placed in this module named **Profit_in_context**).

```
Profit_in_context :=
  Seats_sold[Run=Run_context] * (Fare - Var_cost)
  - Num_planes * Fixed_cost
```

The choice was made here to slice on the current **RunContext** at **Profit**, rather than within **Seats_per_plane**, where each of **Base_demand** and **Elasticity1** could have been sliced. A speed advantage can be obtained by using your context as late in the

computation as possible. In this way, `Seats_sold` is computed once, and does not have to be re-evaluated at each point in the search space.

Case 4. NLP and parametric analysis

What if you want to examine how optimal decisions vary as you change one or more input parameters, such as `Demand`? (See Chapter 3, “Analyzing Model Behavior,” in the *Analytica User Guide* for more on parametric analysis.) In this case, the variables you treat parametrically have multiple values, so you cannot apply `NlpDefine()` to them directly. However, the function `Airline_nlp()` that we just defined comes in handy again. Suppose we define this.

```
Variable Demand_param := [200K, 400K, 600K, 800K, 1M]
Variable NLP_4 := Airline_nlp( Num_planes, Fare,
    Demand_param, Elasticity1)
Decision Optimal_decisions4 := LPSolution(Nlp_4)
```

Because `Airline_nlp()` qualifies its parameters as *atomic*, `NLP_4` generates an array of NLPs, one for each value of `Demand_param`. The result for `Optimal_decisions4` shows corresponding optimal values for each value of `Demand_param`.

	200K	400K	600K	800K	1M
Number of planes	1	2	2	2	4
Plane fare	223.1	223.1	255.4	281.1	240.4

Note how the optimal number of planes increases from 1 to 4 as the demand increases, and the optimal fare varies nonmonotonically.

Once again, the parametric analysis can also be accomplished using the `SetContext` parameter.

```
NlpDefine(Vars: Dvars, X: Decisions, Ctype : Dec_type,
    Obj: Profit_4b, Maximize : True,
    Lb : Lb_decisions, Ub : Ub_decisions,
    SetContext: Demand_param)
```

Case 5. NLP over time using NPV

The most common formulations for optimization over time involve finding a set of decisions to optimize an objective that measures overall performance over multiple time periods, such as the net present value (NPV). In these cases, the objective function returns a single number that aggregates over the time periods, so it poses no problem for direct application of `NlpDefine()`.

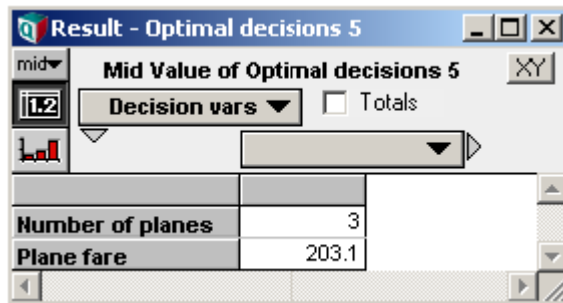
Consider the airline example again. We add an uncertain annual compound growth in demand, define `Time` for years from 2005 to 2010, and compute the resulting `Demand_by_time`.

```
Chance Demand_growth_rate := Triangular(0%, 10%, 20%)
Time := 2005 .. 2010
Variable Demand_by_year := Dynamic(Base_demand,
    Self[Time-1] * (1 + Demand_growth))
```

We now define the objective of the NLP using mean of the net present value (NPV).

```
Variable Nlp_5 :=
  NlpDefine(Vars: Dvars, X: Decisions, Ctype:Dec_type,
    Obj: Mean(NPV(Discount_rate,
      ProfitFn( Num_planes, Fare,
        Demand_by_year, Elasticity1), Time)),
    Maximize:true,
    Lb:Lb_decisions, Ub:Ub_decisions)
```

This causes no array-abstraction issues for the objective since the mean of the NPV is a scalar. Notice that we are finding a single optimal value for the decisions, **Num_planes** and **Fare**, for all time periods: We are assuming that these decisions stay the same over the six years. Because of the growth in demand, the optimal number of planes is three, larger than before.

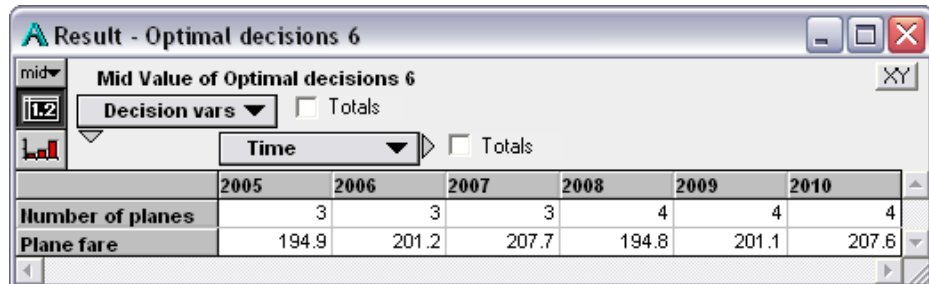


Case 6. Optimize for each year

What if you want to change the decisions **Num_planes** and **Fare** in each time period? One approach is to perform a separate optimization in each time period. This formulation models a process in which the decisions are made at the start of each time period to maximize profit for that time period. In this case, the decisions and objectives (and possibly constraints) are indexed by time. Again, the function **Airline_nlp()**, which we defined earlier, comes in handy.

```
Variable Nlp_6 := Airline_nlp( Num_Planes, Fare,
  Demand_by_year, Elasticity1 )
Decision Optimal_decisions4 := LPSolution(Nlp_4)
```

Since **Demand_by_year** is indexed by **Time**, **Airline_nlp()** creates an array of NLPs over time. **optimal_decisions4** is then computed separately for each year.



Demand_by_year has uncertainty. Since we have formed the objective given the value for **Demand_by_year**, in the probabilistic result we have actually created a separate optimization problem for each Monte Carlo sample at each time period. In terms of the problem, this is saying that we are uncertain today exactly what the demand will be in

future years, but we will know the demand each year before we make our decision. Hence, we are array abstracting across two dimensions, **Run** and **Time**.

In the alternative formulation, **SetContext** can also be used to array abstract across both **Run** and **Time**, as demonstrated in **NLP_6b**. In this case, we need to introduce two context variables, **Run_context** (already introduced in Case 3) and **Time_context**. When computing the objective profit, we need to restrict our inputs to these contexts.

```
Profit_in_context6 := ProfitFn( Num_planes,
    Demand_by_year[Time=Time_context, Run=Run_context],
    Elasticity1[Run=Run_context] )
```

Because we are array-abstracting across two contexts, both appear in the **SetContext** parameter.

```
NlpDefine(Vars: Dvars, X: Decisions, Ctype : Dec_type,
    Obj: Profit_in_context6, Maximize : True,
    LB : Lb_decisions, UB : Ub_decisions,
    SetContext: Run_Context, Time_Context)
```

	2005	2006	2007	2008	2009	2010
Mean Value of Profit by year 6b	26.63M	29.3M	32.31M	35.59M	39.16M	43.27M

Case 7. NLP with optimizations over time

If there are interactions between decisions in different years, you might want to find the decisions in each year that collectively maximize the NPV (or another objective that aggregates over time). In this case, we want to perform only one optimization, but with an expanded set of decisions, that comprises both decisions over all time period. With 2 decisions in each of 6 time periods, we define a **Decisions** vector of 12 elements. Note that **Decisions** must be a one-dimensional vector with 12 elements, not a two-dimensional table with 2 by 6 elements.

In this case, we choose to create a single table with the decision settings initial values, **Ctype**, lower bounds, and upper bounds for all 12 elements.

We derive the **Decisions_by_time** as a slice of this table.

```
Decision Decisions_by_time :=
    Decision_params[Decision_settings='Initial']
```

See the module in the example model for details of how the NLP is defined. We have a single optimization problem here with no array abstraction considerations. Here are sample results for the optimal decisions.

Result - Optimal decisions 7

mid

Mid Value of Optimal decisions 7

Decision vars by time

Totals

Num planes 2005	3
Num planes 2006	3
Num planes 2007	3
Num planes 2008	3
Num planes 2009	4
Num planes 2010	4
Fare 2005	192.4
Fare 2006	198.3
Fare 2007	204.1
Fare 2008	210.3
Fare 2009	198.5
Fare 2010	203.3

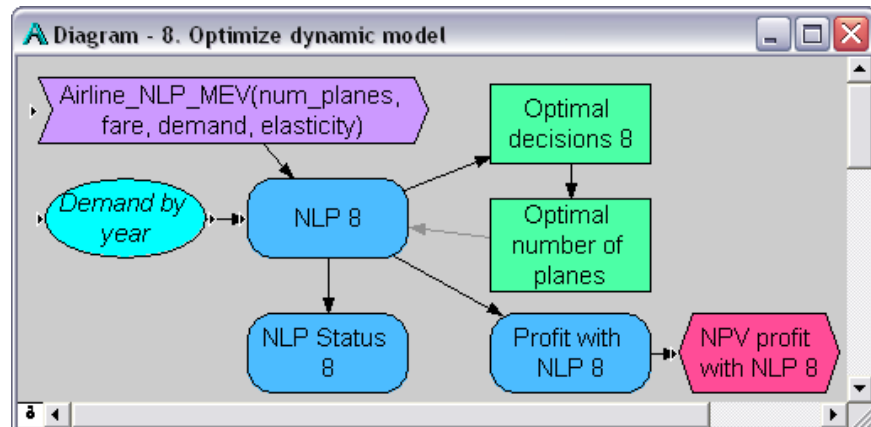
The time to perform NLP optimization typically increases superlinearly with the number of decision variables, so this approach can become time consuming if you have many decision variables and time periods. In general, it takes longer than **Case 6. Optimize for each year**, which is linear in the number of time periods.

Case 8. NLP with a dynamic model

The previous three cases are dynamic in the sense that the model changes over time. However, they do not need to use the **Dynamic()** function explicitly because the decisions in each year do not depend on the results of the previous year. In this final case, the optimization at each time step depends on the results of the optimization at the previous time step, so we *must* use **Dynamic()**. We assume that the planes are on long-term leases — we can lease more planes each year, but cannot decrease them because of the lease agreement. This means that the lower bound on the number of planes decision in each period is the value of the optimal number of planes computed in the previous time period.

```
Variable Nlp_8 := Dynamic(
  Airline_nlp_mev(Num_planes, Fare, Demand,
    Elasticity1),
  Airline_nlp_mev(Optimal_num_planes[Time-1],
    Fare, Demand_by_year, Elasticity1 ))
Decision Optimal_decisions_8 := LPSolution(Nlp_8)
Decision Optimal_num_planes :=
  Optimal_decisions_8[Dvars = 'Number of planes']
```

Note that this creates a dynamic loop, with the time lagged dependence shown in the diagram in gray.



As a general principle, when embedding an NLP inside a dynamic loop only a local variable should ever be used for your decision variable, x , and `SetContext` should not be used. If you were to use a global variable for x or `SetContext`, each optimization would end up invalidating the results of optimizations at other time steps in the dynamic loop. However, you can include your decision variables as global decision variables in your dynamic loop, defining them using `LpSolution()`, but using a local variable for the decision vector passed to parameter x of `NlpDefine()`. This is demonstrated in `nlp_8b`. The expressions for `Obj` and `Lhs` can contain time offsets (in this case they do not). Keep in mind that `self[Time-1]` in one of these expressions refers to an NLP object, not the previous value of the `Obj` or `Lhs` expression. To use the optimum objective value in the previous time step, you would use `LpOpt(self[Time-1])`, for example. A portion of the definition of `Nlp_8b` is as follows.

```
Variable Nlp_8b := Dynamic(
  Var d[DVars] := Array(Dvars,[num_planes, fare]) do
    NlpDefine( Vars: DVars, X:d, ... ),
  Var d[DVars] := Array[Dvars,[num_planes8b[Time-1], fare]) do
    NlpDefine( Vars: DVars, X:d, ...
      Lb: if DVars='Number of planes' Then Num_planes8b[Time-1]
          Else Lb_decisions )
)
Variable Num_planes8b := LpSolution(Nlp_8b)[DVars='Number of
planes']
```

The critical thing to notice here is that the decision variables appearing in the dynamic loop, in this case `Num_planes8b`, are defined using `LpSolution`. The expressions at `Time=t` can make use of the optimal solutions at `Time=t-1` in the definition of `NlpDefine()`, in its parameters, or in the expressions. The decision vector passed to x must be a local variable, and `setContext` cannot be used.

Summary of array abstraction for NLP

These airline problem Cases 1 to 8 shown above illustrate ways to reformulate a problem for NLP to deal with various issues of array abstraction and Intelligent Arrays. Case 1 shows how to combine multiple scalar decisions into a single vector of decisions, as needed for `NlpDefine()`. Case 2 shows that you require no special reformulation for NLP to maximize expected value (or other statistical function of an uncertain objective), since the objective is a scalar, even if the underlying model has uncertainty. Simi-

larly, Cases 5 and 7 illustrate that maximizing the net present value (or another objective that aggregates over time) produces a scalar value for the objective, so you can apply **NlpDefine()** directly. Case 7 shows how to assemble array-valued decisions into a single vector of decisions.

In the other cases, the objective is intrinsically an array of values, indexed by **Run** for uncertainty in Case 3, by a parametric analysis (**Demand**) in Case 4, and by **Time** in Cases 6 and 8. We handle these cases in a similar way — we encapsulate the **NlpDefine()** in a function whose parameters are qualified as **Atomic**, so that each call to **NlpDefine()** is made with the required inputs and hence the objective passed to **X** as scalar. The result of calling these functions is an array of NLPs. Functions of this result, such as the optimal decisions, **LPSolution()**, status, **LPStatusText()**, and optimal value, **LPOpt()**, are therefore similarly indexed by these extra dimensions. For Cases 3, 4, and 6, we've also shown an alternative approach, in which we use the **SetContext** parameter to **NlpDefine()** restrict the model to a restricted slice while the optimization search takes place. In Case 8, with dynamic dependence, we see that an NLP can be embedded in a dynamic loop, referring to previous time points, as long as a local variable is used for the **X** parameter and **SetContext** is not used. If dimensions other than **Time** must be abstracted over, with each NLP depending on the solution to an earlier one, then encapsulating the NLP in a user-defined function is the only viable option.

For more details, look at the example Analytica file that contains these cases in **Example Models/Optimizer Examples/Airline NLP.ana**.

These examples show how to deal with array abstraction for the objective **Obj**. The same approach works for the other parameters that are repeatedly evaluated during an optimization, i.e., **LHS**, **gradient**, and **Jacobian**. *All other parameters array-abstract automatically.*

Solving systems of equations

Solving a system of nonlinear equations is a special case of a nonlinear program. The set of solutions is the set of feasible points. The nonlinear optimizer can be used to find a solution to a system of equations by encoding the system of equations as the set of constraints, using a sense of =. You can omit the objective function (the **Obj** parameter) if you simply care about finding any solution, or you can use the objective to express a preference among solutions when the system of equations has, or might have, multiple solutions.

Other examples

If you haven't already, you might find it useful to follow the steps in the "Quick Start" section for creating a nonlinear optimization model (see "A nonlinear program" on page 14).

The **Example Models/Optimizer Examples** directory, installed with Analytica, contains several models demonstrating nonlinear optimization. These models include:

- **Asset Allocation.ana**: A classic portfolio optimization problem, formulated in four ways. One formulation uses a linear objective with a quadratic constraint, which qualifies as a nonlinear problem. Another formulation maximizes expected utility, thus demonstrating the use of stochastic simulation within a nonlinear optimization. The other two formulations are quadratic programs.
- **NLP with Jacobian.ana**: A very simple nonlinear program demonstrates the use of a **gradient** and **Jacobian**, as well as the use of a local variable for **x**.

- **Optimal can dimensions.ana**: The example is the one used in Chapter 1, “Quick Start,” of this manual. The problem is to find the dimensions for a cylindrical can to hold a given volume using the minimum surface area.
- **Solve using NLP.ana**: A very simple example of using the nonlinear optimizer to solve a nonlinear system of equations.
- **Problems with local optima.ana**: Demonstrates several techniques for overcoming the problem of local optima in nonlinear optimizations.
- **Traveling salesman.ana**: Demonstrates the use of grouped-integer decision variables.

Giving hints to help the Optimizer

The Optimizer tries to identify characteristics of your NLP problem so that it can choose the most efficient and reliable algorithms. In some cases, you can improve its performance by telling it things about the problem that it might not be able to figure out on its own.

Type of dependence If the Optimizer knows that the objective has smooth nonlinear dependence on some or all of the decision variables, it can use much faster gradient-based algorithms than in the general case that allows discontinuous functions. You can provide this information using these two optional parameters to **NlpDefine()**.

```
objN1: Optional Text[Vars]
lhsN1: Optional Text[Vars, Constraints]
```

You should provide each of these parameters with one of these text values:

- **L**: Linear or no dependence
- **Q**: Quadratic dependence
- **N**: Smooth nonlinear dependence
- **D**: Discontinuous

You can provide a single text value to each parameter, e.g., **N**, to specify the same type of dependence for all decision variables and, to **lhsN1**, for all constraints. Otherwise, if the type of dependence varies by variables and constraints, you will probably create a variable defined as an edit table indexed by **Vars** and **Constraints**, to specify each dependency type.

When the objective has linear, quadratic, or smooth nonlinear dependence on continuous decision variables, the Optimizer uses an efficient gradient-based search method. If it knows that the dependence is linear (and so has constant derivative), or quadratic (and so has a constant second derivative) it can drastically speed the search by reducing the number of re-evaluations of the objective. If one or more decision variables are discontinuous, the Optimizer uses a genetic (evolutionary) algorithm, in which multiple candidate solutions are maintained, and the search is performed by mutating and recombining members of the population based on a fitness metric.

If you do not indicate the type of dependence, the Optimizer defaults to smooth nonlinear, and the GRG Nonlinear engine is used. If any discontinuous dependence is indicated, the evolutionary algorithm is selected. In some cases, you might find the Evolutionary engine performs better even on smooth nonlinear problems, in which case you can force use of the Evolutionary engine using the parameter:

```
Engine: "Evolutionary"
```


Gradient and Jacobian functions

If the decision variables are all continuous, you can speed up the Optimizer considerably if you can give it an analytical expression for the gradient of the objective function and/or the Jacobian of the constraint left-hand sides. The gradient and Jacobian enable the Optimizer to avoid most re-evaluations of the objective and LHS expressions, respectively, which it uses estimate the partial derivatives based on small changes to each decision variable.

The **gradient** of the objective function is a vector indexed by **Vars**, where each element is the partial derivative.

$$\frac{\partial}{\partial x_i} f(\hat{x})$$

In this equation, $f(\hat{x})$ is the objective function.

The **Jacobian** of the left-hand side of the constraints is a matrix, indexed by **Vars** and **Constraints**, where each element is the partial derivative.

$$\frac{\partial}{\partial x_i} g_j(\hat{x})$$

In this equation, $g_j(\hat{x})$ is the left-hand side of constraint j .

The **gradient** and **Jacobian** parameters accept an Analytica variable or expression, which should depend on **X**, directly or indirectly. The Optimizer evaluates these parameters deterministically and repeatedly at each step of the search process. Assuming **X** is indexed only by **Vars**, the gradient must be indexed only by **Vars**, and the Jacobian must be indexed only by **Vars** and **Constraints**. See “Intelligent arrays, array abstraction, and NLP” on page 51 for information on coping with these restrictions.

It is important for your gradient and Jacobian expressions to be correct, otherwise you mislead the Optimizer and it might move away from the optimum. Debugging a Jacobian expression can be challenging. However, you can check whether the Jacobian is correct using the optional parameter, **DerivMethod**, to **NlpDefine()**.

```
NlpDefine(..., DerivMethod: "check", ...)
```

When **DerivMethod** is set to **check**, the Optimizer compares the supplied Jacobian expression, with the Jacobian that it estimates using finite differencing. If they are not within a small difference, the optimization stops with **LpStatusNum()** = 67 (“error in evaluating problem functions”). When you have confirmed the supplied Jacobian is correct, remember to reset **DerivMethod** to **Jacobian** so that the Optimizer reaps the benefits of not having to estimate the Jacobian itself at each search point.

Initial guess

If you know the approximate region that contains the optimal solution, you can speed the Optimizer by giving it an initial solution in that region. You specify this starting solution as an array indexed by **Vars** for the optional parameter **guess**.

```
guess: Optional Number[Vars]
```

If you do not provide this parameter, and if you provide a global variable (as opposed to a local variable) for **x**, the Optimizer uses the current value of **x** as its starting solution.

Dealing with local optima

A difficult problem common to many hard nonlinear optimization problems is the existence of local optima. When a local optima is reached, it is impossible for the Optimizer

to know where, or even if any, better solutions exist. If you think you are having problems with local optima, there are several settings that can be manipulated.

MultiStart If you have a continuous nonlinear problem, enabling the **MultiStart** setting of the GRG Nonlinear engine is often the quickest and easiest recourse. This can be tried quickly simply by adding the following parameters to **NlpDefine()**.

```
NlpDefine(..., parameter:"MultiStart", setting:1 )
```

MultiStart requires more search time, as it tries multiple starting points. When using **MultiStart**, you should also specify finite lower and upper variable bounds using the **Lb** and **Ub** parameters to **NlpDefine()**. Narrow bounds produce better results.

If turning **MultiStart** on alone is inadequate, you can further enhance exploration by enabling a topographic search via the **TopoSearch** setting, which improves the selection of starting points, and by increasing the number of starting points with the **Population Size** setting. See “Specifying settings” on page 68.

Engine selection Two nonlinear optimization engines come with Analytica Optimizer:

- GRG Nonlinear: A gradient-descent search.
- Evolutionary: A genetic-algorithm search.

If you have purchased other add-on engines, additional options might be available. To explicitly select the engine to be used, include the **Engine** parameter to **NlpDefine()**:

```
Engine : Optional Text
```

If you have indicated that your problem is discontinuous, the GRG engine cannot be used.

By default, the Evolutionary engine does not use gradient information. However, if the **LocalSearch** setting is on, then it optimizes sample points before adding them to the population using various techniques including gradient-based search.

To view the list of possible engines installed, evaluate the following Analytica expression.

```
SolverInfo( "AvailEngines" )
```


If your problem is highly discontinuous or contains many local optima, then the Evolutionary engine is a better choice. If your problem is relatively smooth with relatively few local optima, then the GRG Nonlinear engine is likely to obtain results more quickly.

Chapter 6

Control Settings

This chapter shows you how to:

- Specify Optimizer engine settings to **LpDefine()**, **QpDefine()**, and **NlpDefine()**
- Determine what settings are available for each engine, defaults, and possible range
- Determine size capacities for installed engines
- Control termination criteria during optimization
- Select search algorithms
- Specify numeric precision



Controlling the search

The optimization engine exposes several settings that you can change to influence how the search for the optimum proceeds and when it terminates. The specific collection of available settings is a function of which engine is used to solve the optimization, so that if you install and use an add-on engine, other than the engine that comes standard with Analytica Optimizer, the possible settings might be different. The **SolverInfo()** function can be used to view all available settings, the range of possible values, their defaults, and their current values for a problem. Settings can be changed for a particular problem by specifying values for the **parameter** and **settings** parameters to **LpDefine()**, **QpDefine()**, or **NlpDefine()**. The first subsection below describes how you specify and view settings, while the subsequent sub-sections detail particular settings used by engines that come standard with Analytica Optimizer.

Selecting the optimization engine

Four optimization engines come standard with Analytica Optimizer:

- **LP/Quadratic**: Used for LPs and QPs with linear constraints.
- **SOCP Barrier**: QPs with quadratic constraints.
- **GRG Nonlinear**: Smooth NLPs: A gradient-descent search.
- **Evolutionary**: NLPs: A genetic-algorithm search.

If you have purchased other add-on engines, other options might also be available to you. You can obtain a full list of installed engines by evaluating the following Analytica expression.

```
SolverInfo( "AvailEngines" )
```

To explicitly select the engine to be used, include the **Engine** parameter to **LpDefine()**, **QpDefine()**, or **NlpDefine()**.

```
Engine : Optional Text
```

For example:

```
NlpDefine( ..., Engine: "Evolutionary" )
```

The following engines can be used with each function:

- **LpDefine()**: LP/Quadratic, SOCP Barrier, GRG Nonlinear, Evolutionary
- **QpDefine()**: LP/Quadratic (but only if constraints are linear, i.e., parameter **LhsQ** not specified), SOCP Barrier, GRG Nonlinear, Evolutionary
- **NlpDefine()**: GRG Nonlinear (but only if any **objNI** or **lhsNI** has been marked **D** for discontinuous), Evolutionary

If you do not specify the engine, Analytica selects an appropriate engine based on the function you used to define your problem, and the properties of the problem that you specified. However, if the engine does not perform satisfactorily on that problem, you might obtain better results with a different engine.

To determine what engine is actually used on a problem, evaluate this Analytica expression.

```
SolverInfo( "Engine", prob )
```

Here **prob** is the object returned by **LpDefine()**, **QpDefine()**, or **NlpDefine()**.

The LP/Quadratic engine uses a dual simplex method combined with branch-and-bound for mixed-integer constraints, with a variety of integer cut-set procedures. This is generally the engine of choice for LPs and mixed-integer LPs. However, for hard mixed-integer LPs, since the Evolutionary engine uses a very different approach, that engine might be worth trying.

SOCP uses a second-order cone programming technique designed specifically for quadratically constrained convex problems. The GRG Nonlinear engine is often a good alternative for problems formulated with **QpDefine()**, especially if the constraints end up being non-convex.

For nonlinear problems, if your problem is highly discontinuous or contains many local optima, then the Evolutionary engine is a better choice. If your problem is relatively smooth with relatively few local optima, then the GRG Nonlinear engine is likely to obtain results more quickly, and if gradients and Jacobians can be analytically computed, it is likely to be dramatically faster. If non-integer values cannot be explored during the intermediate steps of a search, the Evolutionary engine should be used.

By default, the Evolutionary engine does not use gradient information. However, if the **LocalSearch** setting is on, then it optimizes sample points before adding them to the population using various techniques including gradient-based search.

Examining engine capabilities

Information about an installed optimization such as the maximum number of variables or constraints allowed can be accessed using this expression.

```
SolverInfo( Item: "<item>", Engine: <engineName> )
```

engineName is a value returned by **SolverInfo("AvailEngines")**, and **<item>** is one of **MaxVars**, **MaxIntVars**, **MaxConstraints**, or **MaxVarBounds**. These return a result indexed by a local index named **ProblemType**, having elements [LP, QP, QCP, CVX, NLP, NSP], described in the table below.

Element	Description
QP	quadratic objective, linear constraints
QCP	quadratic with convex quadratic constraints
CVX	non-convex quadratic
NLP	smooth nonlinear
NSP	non-smooth nonlinear

For example:

```
Index Engines := SolverInfo("AvailEngines");
SolverInfo( ["Maxvars", "MaxIntVars", "MaxConstraints",
"MaxVarBounds"],
Engine: Engines) [.ProblemType='LP']
```

This returns the result shown below.

	Maxvars	MaxIntVars	MaxConstraints	MaxVarBounds
LP/Quadratic	8000	2000	8000	16K
SOCP Barrier	2000	2000	8000	4000
GRG Nonlinear	500	500	250	1000
Evolutionary	500	500	250	1000

Specifying settings

If you want to change the value for a single control setting, you can specify values for two optional parameters, **parameter** and **setting**, to **LpDefine()**, **QpDefine()**, or **NlpDefine()**, providing the name of the setting to **parameter**, and the value to **setting**. For example, if you want to set the **Scaling** parameter to 1, you would modify your call to **LpDefine()** as follows.

```
LpDefine( .., Parameter: "Scaling", Setting: 1 )
```

To alter more than one control setting, you need to supply arrays to these parameters. The arrays passed to **parameter** and **setting** should have a single common index. If the index of the array passed to **setting** is a list of labels, where the index labels contain the name of each control setting, then you only need to include the **setting** parameter.

It is often convenient to specify control settings in a self-indexed edit table. The following steps illustrate this:

1. Drag a variable node to your diagram, title it **My Lp Settings**.
2. In the definition pane, set the definition type to **Table**.
3. In the **Index Chooser** dialog, select **My Lp Settings (Self)** as the table index.
4. Click the row heading cell, and change **Item 1** to **Scaling**.
5. With the row header still selected, press *down-arrow* to add a row.
6. Change the second row header cell to **MaxTime**.
7. Enter **1** into the first table body cell.
8. Enter **30** into the second body table cell.

Scaling	1
MaxTime	30

9. In your call to **LpDefine()**, insert a setting parameter as follows.

```
LpDefine( ..., setting: My_lp_settings )
```

The Optimizer scales parameters and terminates after 30 seconds if the optimum has not been found. A self-indexed table set up in this fashion makes it easy to adjust multiple control settings if the need arises.

Note: *In Analytica Optimizer 3.1, control settings were specified as optional parameters to **LpDefine()**, **QpDefine()**, and **NlpDefine()**. These legacy parameters are still supported for backward compatibility; however, use of the **setting** parameter is recommended. This change reflects a change in Frontline's architecture, and more readily generalizes to other add-on engines and future Optimizer engine extensions.*

Examining available settings

When you have defined an optimization problem using **LpDefine()**, **QpDefine()**, or **NlpDefine()**, the following function returns the set of control settings used by this engine.

```
SolverInfo("Setting", Lp: myLp )
```

Replace **myLp** with the name of the variable holding the result from **LpDefine()**.

You can also access the range of allowed values for each setting, as well as the default value, using **SolverInfo()**. For this, you need to know the name of the Optimizer engine used on your problem. For linear programs, this is always **LP/Quadratic** unless you have installed an add-on engine. To obtain the name of the engine used in the general case, use this expression.

```
SolverInfo("Engine", Lp: myLp )
```

Tip For quadratic and nonlinear problems, to be certain you get the correct engine, evaluate **SolverInfo("Engine", Lp:myLp)** after you've attempted to find a solution — after **LpSolution()**, **LpStatusText()**, or **LpOpt()** has been evaluated. If you have not specified the **engine** parameter explicitly, the Optimizer might change to a different, non-default engine based on the properties of your problem.

Using the name of the engine, the range (min/max) of possible values for each setting, and the default value, can be obtained using this.

```
SolverInfo( ["MinSetting", "MaxSetting", "Defaults"],
Engine: "LP/Quadratic")
```

Termination controls

Iterations Specifies the maximum number of iterations (pivots) by the simplex algorithm during the optimization. If this is exceeded, **LpStatusNum()** returns 3 (Iterates limit reached. Indicates an early exit of the algorithm.). Maximum number of generations in Evolutionary solver. Maximum number of gradient descent steps by GRG Nonlinear.

Default: no limit

MaxTime Maximum number of seconds the Optimizer spends on the problem. If exceeded, **LpStatusNum()** is 10 (Time out status. Returned when the maximum allowed time has been exceeded. Indicates an early exit of the algorithm.).

Default: no limit

- MaxTimeNoImp** The maximum number of seconds that the Optimizer continues without finding any improvement in the best solution.
Default: 30 seconds
Allowed range: positive
- IntTolerance** In a MIP optimization, if the branch-and-bound algorithm can determine that the best solution found so far is within this relative tolerance of the true optimal, it terminates the search and return the best solution found so far. The bound is relative, meaning a value of 10% guarantees a solution within 10% of the optimal. Often, the branch-and-bound algorithm quickly locates a nearly optimal solution, but then spends a large amount of refining its best solution to the true optimum. Specifying a non-zero gap tolerance can eliminate this additional search, thus in some cases drastically reducing computation time. The gap is computed as the absolute value of the difference between the best solution so far, and the best bound on the optimum, divided by the best bound on the optimum. With zero gap (default), the search continues until the entire search space is eliminated so that the global optimum is reached.
Default: 0%
Allowed range: 0 to 1
- Convergence** The evolutionary solver stops with status “*Solver has converged to the current solution*” when nearly all members in the current population have very similar fitness values. This stopping criteria is satisfied when 99% of the population members all have fitness values within *Convergence* tolerance of each other.
The fitness value is a combination of the objective function value and a penalty for constraints still violated. If you think the evolutionary solver is terminating too quickly, you can make this tolerance smaller, but you might also want to increase **MutationRate** or **PopulationSize** in order to increase the diversity of trial solutions.
Default: 10^{-4}
Allowed range: 0 or 1
- Tolerance** If the relative (i.e., percentage) improvement observed during the previous **MaxTimeNoImp** seconds does not exceed this value, then evolutionary solver terminates. See **MaxTimeNoImp**.
Default: 0
Allowed range: 0 to 1
- MaxTimeNoImp** Controls the amount of time (in seconds) that the evolutionary solver is willing to spend without making any significant progress. If the relative improvement during this time has not exceeded the setting specified by **Tolerance**, it terminates with status (Solver cannot improve the current solution) or (Solver could not find a feasible solution).
Default: 10^{-5}
Allowed range: 10^{-9} to 10^{-4}
- MaxSubProblems** Maximum number of subproblems explored by Evolutionary algorithm before terminating.
Default: no limit
Allowed range: positive
- MaxFeasibleSolutions** The maximum number of feasible solutions found by the Evolutionary algorithm before terminating.

Default: no limit

Allowed range: positive

MaxIntegerSols The Optimizer terminates after this many feasible solutions have been found by the branch and bound algorithm.

Default: no limit

Allowed range: positive

Algorithm selection

Preprocessing

Scaling When this is True, the Optimizer attempts to rescale decision variables and constraints internally for the simplex algorithm, which usually leads to be reliable results and fewer iterations. A poorly scaled model, in which values of the objective, constraints, or intermediate results differ by several orders of magnitude, can result in numeric instabilities within the Optimizer when scaling is turned off, due to the effects of finite precision computer arithmetic.

Default: False

Allowed range: 0 or 1

Presolve When this is True, the LP/Quadratic engine performs a presolve step to detect singleton rows and columns, remove fixed variables and redundant constraints, and tighten bounds, prior to applying the simplex method.

Default: True

Allowed range: 0 or 1

Engine: LP/Quadratic

PreProcess Turns on or off all integer pre-processing (on by default).

Default: 1

Allowed range: 0 or 1

Engine: LP/Quadratic

StrongBranching This setting applies to integer and mixed-integer problems. When this is on, the Optimizer estimates the impact of branching on each integer variable of the objective function prior to beginning the branch and bound search. It does this by performing a few iterations of the dual simplex method after fixing each variable. This “experiment” provides the search with an estimate of which integer variables are likely to be most effective choices during the branch and bound search. Although the time spent in this estimation process can be moderately expensive, the cost is often regained many times over through a reduction in the number of branch-and-bound iterations that must be explored to find an optimal integer solution.

Default: 1

Allowed range: 0 or 1

Engine: LP/Quadratic

Debugging

SolveWithout When this is True, any integer (**Ctype**) constraints are ignored, and the continuous, and the continuous version of the problem is solved instead. The effect is the same as changing the **Ctype** parameter to **c**, but can be more convenient in some cases when debugging.

Default: True

Allowed range: 0 or 1

IISBounds Determines whether variable bounds should be included in the infeasibility search conducted by **LpFindIIS()** or **LpWriteIIS()**. When set to 1, only a subset of constraints along the **Constraints** index is considered. When set to 0, variable bounds can be eliminated in order to find an IIS with a greater number of constraints. This parameter is only used by **FindIIS()** when the second optional parameter, **newLp**, is True. When **newLp** is True, **FindIIS()** returns a new LP object, from which you can use **SolverInfo()** to access the list of constraints and list of variable bounds present in the IIS. When **newLp** is False, since only a subset of the **Constraints** index is returned, **LpFindIIS()** relaxes only constraints, leaving variable bounds in tact.

Default: 0

Allowed range: 0 or 1

Numeric estimation

Derivatives The **Derivatives** setting controls how derivatives are computed. These values are possible:

- **1 = forward:** This is the default if **Jacobian** and **gradient** parameters are not supplied. The Optimizer estimates derivatives using forward differencing, i.e.,

$$\frac{\partial}{\partial(x)} \approx \frac{f(x + \Delta) - f(x)}{\Delta}$$
- **2 = central:** The Optimizer estimates derivatives using central differencing, i.e.,

$$\frac{\partial}{\partial x} \approx \frac{f(x + \Delta) - f(x - \Delta)}{2\Delta}$$
- **3 = jacobian:** The Optimizer computes derivatives using the supplied **Jacobian** and **gradient** expressions. This is the default if these are supplied.
- **4 = check:** The Optimizer computes derivatives using the supplied Jacobian expression and also estimates the Jacobian using finite differencing. If they don't agree to within a small tolerance, the optimization aborts with **LpStatusNum() = 67** ("error in evaluating problem functions"). This option is useful for testing whether the Jacobian is accurate.

StepSize The step size used to estimate derivatives numerically. This is the Δ value in the estimates listed in the preceding **Derivatives** description.

Default: 10^{-6}

Allowed range: 10^{-9} to 10^{-4}

SearchOption Controls how the gradient-based search determines the next point to jump to during search:

- **0 = Newton:** Uses a quasi-Newton method, maintaining an approximate Hessian matrix for the reduced gradient function.
- **1 = Conjugate-gradient:** Use a conjugate gradient method, which does not require the Hessian.

Default: 0

Allowed range: 0 or 1

Estimates The **Estimates** setting controls the method used to estimate the initial values for the basic decision variables at the beginning of each one-dimensional line search:

- **0 = linear:** Uses linear-extrapolation from the line tangent to the reduced objective function.
- **1 = quadratic:** Extrapolates to the extrema of a quadratic fitted to the reduced objective at its current point.

Default: 0

Allowed range: 0 or 1

RecognizeLinear When set to 1, the Optimizer attempts to detect automatically decision variables that influence the objective and constraints in a linear fashion. It can then save time by pre-computing partial derivatives for these variables for the rest of the search. This aggressive strategy can create problems when a dependence changes dramatically throughout the search space, particularly when a decision variable is near linear around the starting point, but the gradient changes elsewhere in the search space. When the solution is reached, the Optimizer recomputes the derivatives and verifies them against the assumed values. If they do not agree, the status text *“The linearity conditions required by this solver engine are not satisfied”* is returned.

Engine: GRG Nonlinear

Default: 0 (select default)

Allowed range: 0 or 1

SOCP barrier search

In addition to the many search control settings available of linear programs (covered in the previous chapter), a few additional settings can be used to control the search when solving quadratically constrained problems using the SOCP Barrier engine.

These parameters are set using the **parameter** and **settings** parameters to **QpDefine()**, as described for **LpDefine()** in the previous chapter.

SearchDirection Controls the search direction on each iteration of the SOCP Barrier engine. The Power class method is a technique with the long-step barrier algorithm leading to a polynomial complexity. The dual scaling method uses HKM (Helmberg, Kojima, and Monteiro) dual scaling in which a Newton direction is found from the linearization of a symmetrized version of the optimality conditions. Either of these can be further modified by a predictor-corrector term.

Default: 0 (off)

Allowed range: 1 = Power class, 2 = Power class with predictor-corrector, 3 = dual scaling, or 4 = dual scaling with predictor-corrector.

Engine: SOCP Barrier

PowerIndex This parameter is used to select a particular search direction when the **SearchDirection** is set to 1 or 2.

Default: 1

Allowed range: non-negative integer

Engine: SOCP Barrier

- StepSizeFactor** The relative step size (between 0 and 1) that the SOCP Barrier engine can take towards the constraint boundary at each iteration.
Default: 0.99
Allowed range: 0.00 to 0.99
Engine: SOCP Barrier
- GapTolerance** The SOCP Barrier Solver uses a primal-dual method that computes new objective values for the primal problem and the dual problem at each iteration. When the gap or difference between these two objective values is less than the gap tolerance, the SOCP Barrier Solver stops and declares the current solution optimal.
Engine: SOCP Barrier
Default: 10^{-6}
Allowed range: 0 to 1
- FeasibilityTolerance** The SOCP Barrier engine considers a solution feasible when the constraints are satisfied to within this relative tolerance.
Engine: SOCP Barrier
Default: 10^{-6}
Allowed range: 0 to 1

Evolutionary search controls

- PopulationSize** Controls the population size of candidate solutions maintained by the Evolutionary engine, or the number of starting points for **MultiStart** in the GRG Nonlinear engine. **MultiStart** has a minimum population size of 10. If you specify 0, or any number smaller than 10, then the number of starting points used is 10 times the number of decision variables, but no more than 200.
Engine: GRG Nonlinear, Evolutionary
Default: 0 (automatic)
Allowed range: 0, or integer ≥ 10
- MutationRate** The probability that the Evolutionary Optimizer engine, on one of its major iterations, will attempt to generate a new point by “mutating” or altering one or more decision variable values of a current point in the population of candidate solutions.
Engine: Evolutionary
Default: 0.075
Allowed range: 0 to 1
- ExtinctionRate** This determines how often the Evolutionary engine throws out its entire population, except for the very best candidate solutions, and starts over from scratch.
Engine: Evolutionary
Default: 0.5
Allowed range: 0 to 1
- RandomSeed** Both engines use a pseudo-random component in their search for an optima. Thus, the final result can differ each time an optimization of the exact same problem is performed. By setting the random seed, you can ensure that the same sequence of pseudo-random numbers is used, so that the same result obtains every time the same problem is

re-evaluated. If you do not specify the random seed, Analytica uses its internal random seed, so that when you first load a model and evaluate results in a fixed order, you get a predictable result. Setting **RandomSeed** to 0 causes the pseudo-random generated to be seeded using the system clock. Any positive value sets the initial seed to a fixed number.

Engine: GRG Nonlinear, Evolutionary

Default: (use Analytica's random seed)

Allowed range: non-negative integer

Feasibility When set to 1, the Evolutionary engine throws out all infeasible points, and keeps only feasible points in its population. When set to 0, it accepts feasible points in the population with a high penalty in the fitness score, which tends to be useful when it has a hard time finding feasible points.

Default: 0

Allowed range: 0 or 1

LocalSearch Selects the local search strategy employed by the Evolutionary engine. In one step, or generation, of the algorithm, a possible mutation and a crossover occur, followed by a local search in some cases, followed by elimination of unfit members of the population. This parameter controls the method used for this local search. The decision for whether to apply a local search at a given generation is determined by two tests. First, the objective value for the starting point must exceed a certain threshold, and second, the point must be sufficiently far from any already identified local extrema. The threshold is based on the best objective found so far, but is adjusted dynamically as the search proceeds. The distance to local optima threshold is based on distance travelled previous times the local optima was reached.

There is a computational trade-off between the amount of time spent in local searches, versus the time spent in more global searches. The value of local searches depends on the nature of your problem. Roughly speaking, the **Randomized** method is the least expensive and the **gradient** method tends to be the most expensive (i.e., with more time devoted to local searches rather than global search).

Engine: Evolutionary

Default: 0

Allowed range: 0 to 3

1 = Randomized Local Search: Generates a small number of new trial points in the vicinity of the just-discovered "best" solution. Improved points are accepted into the population.

2 = Deterministic Pattern Search: Uses a deterministic "pattern search" method to seek improved points in the vicinity of the just-discovered "best" solution. Does not make use of the gradient, and so is effective for non-smooth functions.

3 = Gradient Local Search: Uses a quasi-Newton gradient descent search to locate an improved point to add to the population.

FixNonSmooth Determines how non-smooth variables (see "Type of dependence" on page 61 for information on **NlpDefine()** parameters **objNI='D'** and **lhsNI='D'**) are handled during the local search step. If set, then only linear and nonlinear smooth variables are allowed to vary during the local search. Because gradients often exist at most points, even for discontinuous variables, leaving this off can still yield useful information in spite of the occasional invalid gradient.

Engine: Evolutionary

Default: 0

Allowed range: 0 or 1

Mixed-integer controls

Integer branch and bound

IntCutoff If you can correctly bound the objective function value for the optimal solution in advance, this can drastically reduce the computation time for MIP problems, since the branch-and-bound algorithm to prune entire branches from the search space without having to explore them at all. For a maximization problem, specify a lower bound, and for a minimization problem, specify an upper bound. If you specify this parameter, you need to be sure that there is an integer solution with an objective value at least this good, otherwise the Optimizer might skip over, and thus never find, an optimal integer solution.

Default: no bounding

UseDual When True, the LP/Quadratic engine uses the *dual simplex method*, starting from an advanced basis, to solve subproblems generated by the branch-and-bound method. When False, it uses the *primal simplex method* to solve subproblems. Use of dual simplex often speeds up the solution of mixed-integer problems.

The subproblems of an integer programming problem are based on the relaxation of the problem, but have additional or tighter bounds on the variables. The solution of the relaxation (or of a more direct “parent” of the current sub problem) provides an “advanced basis” which can be used as a starting point for solving the current subproblem, potentially in fewer iterations. This basis might not be primal feasible due to the additional or tighter bounds on the variables, but it is always dual feasible. Because of this, the dual simplex method is usually faster than the primal simplex method when starting from an advanced basis.

Default: 2

Allowed range:

1 = Primal

2 = Dual

ProbingFeasibility Probing is a pre-processing step during which the solver attempts to deduce the values for certain binary integer variables based on the settings of others, prior to actually solving a subproblem. While solving a mixed-integer problem, probing can be performed on each subproblem before running a constrained simplex. As branch-and-bound fixes one variable to a specific binary value, this can cause the values for other binary variables to become determined. In some cases, probing can identify infeasible subproblems even before solving them. In certain types of constraint satisfaction problems, probing can reduce the number of subproblems by orders of magnitude.

Default: 0

Allowed range: 0 or 1

BoundsImprovement This strategy attempts to tighten bounds on variables that are not 0-1 or binary variables, based on values that have been derived for binary variables, before subproblems are solved.

Default: 0

Allowed range: 0 or 1

OptimalityFixing This strategy attempts to fix the values of binary integer variables before each subproblem is solved, based on the signs of coefficients in the objective and constraints. As with **BoundsImprovement** and **ProbingFeasibility**, this can result in faster pruning of branches by the branch-and-bound search; however, *in some cases optimality fixing can yield incorrect results*. Specifically, optimality fixing creates incorrect results when the set of inequalities imply an equality constraint. Here is an example.

```
Lhs[ Constraints=1 ] = Lhs[ Constraints=2 ]
```

In this example, these are the constraints for all **Vars**.

```
Sum( Lhs[Constraint=1] * X, Vars ) <= 10
Sum( Lhs[Constraint=2] * X, Vars ) >= 10
```

This implies an =10 constraint. You must also watch out for more subtle implied equalities, such as where it is possible to deduce the value of a variable from the inequalities. Such equalities must be represented explicitly as equalities for **OptimalityFixing** to work correctly.

Default: 0

Allowed range: 0 or 1

PrimalHeuristic This strategy attempts to discover a feasible integer solution early in the branch-and-bound process by using a heuristic method. The specific heuristic used by the LP simplex solver is one that has been found to be quite effective in the “local search” literature, especially on 0-1 integer programming problems, but which not guaranteed to succeed in all cases in finding a feasible integer solution. If the heuristic method succeeds, branch-and-bound starts with a big advantage, allowing it to prune branches early. If the heuristic method fails, branch and bound begins as it normally would, but with no special advantage, and the time spent with the heuristic method is wasted.

Default: 0

Allowed range: 0 or 1

LocalHeur, RoundingHeur, LocalTree These strategies look for possible integer solutions in the vicinity of known integer solution using a local heuristic (“local search heuristic” or “rounding heuristic”), adjusting the values of individual integer variables. As with the **PrimalHeuristic**, finding an integer solution can help improve bounds used by the search, and thus prune off portions of the search tree.

Engine: LP/Quadratic

Default: 0

Allowed range: 0 or 1

FeasibilityPump An incumbent finding heuristic used by branch-and-bound to find good incumbents quickly.

Engine: LP/Quadratic

Default: 1

Allowed range: 0 or 1

GreedyCover Another incumbent finding heuristic used by branch-and-bound to find good incumbents quickly.

Engine: LP/Quadratic

Default: 0

Allowed range: 0 or 1

Cut generation control

Cut generation options are available for the LP simplex method and is used when solving integer or mixed-integer LP problems.

A cut is an automatically generated constraint that “cuts off” some portion of the feasible region of an LP subproblem without eliminating any possible integer solutions. Many different cut methods are available each of which are capable of identifying different forms of constraints among integer variables that can be leveraged to quickly reduce the feasible set, and thus prune the branch-and-bound search tree. However, each of these methods requires a certain amount of work to identify cut opportunities, so that when opportunities are not identified, that effort can be wasted. The defaults are set in ways that represent a reasonable trade-off for most problems, but for hard integer problems, you can experiment with these to find the best settings for your own problem. You might find that some methods are more effective than others on your particular problem.

MaxRootCutPasses Controls the maximum number of cut passes carried out immediately after the first LP relaxation is solved. This has an effect only if one of the cut method options is on. If this is set to a value of -1, the number of passes is determined automatically. The setting **MaxTreeCutPasses** is used for all iterations after the first.

Engine: LP/Quadratic

Default: -1 (automatically determined)

Allowed range: -1 or more

MaxTreeCutPasses Controls the maximum number of cut passes carried out at each step of the solution process with the exception of the first cycle. This setting is used only if at least one cut method is on. Each time a cut is added to a problem, this can produce further opportunities for additional cuts, hence cuts can continue to be added until no more cuts are possible, or until this maximum bound is reached.

Engine: LP/Quadratic

Default: 10

Allowed range: 0 or more

GomoryCuts Gomory cuts are generated by examining the inverse basis of the optimum solution to a previous solved LP relaxation subproblem. The technique is sensitive to numeric rounding errors, so when used, it is important that your problem is well-scaled. It is recommended that you set the **Scaling** settings to 1 when using Gomory cuts.

Engine: LP/Quadratic

Default: 0

Allowed range: 0 or 1

MaxGomoryCuts This is the maximum Gomory cuts that should be introduced into a given subproblem.

Default: 20

Allowed range: non-negative

GomoryPasses The number of passes to make over a given subproblem looking for possible Gomory cuts. Each time you add a cut, this can present opportunities for new cuts. It is actually possible to solve an LP/MIP problem simply by making continual Gomory passes until the problem is solved, but typically this is less efficient than branch and bound. However, that can be different for different problems.

Default: 1

- Allowed range:** non-negative
- KnapsackCuts** Knapsack cuts are only used with grouped-integer variables (whereas Gomory cuts can be used with any integer variable type). These are also called *lifted cover inequalities*. This setting controls whether knapsack cuts are used.
- Engine:** LP/Quadratic
- Default:** 0
- Allowed range:** 0 or 1
- MaxKnapsackCuts** The maximum number of knapsack cuts to introduce into a given subproblem.
- Default:** 20
- Allowed range:** non-negative
- KnapsackPasses** The number of passes the solver should make over a given subproblem, looking for knapsack cuts.
- Default:** 1
- Allowed range:** non-negative
- ProbingCuts** Controls whether probing cuts are generated. Probing involves setting certain binary integer variables to 0 or 1 and deriving values for other binary integer variables, or tightening bounds on the constraints.
- Engine:** LP/Quadratic
- Default:** 1
- Allowed range:** 0 or 1
- OddHoleCuts** Controls whether *odd hole cuts* (also called *odd cycle cuts*) are generated. This uses a method due to Grotschel, Lovasz, and Schrijver that apply only to constraints that are sums of binary variables.
- Engine:** LP/Quadratic
- Default:** 0
- Allowed range:** 0 or 1
- MirCuts, TwoMirCuts** Mixed-integer rounding cuts and two mixed-integer rounding cuts.
- Engine:** LP/Quadratic
- Default:** 0
- Allowed range:** 0 or 1
- RedSplitCuts** Reduce and split cuts are a variant of Gomory cuts.
- Engine:** LP/Quadratic
- Default:** 0
- Allowed range:** 0 or 1
- SOSCuts** Special ordered sets (SOS) refer to constraints consisting of a sum of binary variables equal to 1. These arise common in certain types of problems. In these constraints, in any feasible solution exactly one of the variables in the constraint must be 1, and all the others zero, such that only n permutations need to be considered, rather than 2^n .
- Engine:** LP/Quadratic
- Default:** 0

- FlowCoverCuts** Controls whether flow cover cuts are used.
Engine: LP/Quadratic
Default: 0
Allowed range: 0 or 1
- CliqueCuts** Controls whether clique cuts can be used, using a method due to Hoffman and Padberg. Both row clique cuts and start clique cuts are generated.
Engine: LP/Quadratic
Default: 0
Allowed range: 0 or 1
- RoundingCuts** A rounding cut is an inequality over all integer variables formed by removing any continuous variables, dividing through by the greatest common denominator of the coefficients, and rounding down the right-hand side.
Engine: LP/Quadratic
Default: 0
Allowed range: 0 or 1
- LiftAndCoverCuts** Lift and cover cuts are fairly expensive to compute, but when they can be generated, they are often very effective in cutting off portions of the LP feasible region, improving the speed of the solution process.
Engine: LP/Quadratic
Default: 0
Allowed range: 0 or 1

Coping with local optima

- MultiStart** When turned on, the GRG engine restarts at multiple starting points, following the gradient from each to its corresponding local optima. Starting points are selected randomly between the specified lower and upper variable bounds, and clustered using a method called multi-level single linkage. The solver selects a representative point from each cluster, and then continues to successively smaller clusters based on the likelihood of capturing undiscovered local optima. Best results are obtained from **MultiStart** when your variable upper and lower bounds are finite with as narrow range as possible. If finite bounds are not specified, you must set **RequireBounds** to 0. **PopulationSize** controls the number of starting points. **TopoSearch** can be set for a more sophisticated method of selecting starting points.
Engine: GRG Nonlinear
Default: 0 (off)
Allowed range: 0 or 1
- RequireBounds** When **MultiStart** is used to select random starting positions, points between the bounds specified for each variable are sampled. If finite bounds on some variables are not specified, then **MultiStart** can still be used, but is likely to be less effective because starting value must be selected from an infinite range, which is unlikely to cover all possible starting points, and thus is unlikely to find all the local optima. When **RequireBounds** is on, as it is by default, an error results if you have not specified finite bounds

on variables and have selected the **MultiStart** method, so as to remind you to specify bounds. If you really intend to use Multistart without finite bounds on the variables, you must explicitly set **RequireBounds** to 0.

When using the Evolutionary engine, finite bounds are also important in order to ensure a appropriate sampling for an initial population. Although it can still function without bounds, the infinite range that must be explored can dramatically slow down amount required to find a solution, and thus it is recommended that you always specify finite upper and lower bounds when using the Evolutionary engine. If **RequireBounds** is 1 (the default) when no bounds are specified, an error is reported in order to encourage the use of bounds.

Engine: GRG Nonlinear, Evolutionary

Default: 1 (on)

Allowed range: 0 or 1

TopoSearch Only used when **MultiStart** is 1. When set to 1, the **MultiStart** method uses a topographic search method that fits a topographic surface to all previously sampled starting points in order to estimate the location of hills and valleys in the search space. It then uses this information to find a better starting points. Estimating topography takes more computing time, but in some problems that can be more than offset from the improvements in each GRG search.

Engine: GRG Nonlinear

Default: 0 (off)

Allowed range: 0 or 1

Numeric tolerance and precision

ReducedTol The optimal or reduced cost tolerance. The simplex method looks for a variable to enter the basis that has a negative reduced cost. Decision variables whose reduced cost is less than the negative of this tolerance are candidates for entering the basis during the simplex search.

Default: 10^{-5}

Allowed range: 10^{-9} to 10^{-4}

PivotTol During the simplex algorithm, elements in the solution matrix must have an absolute value greater than this value to be candidates for pivoting.

Default: 10^{-5}

Allowed range: 10^{-9} to 10^{-4}

Precision This value specifies how closely the calculated values on the left-hand side of constraints must match the right-hand sides in order for the constraint to be satisfied. Because of the finite precision arithmetic, a left-hand side that would ideally evaluate to 7.0 might compute as 6.9999999. With a precision of 10^{-6} , the constraint $A1 \geq 7$ would be considered satisfied in this case.

Default: 10^{-6}

Allowed range: 10^{-9} to 10^{-4}

PrimalTolerance The maximum amount by which the constraints can be violated and still considered feasible.

Engine: LP/Quadratic

Default: 10^{-7}

Allowed range: 0 to 1

DualTolerance The maximum amount by which the dual constraints and still considered feasible.

Engine: LP/Quadratic

Default: 10^{-7}

Allowed range: 0 to 1

Unused

There are a few Optimizer settings that are not used by the standard engines in Analytica Optimizer, even though they do show up on the list of settings. Some of these are used by add-on engines (add-on engines have their own set of additional parameters in general).

Crashing

IntCutoffHigh, deprecated, used **IntCutoff**

IntCutoffLow, deprecated, use **IntCutoff**

PrecisionTol

SolutionAccuracy

SolutionResolution

SolutionTol

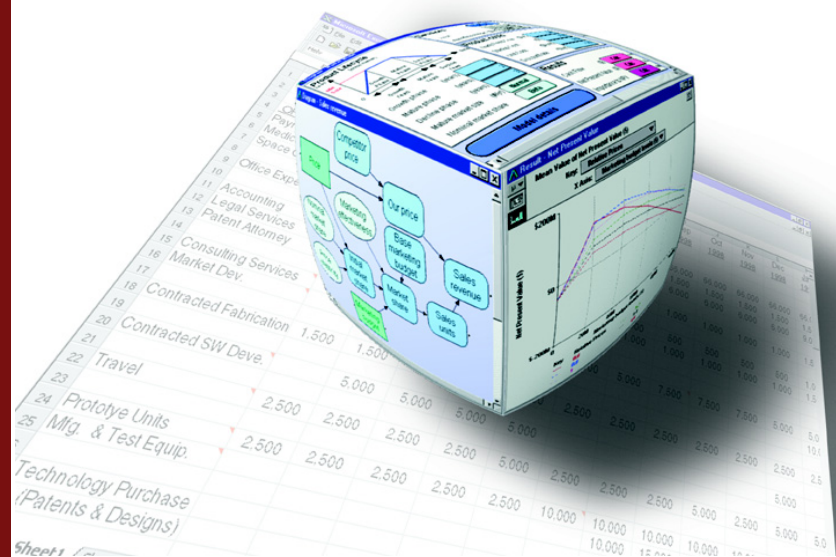
VariableReordering

Chapter 7

Debugging Optimization Problems

This chapter shows you how to:

- Write a linear or quadratic optimization formulation to a file for inspection
- Diagnose conflicting constraints
- Debug a nonlinear optimization



Writing and reading from a file

A linear or quadratic optimization formulation can be written to (and read from) a text file using these functions.

```
LpWrite()      LpWrite(lp: LpType; filename: Text )
LpRead()      LpRead(filename: Text)
```

LpWrite() returns the full filename path written to. **LpRead()** returns an <<LP>> or <<QP>> object. Viewing the resulting file can sometimes be useful for detecting problems with your call to **LpDefine()** or **QpDefine()**. These functions cannot be used on a nonlinear optimization. The *filename* variables are interpreted relative to the current Analytica data directory.

Both functions accept an optional parameter.

```
format : optional text
```

This parameter accepts the value of **LP**, **MPS**, or **PFML**. These are standard file formats used for exchanging problem specifications between other optimization software.

LpRead() also optionally accepts two indexes.

```
Vars,Constraints : optional Index
```

If specified, the length of each of these indexes must match the number of variables and number of constraints in the file being read exactly. When they are not specified, local indexes are created.

Diagnosing conflicting constraints

If you have conflicting constraints in your formulation, there is no feasible solution. When you have many constraints, you can find the conflicting constraints by computing an *irreducibly infeasible subset (IIS)* of constraints using one of the following functions.

```
LpFindIIS()   LpFindIIS(lp: LpType)
LpWriteIIS()  LpWriteIIS(lp: LpType; filename: Text)
```

An IIS of constraints is a subset of your constraints that contains no feasible solution, but that has the property that if any single constraint is removed, there is feasible solutions. Thus, it is a minimal set of conflicting constraints.

LpFindIIS() returns a subset of your **Constraints** index. This can be used on linear constraints defined from **LpDefine()** or **QpDefine()**.

LpWriteIIS() writes the IIS to an indicated file and returns the full file path. This function can be used with linear and quadratic optimizations with linear constraints, but not with nonlinear optimization problems. The file format is the same as that used by **LpWrite()**. An optional *format* parameter of **LP**, **MPS**, or **LPFML** can also be included.

When finding an IIS, there is an option of whether to only remove constraints, or whether variable bounds can also be removed in order to find an IIS. By removing variable bounds, it might be possible to find an IIS with a larger number of constraints. By default, **LpFindIIS()** removes only constraints, leaving variable bounds alone, which is necessary since a subset of the **Constraints** index is returned. To allow variable bounds to be relaxed, you must include a second parameter to **LpFindIIS()**.

```
LpFindIIS( lp, newLp: true )
```

When this second parameter is included, **LpFindIIS()** returns a new <<LP>> object, the same type of object returned by **LpDefine()**. The new object is still infeasible, but using it you can examine the reduced set of constraints and reduced set of variable lower and upper bounds using these expressions, where **lp** is the object returned by **LpFindIIS()**.

```
SolverInfo( "Constraints", lp )
SolverInfo( "lb", lp )
SolverInfo( "ub", lp )
```

By default, **LpWriteIIS()** relaxes both variable bounds and constraints. The setting **IIS-Bounds** can be used to override this behavior for both **LpWriteIIS()** and **LpFindIIS(..., newLp:true)**, see “IISBounds” on page 72.

Debugging a nonlinear optimization

After formulating a nonlinear problem, you might find that the optimization runs and returns something other than what you expect. After viewing the **LpStatusText()**, it might not be clear why it terminated where it did, or why it didn't succeed in solving your optimization as you desire. In these cases, you might need to monitor the optimization while it is searching in order to debug why it is doing what it is doing.

TraceFile This is an optional parameter to **NlpDefine()**.

TraceFile : optional Text

TraceFile can be given a filename to write a log of all points visited during the optimization search. Written to the file are the values of the decision variables at each evaluation, the value computed for the objective, and the computed Jacobian and gradient values if those expressions are also provided to **NlpDefine()**. The file can then be viewed in a text editor such as TextPad or Notepad. The values are tab-separated, so adjusting tab width in your text editor can help with readability. By studying how the search progressed, it often becomes evident why the Optimizer is behaving as it is. When this is understood, this can help to uncover errors in your problem formulation, or suggest approaches to improve the search.

Using MsgBox() to Debug

Another convenient “trick” is to simply peek at what values Optimizer is plugging in for **x** in a more interactive fashion while the search is taking place. You can do this by inserting **MsgBox()** inside the expression that computes your objective (or in any node downstream of **x** and upstream of your objective expression). For example, if your objective expression is this

```
obj: Sum(Exp(-a*x), Vars)
```

you can modify this to read as follows.

```
obj: MsgBox(x, 0, "X="): Sum(Exp(-a*x), Vars)
```

Then each time the Optimizer evaluates the objective, a message box appears on the screen, allowing you to view progress. Seeing the Optimizer in action often gives you an understanding of what it would take to improve the search.

There are a few quirks to be aware of when using **MsgBox()** in this fashion. First, the **MaxTimeNoImp** parameter specifies a maximum time the Optimizer works with no improvement in the best feasible solution found so far. Time spent staring at the message box counts toward time spent, and can result in an earlier termination. If this happens, you might want to explicitly set this parameter in your call to **NlpDefine()** to something large.

A second quirk is that if you decide to print out multiple pieces of information with a message box, you must consider how they will array abstract. **MsgBox()** prints out a

description of your entire array result, but its parameter is evaluated before it even considers printing it.

So, if you call **MsgBox()** using `MsgBox("x=" & x)` when `x` is array-valued, you see something like `{1} Array(Vars, [X=0.2, X=0.5, X=-0.3])` rather than `X=Array(Vars, [X=0.2, 0.5, -0.3])` as you might have expected.

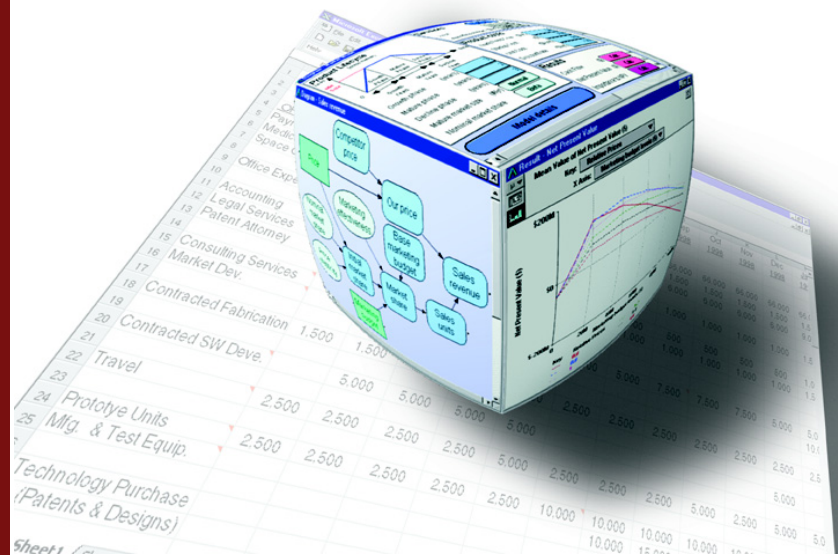
If you plan on displaying multiple variables in the same message box, consider using expressions such as `MsgBox("X=[" & JoinText(X, Vars, ",") & "]")` which outputs `[X=0.2, 0.5, -0.3]`.

You can scatter **MsgBox()** calls throughout expressions to peek at the optimization at various points as it progresses.

Chapter 8

Optimizer Function Reference

This chapter lists and defines all the Analytica optimization functions.



Problem definition functions

When defining an optimization problem, we highly recommend using a named-parameter syntax, rather than relying on the parameter being the first, second, or third parameter, etc. In a named-parameter syntax, you type the parameter name, followed by a colon (:), followed by the parameter value. Here is an example.

```
NlpDefine( X:d, Obj: F(d) )
```

In the descriptions of **LpDefine()**, **QpDefine()**, and **NlpDefine()** that follow, the parameters are shown in a logical grouping, but not in the actual order. We assume you will use named-parameter syntax. You can view the full parameter declarations from Analytica, in the actual parameter order, by selecting **Definition > Optimizer > <function>** from the Analytica menu which viewing an influence diagram with nothing selected.

LpDefine()

The **LpDefine()** function defines a linear optimization program. See Chapter 3, “Linear Optimization,” for a description of usage and parameters, and Chapter 6, “Control Settings,” for possible engine settings. Parameters are described in the table below.

Parameters of LpDefine()

Parameter	Type
Decision variables	
vars	<i>Index</i>
lb	Optional <i>Number[vars]</i>
ub	Optional <i>Number[vars]</i>
Ctype	Optional <i>Text[vars] = 'C'</i>
group	Optional <i>Number[vars]</i>
Objective function	
ObjCoef	<i>Number[vars]</i>
maximize	Optional <i>Boolean = false</i>
Constraint specification	
constraints	<i>Index</i>
lhs	<i>Number[vars, constraints]</i>
sense	Optional <i>Text[constraints] = '<='</i>
rhs	<i>Number[constraints]</i>
Engine settings	
engine	Optional <i>Text</i>
parameter	Optional <i>Text</i>
setting	Optional <i>Number</i>
Deprecated	
ItLimit	Optional <i>Positive</i>
NdLimit	Optional <i>Positive</i>
MipLimit	Optional <i>Positive</i>
TimeLimit	Optional <i>Positive</i>
optTolerance	Optional <i>Positive</i>
pivotTolerance	Optional <i>Positive</i>
feasTolerance	Optional <i>Positive</i>

Parameter	Type
gapTolerance	Optional <i>Positive</i>
OptLb	Optional <i>Number</i>
OptUb	Optional <i>Number</i>
scaling	Optional <i>Number</i>

QpDefine()

The **QpDefine()** function defines a quadratic optimization program. See Chapter 4, “Quadratic Optimization,” for a description of usage and parameters, and Chapter 6, “Control Settings,” for possible engine settings.

Parameters of QpDefine()

Parameter	Type
Decision variables	
vars	<i>Index</i>
vars2	<i>Index</i>
lb	Optional <i>Number[vars]</i>
ub	Optional <i>Number[vars]</i>
Ctype	Optional <i>Text[vars] = 'C'</i>
group	Optional <i>Number[vars]</i>
guess	Optional <i>Number[vars]</i>
Objective function	
c	<i>Number[vars]</i>
Q	<i>Number[vars, vars2]</i>
maximize	Optional <i>Boolean = false</i>
Constraint specification	
constraints	<i>Index</i>
lhs	Optional <i>Number[vars, constraints]</i>
lhsQ	Optional <i>Number[vars, vars2, constraints]</i>
sense	Optional <i>Text[constraints] = '<='</i>
rhs	<i>Number[constraints]</i>
Engine settings	
engine	Optional <i>Text</i>
parameter	Optional <i>Text</i>
setting	Optional <i>Number</i>
Deprecated	
warnIndefinite	Optional <i>Boolean</i>
ItLimit	Optional <i>Positive</i>
NdLimit	Optional <i>Positive</i>
MipLimit	Optional <i>Positive</i>
TimeLimit	Optional <i>Positive</i>
optTolerance	Optional <i>Positive</i>
pivotTolerance	Optional <i>Positive</i>
feasTolerance	Optional <i>Positive</i>
gapTolerance	Optional <i>Positive</i>

Parameter	Type
optLb	Optional <i>Number</i>
OptUb	Optional <i>Number</i>
scaling	Optional <i>Number</i>

NlpDefine()

The **NlpDefine()** function defines a nonlinear optimization problem. See Chapter 5, “Nonlinear Optimization,” for a description of usage and parameters, and Chapter 6, “Control Settings,” for possible engine settings.

Parameters of NlpDefine()

Parameter	Type
Decision variables	
vars	Optional <i>Index</i>
x	<i>LVarType</i> { <i>global or local variable</i> }
lb	Optional <i>Number</i> [vars]
ub	Optional <i>Number</i> [vars]
Ctype	Optional <i>Text</i> [vars] = 'C'
group	Optional <i>Number</i> [vars]
guess	Optional <i>Number</i> [vars]
Objective function	
obj	Optional <i>expression</i> { <i>atomic</i> }
maximize	Optional <i>Boolean</i> = <i>false</i>
objNI	Optional <i>Text</i> [vars] = 'N'
gradient	Optional <i>Expression</i> { <i>vars</i> }
Constraint specification	
constraints	Optional <i>Index</i>
lhs	Optional <i>Expression</i> { <i>constraints</i> }
sense	Optional <i>Text</i> [constraints] = '<='
rhs	Optional <i>Number</i> [constraints] = 0
lhsNI	Optional <i>Text</i> [vars, constraints] = 'N'
jacobian	Optional <i>Expression</i> { <i>constraints, vars</i> }
Engine settings	
engine	Optional <i>Text</i>
parameter	Optional <i>Text</i>
setting	Optional <i>Number</i>
Deprecated	
itLimit	Optional <i>Positive</i>
noImpSeconds	Optional <i>Positive</i>
timeLimit	Optional <i>Positive</i>
convTolerance	Optional <i>Positive</i>
mutate	Optional <i>Positive</i>
linVar	Optional <i>Scalar</i>
DerivMethod	Optional <i>Text</i>

Parameter	Type
EstimMethod	Optional <i>Text</i>
DirecMethod	Optional <i>Text</i>
SampSz	Optional <i>Positive</i>

Other functions

LpFindIIS(lp: LpType : newLp : optional boolean)

Computes and returns the *irreducibly infeasible subset (IIS)* of the constraints. This is meaningful when `LpStatus(lp)=2` (“no feasible solution”), and is useful for identifying what portions of your constraint formulation make the problem infeasible.

When the optional parameter, **newLp**, is specified, returns a new <<LP>> object having the subset of constraints (still infeasible). The components of this object can be accessed using `SolverInfo()`.

LpObjSa(lp: LpType; v: Optional)

Returns the sensitivity ranges for the objective function coefficients for a linear program **lp** for decision variable(s) **v**, which should be one of or a subset of decision variables, **Vars**. If **v** is omitted, it computes the sensitivity for all **Vars**.

LpOpt(lp: LpType)

Returns the value of the objective function at the optimum.

LpRead(filename: Text; vars, constraints: optional Index; format: optional Text)

Reads a linear or quadratic program definition from file **filename**, previously written by `LpWrite()` and returns an opaque <<LP>> or <<QP>> object. The optional **Vars** and **Constraints** are the corresponding indexes for the LP, and must be of the same size as the problem read in. The optional **format** parameter can be `LP` (default), `MPS`, or `LPFML` to indicate the type of file being read.

LpReducedCost(lp: LpType)

Returns the reduced costs (dual values) of each variable as an array indexed by **Vars**.

LpRhsSa(lp: LpType; constraint: Optional)

Returns the sensitivity ranges for the **RHS** values. The default is to compute sensitivities for all **RHS** values, with the result indexed by **Constraints**. If you specify the optional second parameter, it returns the sensitivity for only that constraint or subset of constraints.

LpShadow(lp: LpType)

Returns the shadow prices (dual values of the constraints) as an array indexed by **Constraints**.

LpSlack(lpv)

Returns the slack or surplus values at the optimal solution as an array indexed by **Constraints**. If it cannot find an optimal solution, it generates an appropriate error.

LpSolution(lp: LpType)

Returns the optimal solution to the linear, quadratic, or nonlinear programming problem **lp** defined by **LpDefine()**, **QpDefine()**, or **NlpDefine()**. The result is an array of decision variables indexed by **Vars**. If it cannot find an optimal solution, **LpSolution()** returns the best values found during the search so far, and **LpStatusNum()** and **LpStatusText()** indicate why it has not found an optimal solution.

LpStatusNum(lp: LpType) LpStatusText(lp: LpType)

Returns the status number as an integer and corresponding text message, respectively, of the optimization problem **lp**. It is wise to examine the status before evaluating **LpSolution()** to avoid an error message. Possible results are shown in the table below.

Status Number	Status Text
-3	Invalid status.
-2	Ignore status. Used when dummy result code needs to be overridden.
-1	Invalid license status. (License expired, missing, invalid, etc.)
0	Optimal solution has been found.
1	The Solver has converged to the current solution.
2	"No remedies" status. (All remedies failed to find better point.)
3	Iterates limit reached. Indicates an early exit of the algorithm.
4	Optimizing an unbounded objective function.
5	Feasible solution could not be found.
6	Optimization aborted by user. Indicates an early exit of the algorithm.
7	Invalid linear model. Returned when a linearity assumption renders incorrect.
8	Bad data set status. Returned when a problem data set renders inconsistent.
9	Float error status. (Internal float error.)
10	Time out status. Returned when the maximum allowed time has been exceeded. Indicates an early exit of the algorithm.
11	Memory dearth status. Returned when the system cannot allocate enough memory to perform the optimization.
12	Interpretation error. (Parser, Diagnostics, or Executor error.)
13	Fatal API error. (API not responding.)
14	The Solver has found an integer solution within integer tolerance.
15	Branching and bounding node limit reached. Indicates an early exit of the algorithm.
16	Branching and bounding maximum number of incumbent points reached. Indicates an early exit of the algorithm.
17	Probable global optimum reached. Returned when MSL (Bayesian) global optimality test has been satisfied.
18	Missing bounds status. Returned for EV/MSL Require Bounds when bounds are missing.
19	Bounds conflict status. Indicates <=, =>, = bounds conflict with existing binary or all different constraints.
20	Bounds inconsistency status. Returned when the lower bound value of a variable is greater than the upper bound value, i.e., lb[i] > ub[i] for some variable bound i.
21	Derivative error. Returned when API_Jacobian has not been able to compute gradients.
22	Cone overlap status. Returned when a variable appears in more than one cone.
999	Exception occurred status. Returned when an exception has been caught by try/catch top-level.
1000	Custom base status. (Base for Solver engine custom results.)
1102	The quadratic constraints are non-convex, the SOCP engine cannot solve this problem.

Note: *The possible status numbers and messages returned in Analytica Optimizer 4.0 have changed since Analytica 3.1, due to changes in the underlying Frontline solver. If you have legacy models that test against specific status numbers, you need to adjust these accordingly.*

LpWrite(lp: LpType; filename: Text; format: optional Text)

Writes a text description of a linear or quadratic program, **lp**, defined using **LpDefine()** or **QpDefine()**, to a file with the specified filename. Note that if **lp** is an array of LP problems, and the filename does not share the same dimension, the file written by **LpWrite()** contains the result of only the last **lp**.

LpWriteIIS(lp: LpType; filename: Text; format: optional Text)

Writes an irreducibly infeasible subset (IIS) of a linear or quadratic program to a file, including only a subset of constraints that is infeasible, but with the property that if any single constraint is removed, the resulting problem will be feasible. The format is the same as that used by **LpWrite()**.

SolverInfo(item: Text; lp: optional LpType; engine: optional Text)

Returns general Optimizer information, internals of a specific optimization problem, or information about an engine. There are three styles of use.

SolverInfo(item) Information about the Optimizer. Possible values for item include:

- **AvailEngines:** Returns a list of installed Optimizer engines.

SolverInfo(item, lp) Returns information about an optimization problem definition. Possible values for item include the following, where dimensionality of the result is shown in brackets:

- **type:** Problem type, one of **LP**, **QP**, **QCP**, or **NLP**.
- **Vars** [vars]: Elements of the vars index, i.e., variable names.
- **lb** [vars]: Lower bound for each variable. Indexed by vars.
- **ub** [vars]: Upper bound for each variable. Indexed by vars.
- **Ctype** [vars]: Integer type for each variable. One of **C**, **I**, **B**, or **G**.
- **group** [vars]: Group number for grouped-integer variables.
- **Maximize:** 0 for minimization, 1 for maximization problem.
- **ObjCoef** [vars]: (LP, QP) Objective coefficients.
- **Q** [vars,vars2]: (QP) Objective function quadratic matrix.
- **Ihs** [constraints,vars]: (LP, QP) The linear constraint coefficients.
- **IhsQ** [constraints,vars,vars2]: (QP) The quadratic constraint matrices.
- **sense** [constraints]: One of **>=**, **<=**, or **=** for each constraint.
- **rhs** [constraints]: Right-hand side coefficient for each constraint.
- **constraintUb** [constraints]: Upper bound for each constraint.
- **constraintLb** [constraints]: Lower bound for each constraint.
- **engine:** The engine name used to solve the problem.


- **setting** [parameter]: The engine settings for this problem. The index is a local index, whose elements depend on the engine used for the problem.

Solver(item, engine: "engineName") Returns information about an installed Optimizer engine. The items providing engine capabilities are indexed by a local index, **ProblemType**, which includes [**LP**, **QP**, **QCP**, **CVX**, **NLP**, and **NSP**].

Chapter 9

Logistic Regression Function Reference

This chapter describes the Analytica logistic regression functions.



Logistic regression functions

The **Generalized Regression.ana** library contains functions that you can use to estimate the probability (or probability distribution) of a dependent (output) variable as a function of known values for independent (input) variables. This is similar to linear regression, which predicts the value of a dependent variable as a function of known values for independent variables. Logistic regression is the best known example generalized regression, so even though the term logistic regression technically refers to one specific form of generalized regression (with prob and poisson regression being other instances), it is also not uncommon to hear the term logistic regression functions used synonymously with generalized regression, as we have done with the title of this chapter.

To use the functions described in this chapter, you must have Analytica Optimizer and you must add the **Generalized Regression.ana** library to your model using the **Add Library** option from the **File** menu.

Logistic_regression(y, b, i, k)

Logistic regression is a technique for predicting a Bernoulli (i.e., 0,1-valued) random variable from a set of continuous dependent variables. See the Wikipedia article on logistic regression (http://en.wikipedia.org/wiki/Logistic_regression) for a simple description. Another generalized logistic model that can be used for this purpose is the **Probit_regression()** model. These differ in functional form, with the logistic regression using a **logit** function to link the linear predictor to the predicted probability, while the probit model uses a cumulative normal for the same.

The **Logistic_regression()** function returns the best-fit coefficients, **c**, for a model of this form given a data set basis **b**, with each sample classified as **y_i**, having a classification of 0 or 1.

$$\ln\left(\frac{p_i}{1-p_i}\right) = \sum_k c_k b_{i,k}$$

The syntax is the same as for the **Regression()** function. The basis can be of a generalized linear form, that is, each term in the basis can be an arbitrary nonlinear function of your data; however, the **logit** of the prediction is a linear combination of these.

When you have used the **Logistic_regression()** function to compute the coefficients for your model, the predictive model that results returns the probability that a given data point is classified as 1.

Example Suppose you want to predict the probability that a particular treatment for diabetes is effective given several lab test results. Data is collected for patients who have undergone the treatment, as follows, where the variable **Test_results** contains lab test data and **Treatment_effective** is set to 0 or 1 depending on whether the treatment was effective or not for that patient.

Patient ID	Fasting Plasma Glucose	Test Plasma Glucose	Test insulin level	Steady-state glucose	Treatment effective
1	130	670	44	167	1
2	100	429	201	194	0
3	114	540	188	211	0
4	82	390	375	273	0
5	92	580	132	155	0
6	93	391	221	103	0
7	91	436	148	167	0
8	87	360	292	128	0
9	213	1001	42	297	1
10	97	379	142	98	0
11	90	345	123	50	0

Using the data directly as the regression basis, the logistic regression coefficients are computed using this.

```
Variable c := Logistic_regression( Treatment_effective,
                                Test_results, Patient_ID, Lab_test )
```

We can obtain the predicted probability for each patient in this testing set this.

```
Variable Prob_Effective := InvLogit( Sum( c*Test_results,
                                         Lab_Test ) )
```

If we have lab tests for a new patient, say `New_Patient_Tests`, in the form of a vector indexed by `Lab_Test`, we can predict the probability that treatment will be effective this.

```
InvLogit( Sum( c*New_patient_tests, Lab_test ) )
```

Probit_regression(y, b, i, k)

A probit model relates a continuous vector of dependent measurements to the probability of a binomial (i.e., 0,1-valued) outcome. In econometrics, this model is sometimes called the *Harvard model*. The **Probit_regression()** function infers the coefficients of the model from a data set, where each point in the training set is classified as 0 or 1.

Probit regression is very similar to **Logistic_regression()**. Both are used to fit a binomial outcome based on a vector of continuous dependent quantities. They differ in their use of the **link** function.

Given a set of data points, indexed by **i**, with each point classified as **0,1** in the **Y** parameter, and a set of basis terms, **b**, containing the dependent variables (where the vector of dependent variables is indexed by **k**), the **Probit_regression()** function finds and returns the set of coefficients for the probit model where Φ is the inverse cumulative normal distribution function.

$$p_i = \Phi\left(\sum_k c_k b_k\right)$$

The basis, **b**, is a function of the dependent variables in your data. Each element along **k** of the basis vector can be an arbitrary, even nonlinear, combination of the data in your data set. However, the number of terms in the basis should be kept small relative to the number of data point in your data set.

Example Probit regression can be applied to the same prediction problem example shown above for logistic regression. The probit coefficients are obtained using this.

```
Variable c2 := Probit_regression( Treatment_effective, Test_results,
                                Patient_ID, Lab_test )
```

The predicted probability for a new patient (with lab tests given by `New_patient_tests`) is given by this.

```
CumNormal( Sum( c2*New_patient_tests, Lab_test ) )
```

Library `Generalized Regression.ana`

Poisson_regression(y, b, i, k)

A Poisson regression model is used to predict the number of events that occur, **y**, from a vector independent data, **b**, indexed by **k**. The **Poisson_regression()** function computes the coefficients, **c**, from a set of data points, (**b**, **y**), both indexed by **i**, such that the expected number of events is predicted by this formula.

$$E(Y) = \exp\left(\sum_k c_k b_k\right)$$

The random component in the prediction is assumed to be Poisson-distributed, so that given a new data point **b**, the distribution for that point is shown below.

```
Poisson(sum(c*B,K)
```

If your dependent variable is continuous, with normally distributed error, use **Regression** or **RegressionDist**². If your dependent variable is binomially distributed (i.e., 0,1-valued), use **Logistic_Regression()** or **Probit_Regression()**. If your dependent variable models a count, such as the number of events that occur, use **Poisson_Regression()**.

***Note:** The distribution here accounts for data variation only, and does not include error in the coefficients **c**, as the **RegressionDist()** function does, for example.*

Library `Generalized Regression.ana`

2. To use **RegressionDist**, add the `Multivariate Distributions.ana` library to your model.

A

- add-on engines, installing 5
- ADE licenses 3
- Airline NLP.ana
 - array abstraction 60
 - introduction 47
- Airline_nlp() function
 - defining 53
 - optimizing for each year 56
 - parametric analysis 55
- algorithms
 - branch-and-bound, *see* branch-and-bound algorithm
 - cuts 31
 - debugging 72
 - dual simplex 76
 - genetic 47
 - long-step barrier 73
 - mixed-integer controls 76
 - numeric estimation 72
 - preprocessing 71
 - primal simplex 76
 - selecting 71
 - simplex, *see* simplex algorithm
 - StrongBranching 33
- Analytica Decision Engine (ADE) licenses 3
- Analytica Optimizer
 - about 2
 - edition 3
 - new features in 4.0 5
- Analytica Power Player with Optimizer 3
- array abstraction
 - about 34
 - in NLPs 51
 - nonlinear programs 45
 - restrictions on in NLPs 42
 - summary for NLPs 59
- Asset Allocation.ana 39, 60
- atomic parameters 53
- Automobile production.ana 30

B

- Beasley, J.E. 9
- Bernoulli random variable 96
- Big Mac Attack.ana 30
- binary (Boolean) decision variables 20
- bounds
 - about 20
 - constraints 20
 - illustration 20
 - nonlinear programs 45
 - specifying upper and lower 25
 - tightening 76
- BoundsImprovement 76

Index

- branch-and-bound algorithm
 - controlling searches 31
 - mixed-integer constraints 67
 - mixed-integer controls 76
 - settings 33

C

- calling convention, named 37
- Capital Investment.ana 30
- chance variables 52
- CliqueCuts 33, 80
- coefficient sensitivity 28
- conjugate-gradient search 72
- constants
 - about 20
 - linear programs 24
- constraints
 - about 20
 - convex 38
 - debugging conflicting 84
 - hard vs. soft 47
 - illustration 20
 - LHS 36
 - linear programs 24
 - maximum 67, 81
 - mixed-integer 67
 - multiple 44
 - nonlinear programs 42
 - optimization 2
 - optimization without 43
 - quadratic 5
 - quadratic programs 21, 38
 - representing in NLPs 44
 - scaling 33
 - sense 25
 - setting up 15
 - slack or surplus 27
 - special ordered sets (SOS) 79
 - zero slack 28
- Constraints index
 - LHS and RHS 43
 - linear programs 24
 - multiple constraints 44
 - nonlinear programs 43
 - omitting 6
- continuity type parameter 13
- continuous decision variables 20
- Convergence 70
- convex constraints 38
- convex quadratic optimization 21
- costs, reduced 29
- covariance matrices 38
- Crashing 82
- crossovers 75

- Ctype parameter
 - identifying decision variable type 30
 - integer constraints 47
 - linear programs 13
 - optimizations over time 57
 - single text character 30
 - SolveWithout algorithm 72
 - specifying continuity type 46

cuts

- about 31
- controlling number of passes 33
- description 78
- generation control 33, 78–80
- Gomory 78
- LP/Quadratic engine 67
- methods 34

D

- decision variables
 - about 20
 - binary (Boolean) 20, 30, 46
 - bounds 45
 - continuous 20, 30, 62
 - grouped-integer 30, 46
 - identifying 9
 - integer 20, 30, 46
 - integer-valued 47
 - lower and upper bounds 25
 - mixed-integer 20, 30
 - obtaining solution 25
 - optimization problems 20
 - reformulating 49
 - scaling 33
- Decisions variable 49
- dependence, type 42, 61
- Derivatives setting 72
- derivatives, partial 42
- DerivMethod parameter 62
- deterministic pattern search 75
- discontinuous optimization problems 61, 63
- dual scaling method 73
- dual simplex algorithm 76
- dual values 29, 91
- DualTolerance 82
- dynamic loops
 - embedding linear programs 34
 - NLP with dynamic models 59
- dynamic models
 - array abstraction 42
 - NLP with 58
- Dynamic() function 58

E

- edit tables

- filling in 11
 - self-indexed 32
 - setting dimension 15
 - specifying control settings 68
 - EigenDecomp() function 38
 - Eigenvalues 38
 - Engine parameter
 - examining available settings 69
 - forcing use of Evolutionary 61
 - integer constraints 47
 - linear optimization settings 32
 - nonlinear programs 63
 - quadratic programs 37
 - selecting the optimization engine 66
 - engines
 - Evolutionary, *see* Evolutionary engine
 - examining available settings 69
 - GRG Nonlinear, *see* GRG Nonlinear engine
 - listing installed 63
 - LP/Quadratic, *see* LP/Quadratic engine
 - maximum variables and constraints 2, 67
 - nonlinear 6
 - selecting 63, 66
 - SOCP Barrier, *see* SOCP Barrier engine
 - standard 66
 - viewing capabilities with SolverInfo() 67
 - equations
 - parameters and constraints 11
 - representation 44
 - simultaneous 21
 - solving nonlinear 60
 - Estimates 73
 - Evolutionary engine
 - constraint types 47
 - engine selection 63
 - gradient information 63
 - search control algorithms 74
 - selecting 66
 - type of dependence 61
 - Example Models directory 14, 30, 47, 60
 - expected value, maximizing 51
 - expressions, gradient and Jacobian 42
 - ExtinctionRate 74
- F**
- Feasibility 75
 - FeasibilityPump 77
 - FeasibilityTolerance 74
 - Fermat's last theorem 21
 - files
 - format parameter 91
 - standard formats 84
 - writing and reading 84
 - fitness values 70
 - FixNonSmooth 75
 - FlowCoverCuts
 - cut generation 33
 - description 80
 - Frontline Solver engine 21
 - functions
 - browsing 8
 - logistic regression reference 96
 - problem definition reference 88
- G**
- gap tolerance 31, 74
 - generalized regression 96
 - Generalized Regression.ana library 96
 - genetic algorithm 47
 - genetic algorithm, *see* Evolutionary engine
 - Getfrac() function 52
 - GoalSeek function 44
 - gomory cuts
 - description 78
 - maximum number 78
 - number of passes 78
 - using 78
 - GomoryCuts 33
 - gradient local searches 75
 - gradient parameter
 - array abstraction 60
 - computing derivatives 42
 - derivatives 72
 - description 62
 - indexing 51
 - nonlinear programs 42
 - gradient-based searches 61
 - GreedyCover 77
 - GRG Nonlinear engine
 - description 67
 - engine selection 63
 - quadratic programs 37
 - selecting 66
 - group parameter
 - about 31
 - nonlinear programs 47
 - omitting 47
 - quadratic programs 37
 - guess parameter
 - indefinite objectives 38
 - initial guess 62
 - quadratic programs 37
- H**
- hard-integer constraints 47
 - Helmberg, Kojima, Monteiro dual scaling 73
 - Hessian of objective function 37
 - heuristics

Index

- local 77
- primal 33, 77
- rounding 77

HKM dual scaling 73
Hoffman and Padberg 80

I

IISBounds 72
indefinite objectives 38
indexes, constraint 10
initial guess, *see* guess parameter
installation, add-on engines 5
IntCutoff 33, 76
IntCutoffHigh 82
IntCutoffLow 82
integer constraints

- decision variables 20
- hard vs. soft 47
- see also* decision variables, integer

integer cut-set procedures, *see* cuts
Intelligent Arrays

- abstraction 51
- support for 21

IntTolerance 33, 70
irreducibly infeasible subset (IIS)

- diagnosing conflicting constraints 84
- writing to a file 93

Iterations 33, 69

J

Jacobian parameter

- array abstraction 60
- computing derivatives 42
- derivatives 72
- description 62
- indexing 51
- nonlinear programs 42

K

KnapsackCuts 33, 79
KnapsackPasses 79
KNITRO add-on 5

L

large-scale add-ons 5
lb parameter

- about 11
- quadratic programs 37
- setting bounds 45
- specifying 25

left-hand side, *see* LHS and Lhs parameter
LHS

- linear programs 24

- nonlinear programs 43

Lhs parameter

- about 11
- array abstraction 45
- illustration 20
- indexing 51

lhsNI parameter 61
licenses

- activating 3, 4
- obtaining 3

LiftAndCoverCuts 33, 80
lifted cover inequalities 79
LINDO package 29
linear dependence 61
linear optimization

- about 21
- debugging 84
- defining problems 24
- standard formulation 24
- working with 24–34
- writing and reading text files 84

linear programs

- branch and bound settings 33
- constraints 24
- controlling searches 31
- cut generation settings 33
- dynamic loops 34
- example optimization models 30
- LHS 24
- objective coefficients 24
- obtaining solutions 25
- optimization 21
- preprocessing settings 33
- RHS 24
- secondary solutions 27
- settings 33
- termination controls 33
- tolerance and precision settings 34
- using 8–14
- working with 24–34

linear regression 96
local heuristics 77
local optima, working with 62, 80
LocalHeur 77
LocalSearch 47, 67, 75
LocalTree 77
Logistic_regression() function 96
long-step barrier algorithm 73
LP file format 84, 91
LP, *see* linear programs
LP/Quadratic engine

- control settings 32
- cut generation control 78
- description 67

- dual simplex method 76
 - mixed-integer control 77
 - numeric tolerance 81
 - preprocessing 71
 - quadratic programs 37
 - selecting 66
 - LpDefine() function
 - array abstraction 34
 - control settings 31
 - controlling Optimizer behavior 31
 - core parameters 24
 - description and syntax 88
 - engines to use with 66
 - name-based calling syntax 12
 - optional parameters 25
 - selecting the optimization engine 66
 - specifying settings 68
 - variable bounds 25
 - LpFindIIS() function
 - description 91
 - using 84
 - variable bounds 72
 - LPFML file format 84, 91
 - LpObjSa() function
 - description 91
 - secondary solutions 28
 - LpOpt() function
 - description 91
 - obtaining the solution 26
 - LpRead() function
 - description 91
 - using 84
 - LpReducedCost() function
 - description 91
 - dual values 29
 - quadratic programs 39
 - LpRhsSa() function
 - description 91
 - quadratic programs 39
 - secondary solutions 28
 - LpShadow() function
 - description 91
 - quadratic programs 39
 - shadow prices 29
 - LpSlack() function
 - description 92
 - quadratic programs 39
 - secondary solutions 27
 - LpSolution() function
 - description 92
 - obtaining the solution 26
 - solving quadratic programs 39
 - LpStatusNum() function
 - description 92
 - obtaining the solution 26
 - solving quadratic programs 39
 - success finding solutions 31
 - termination controls 69
 - LpStatusText() function
 - description 92
 - listing of all codes 27, 92
 - obtaining the solution 26
 - solving quadratic programs 39
 - success finding solutions 31
 - LpWrite() function
 - description 93
 - using 84
 - LpWriteIIS() function
 - description 93
 - using 84
 - variable bounds 72
 - Lumina contact info and web site 3
- M**
- MaxFeasibleSolutions 70
 - MaxGomoryCuts 78
 - Maximize parameter
 - linear programs 25
 - nonlinear programs 43, 44
 - optimization 20
 - quadratic programs 37
 - MaxIntegerSols 71
 - MaxKnapsackCuts 79
 - MaxRootCutPasses 33, 78
 - MaxSubProblems 70
 - MaxTime 33, 69
 - MaxTimeNoImp 33, 70, 85
 - MaxTreeCutPasses 33, 78
 - mean, maximizing 51
 - MirCuts 33, 79
 - mixed-integer controls 76
 - mixed-integer decision variables 20
 - mixed-integer programs 30
 - Monte Carlo sample 52
 - MOSEK add-on 5
 - MPS file format 84, 91
 - MsgBox() 85
 - MultiStart 74, 80
 - MultiStart setting 63
 - MutationRate 70, 74
- N**
- name-based calling syntax 12
 - named calling convention 37
 - named-parameter syntax 88
 - negative semi-definiteness 38
 - negative-definiteness 37
 - net present value (NPV) 55

Index

Newton gradient descent 75
NLP with Jacobian.ana 60
NLP, see nonlinear programs
NlpDefine() function
 array abstraction 45
 core parameters 43
 DerivMethod parameter 62
 description and syntax 90
 engines to use with 66
 giving hints to Optimizer 61
 GoalSeek 44
 MultiStart 63
 NLP using NPV 55
 optional parameters 44
 selecting the optimization engine 66
 specifying settings 68
 TraceFile parameter 85
 type of dependence 61
nonlinear equations, solving systems of 60
nonlinear optimization
 debugging 85
 standard formulation 42
 without constraints 43
nonlinear programs
 array abstraction 45, 59
 constraints 42
 defining and solving 14–17
 dynamic models 58
 maximizing value 51
 multiple constraints 44
 no constraints 43
 objectives 42
 optimization 21
 optimizations over time 57
 optimizing for each year 56
 over time using NPV 55
 parametric analysis 55
 probabilistic optimization 52
 problem formulation 42
 reformulating decision variables 49
 representing constraints 44
 simple optimization 50
 single decision variable 43
 working with 42–63
numeric estimation 72

O

Obj parameter
 array abstraction 45
 nonlinear programs 43, 51
ObjCoef parameter 24
objective coefficients 24
objective functions 20
objectives

 about 20
 LHS 36
 maximizing 25
 nonlinear programs 42
 quadratic program 21
objNI parameter 61
odd cycle cuts 79
OddHoleCuts 33, 79
optima, local 62, 80
Optimal can dimensions.ana 14, 61
optimal cost tolerance 81
Optimal production planning.ana 30
OptimalityFixing 77
optimization
 choosing type 21
 convex quadratic 21
 engines, see engines
 formulating problems 19
 linear, see linear programs
 nonlinear, see nonlinear programs
 parts of a problem 20
 possible outcomes 26
 quadratic, see quadratic programs
 simultaneous equations 21
OptQuest add-on 5
Over parameter 45

P

parameters
 analysis 42, 60
 atomic 53
 optional 68
parametric analysis 55
pivots, about 31
PivotTol 34, 81
Poisson_regression() function 98
PopulationSize 70, 74
portfolio allocation 38
positive semi-definiteness 38
positive-definiteness
 about 37
 covariance matrix 39
PowerIndex 73
Precision
 allowed range 81
 parameter description 34
precision settings
 linear programming 34
 numeric 81
PrecisionTol 82
Premium Solver 2
preposterior analysis 52
PreProcess 71
preprocessing

- algorithms 71
 - settings 33
- Presolve 33, 71
- prices, shadow 29
- primal problems 29
- primal simplex algorithm 76
- PrimalHeuristic 33, 77
- PrimalTolerance 34, 81
- probabilistic optimization 52
- ProbingCuts 33, 79
- ProbingFeasibility 33, 76
- Probit_regression() function 97
- problem formulations, debugging 84
- Problems with local optima.ana 61
- ProblemType index 67
- Production Planning LP.ana 30
- ProfitFn() function 52
- programs
 - linear, *see* linear programs
 - nonlinear, *see* nonlinear programs
 - optimization 21
 - quadratic, *see* quadratic programs

Q

- Q matrices 36
- QP, *see* quadratic programs
- QpDefine() function
 - description and syntax 89
 - engines to use with 66
 - optional parameters 36, 37
 - required parameters 36
 - selecting the optimization engine 66
 - SOCP barrier searches 73
 - solution properties 37
 - specifying settings 68
- quadratic dependence 61
- quadratic matrices 36
- quadratic optimization
 - about 21
 - standard formulation 36
 - writing and reading text files 84
- quadratic programs
 - common applications 38
 - constraints 21, 38
 - defining 36
 - objectives 21
 - obtaining solutions 39
 - optimization 21
 - quadratically constrained 36
 - search control settings 39
 - solution properties 37
 - working with 36–39
- quadratic terms 36
- quadratically constrained problems 67

- quasi-Newton
 - gradient descent 75
 - method 72

R

- randomized local search 75
- RandomSeed 74
- real-valued numbers 30
- RecognizeLinear 73
- RedSplitCuts 33, 79
- reduced costs
 - finding 91
 - tolerance 81
- ReducedTol 34, 81
- RequireBounds 80
- RHS
 - array abstraction 34
 - linear programs 24
 - nonlinear programs 43
- Rhs parameter
 - about 11
 - illustration 20
- right-hand side, *see* RHS and Rhs parameter
- rounding heuristics 77
- RoundingCuts 33, 80
- RoundingHeur 77
- Run system variable 34, 52, 57
- Run_context variable 57

S

- saddle points 38
- Scaling
 - algorithm selection 71
 - Gomory cuts 78
 - linear programming 33
 - specifying control settings 32, 68
- SearchDirection 73
- searches
 - control settings 39
 - controlling 31, 66
 - cut generation control 78–80
 - debugging control settings 72
 - deterministic pattern 75
 - evolutionary search controls 74
 - gradient local 75
 - gradient-based 61
 - mixed-integer controls 76
 - not used by Optimizer 82
 - numeric estimation settings 72
 - randomized local 75
 - search space 20
 - selecting algorithms 71
 - SOCP Barrier search algorithms 73
 - termination controls 69

Index

- topographic 81
 - viewing control settings 31
 - SearchOption 72
 - secondary solutions 27
 - second-order cone programming 67
 - semi-definiteness 38
 - sense
 - about 11
 - changing 25
 - constraints 44
 - linear programs 24, 25
 - nonlinear programs 45
 - optimization 20
 - parameter description 25
 - quadratic programs 37
 - SetContext parameter
 - array abstraction 45, 57
 - NLP 51
 - using 54
 - when not to use 59
 - setting parameter
 - control settings 32
 - using 68
 - settings
 - control 31
 - controlling searches 31
 - examining 69
 - specifying 68
 - shadow prices 29, 91
 - simplex algorithm
 - controlling searches 31
 - dual 76
 - LP quadratic engine 67
 - numeric tolerance and precision 81
 - primal 76
 - simultaneous equations 21
 - slack values
 - constraints 28
 - locating 92
 - smooth nonlinear dependence 61
 - SOCP Barrier engine
 - description 67
 - quadratic programs 37
 - search algorithms 73
 - selecting 66
 - solving problems using 73
 - soft-integer constraints 47
 - SolutionAccuracy 82
 - SolutionResolution 82
 - solutions
 - linear programs 25
 - nonlinear programs 44
 - obtaining 25
 - secondary aspects 27
 - tolerance 31
 - SolutionTol 82
 - Solve using NLP.ana 61
 - SolverInfo() function
 - AvailEngines 66
 - constraints 85
 - control settings 32
 - description 93
 - determining engine used for optimization 66
 - engine capabilities 67
 - examining available settings 69
 - listing installed engines 63
 - viewing settings 66
 - SolveWithout 72
 - SOSCuts 33, 79
 - special ordered sets (SOS) 79
 - status messages 27, 92
 - StepSize 72
 - StepSizeFactor 74
 - StrongBranching 33, 71
 - subproblems
 - controlling searches 31
 - UseDual settings 76
 - Sudoku with Optimizer.ana 30
 - surplus values 92
 - surplus, constraints 28
 - syntax
 - name-based calling 12
 - named-parameter 88
- ## T
- termination controls 33, 69
 - Time system variable
 - array abstraction 57
 - NLP using NPV 55
 - Time_context variable 57
 - tolerance settings
 - controlling searches 31
 - numeric 81
 - precision 34
 - termination controls 33, 70
 - TopoSearch 81
 - TraceFile parameter 85
 - Traveling salesman.ana 47, 61
 - trough 38
 - Two Mines Model.ana 30, 9
 - TwoMirCuts 33, 79
- ## U
- ub parameter
 - about 12
 - quadratic programs 37
 - setting bounds 45
 - specifying 25

upgrades 3
UseDual 33, 76

V

values

 dual 29

 obtaining 25

VariableReordering 82

variables, *see* decision variables

Vars parameter

 linear optimization 24

 nonlinear optimization 43

X

X variable 43

XPRESS add-on 5

Index