

A Decision-Support Model for Managing the Fuel Inventory of a Panamanian Generating Company

by

Roberto Perez-Franco

Electromechanical Engineer, Panama Technological University, 2001

Submitted to the Engineering Systems Division
in Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Logistics
at the
Massachusetts Institute of Technology

May 7th, 2004
[June 2004]

(c) 2004 Roberto Perez-Franco, All rights reserved

The author hereby grants to M.I.T. permission to reproduce and distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Signature of Author _____


Engineering Systems Division
May 9th, 2004

Certified by _____

Stephen Graves
Abraham J. Siegel Professor of Management,
Sloan School of Management Professor of Engineering Systems
Thesis Supervisor

Accepted by _____

Yosef Sheffi
Director, Center for Transportation & Logistics
Professor of Engineering Systems
and of Civil and Environmental Engineering



BARKER

A Decision-Support Model for Managing the Fuel Inventory of a Panamanian Generating Company

by
Roberto Perez-Franco

Submitted to the Engineering Systems Division on May 7th, 2004
in Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Logistics at the Massachusetts Institute of Technology

ABSTRACT

Bahia Las Minas Corp (BLM) is a fuel-powered generating company in the Panamanian power system. The purpose of this thesis is to design and evaluate a decision-support model for managing the fuel inventory of this company. First, we research BLM and its fuel replenishment methods. Then we define the problem, its objective function, assumptions, parameters and constraints. After identifying the most important given information (fuel price forecast, demand forecast, and current inventory levels), we define the equations that relate these inputs with the order sizes, and the availability and reserve constraints. Due to the large number of constraints, we devise a mechanism to calculate lower limits for the aggregate order sizes that prevent violations of the constraints beyond user-defined limits. We prepare a model in Excel for use with a single fuel type. This model takes stochastic forecasts of demand and fuel prices, and determines the best size for the weekly fuel order. After testing the model under several different scenarios, we conclude that it responds correctly to changes in price and demand. The complete discussion of these results can be found in the body of the thesis. Finally, we present some recommendations for BLM, both in relation to this replenishment problem and to its supply chain in general.

Thesis Supervisor: Stephen Graves
Title: Abraham J. Siegel Professor of Management, Sloan School of Management
Professor of Engineering Systems

to my loving wife, Monica,
and my parents, Tito and Eka,
optimal solution to the complex problem of my existence

Table of Contents

List of Figures, 6

Acknowledgements, 7

1. Introduction, 8

1.1. Thesis overview, 8

1.2. Introduction to BLM Corp, 9

1.2.1. Brief history of BLM Corp, 9

1.2.2. Brief description of BLM's current generation capacity, 10

1.2.3. A word about BLM's future generation capacity, 11

1.3. Why optimize BLM's fuel inventory management? 12

1.4. Literature review, 13

1.4.1. On fuel inventory management for thermal generating plants, 13

1.4.2. On replenishment and inventory management in general, 14

1.4.3. On the economics and operation of power systems in general, 14

2. Understanding the Problem, 16

2.1. Problem Definition, 16

2.1.1. Objective, 16

2.1.2. Decision variables, 16

2.1.3. Constraints, 16

2.1.4. Assumptions, 17

2.1.5. Parameters, 19

2.2. Fuel demand, 22

2.2.1. Forecasting the required generation, 23

2.2.2. Translating generation into fuel consumption, 25

2.3. Fuel prices, 27

2.3.1. Fuel prices structure, 27

2.3.2. Fuel prices variability, 28

3. Building the model, 29

3.1. Equations, objective function and constraints, 29

3.1.1. Nomenclature, 29

3.1.2. Inputs, 32

3.1.3. General equations, 33

3.1.4. Weeks of relevance of the general equations, 34

3.1.5. Cost equations, 34

3.1.6. Objective function, 36

3.1.7. Original constraints, 36

3.1.8. Simplified constraints, 37

3.2. Implementing the model in Excel, 39

3.2.1. The spreadsheet, 40

- 3.2.2. The Macro, 46
- 3.2.3. Solver parameters and options, 49

4. Analysis of Results, 53

- 4.1. Using the model, 53
- 4.2. Example #1, 55
- 4.3. Example #2, 57
- 4.4. Example #3, 59
- 4.5. Example #4, 61
- 4.6. Example #5, 63
- 4.7. Example #6, 68
- 4.8. Example #7, 70
- 4.9. Example #8, 72
- 4.10. Example #9, 74
- 4.11. Unexpected results, 76
 - 4.11.1. No added value in price forecast being stochastic, 76
 - 4.11.2. Reserve violations are not constraining, 76

5. Conclusions and recommendations, 78

- 5.1. Conclusions, 78
- 5.2. Recommendations, 79

Appendix A - The Panamanian power system, 82

- A.1. Origin of the Panamanian power system, 82
- A.2. Privatization of the Panamanian power system, 84
- A.3. Current structure of the Panamanian power system, 85
 - A.3.1 Generation, 85
 - A.3.2 Distribution, 85
 - A.3.3 Retail Sales, 86
 - A.3.4 Transmission, 86
 - A.3.5 System operation, 86
 - A.3.6 Market operation, 87
 - A.3.7 Regulatory Institution, 87

Biographical Reference, 88

Bibliographical Reference, 89

List of Figures

- Figure 1 - Transformations to translate generation into fuel consumption, 25
- Figure 2 – Screenshot of the whole spreadsheet, with labels over the four different areas, 40
- Figure 3 – Screenshot of the Input-Output Area of the model, 41
- Figure 4 – Screenshot of the Price-Related Area of the model, 43
- Figure 5 – Screenshot of the extreme-left of the Demand-Related Area of the model, 43
- Figure 6 – Screenshot of the center-left of the Demand-Related Area of the model, 44
- Figure 7 – Screenshot of the center-right of the Demand-Related Area of the model, 45
- Figure 8 – Screenshot of the extreme-right of the Demand-Related Area of the model, 45
- Figure 9 – Screenshot of the Useful Metrics Area of the model, 46
- Figure 10 – Excel Solver Parameters for the model – First six constraints are shown, 51
- Figure 11 – Excel Solver Parameters for the model – Remaining constraints are shown, 51
- Figure 12 – Excel Solver Options settings for the model, 52
- Figure 13 – Summary of the most important variables of Example #1, 55
- Figure 14 – Graphical representation of demand, price and orders of Example #1, 55
- Figure 15 – Summary of the most important variables of Example #2, 57
- Figure 16 – Graphical representation of demand, price and orders of Example #2, 57
- Figure 17 – Summary of the most important variables of Example #3, 59
- Figure 18 – Graphical representation of demand, price and orders of Example #3, 59
- Figure 19 – Summary of the most important variables of Example #4, 61
- Figure 20 – Graphical representation of demand, price and orders of Example #4, 61
- Figure 21 – Summary of the most important variables of Example #5, 63
- Figure 22a – Graphical representation of demand, price and orders of Example #5, 63
- Figure 22b – Total cost (\$) versus the order size of Q2 (BBL), keeping Q1+Q2 constant, 65
- Figure 23 – Summary of the most important variables of Example #6, 68
- Figure 24 – Graphical representation of demand, price and orders of Example #6, 68
- Figure 25 – Summary of the most important variables of Example #7, 70
- Figure 26 – Graphical representation of demand, price and orders of Example #7, 70
- Figure 27 – Summary of the most important variables of Example #8, 72
- Figure 28 – Graphical representation of demand, price and orders of Example #8, 72
- Figure 29 – Summary of the most important variables of Example #9, 74
- Figure 30 – Graphical representation of demand, price and orders of Example #9, 74

Acknowledgements

Above all, I want to thank God for giving me this wonderful, undeserved opportunity to study among the best.

I want to thank my thesis advisor, Stephen Graves, who patiently guided me through the development of this project.

I also thank the staff of BLM Corp, who generously shared with me the information I required to prepare the model.

I am forever in debt with the Fulbright and Barsa scholarships, which supported my learning experience at MIT.

1. Introduction

1.1. Thesis overview

Bahia Las Minas Corp (BLM) is a generating company in the Panamanian power system. It runs a combined cycle powered by Marine Diesel and three steam turbines powered by Fuel Oil. In the future BLM expects to convert the steam turbines to coal. Every week, the managers of BLM decide on whether to place fuel orders from its supplier, and for how many barrels of each fuel type. These purchases represent tens of millions of dollars per year. Currently, there is no decision-support system to guide the decision-makers as to how many barrels of each type of fuel should be ordered at the time of placing the orders.

Our thesis is about designing a decision-support model for managing these weekly replenishment decisions. In order to avoid the need for new software, the model was designed to run in software with which BLM's staff is already familiar: Microsoft Excel 2002.

In Chapter 1, we introduce the reader to BLM Corp, review its history and its inventory management. We also explain briefly the purpose of a decision-support model for replenishment, and the importance it has for achieving a good inventory management, with high service level at the lowest cost.

In Chapter 2, we explore the replenishment problem, including the objective, variables, assumptions, parameters and constraints considered by the model. Among the variables presented in this section are two stochastic forecasts: fuel price and fuel demand.

In Chapter 3, we define the equations that describe the model, the objective function that is minimized, and the constraints that the model considers. Then we create the model, by defining in an Excel spreadsheet the variables that are involved in the calculation and the equations that relate these variables. The stochastic forecasts for fuel demand and fuel prices will also be included. We use the Excel Solver to find the order sizes that minimize the

relevant costs (storage cost, money cost and fuel cost), while respecting the constraints and considering the stochastic nature of the forecasts.

In Chapter 4, we analyze the performance of the model through nine examples. Finally, in Chapter 5 we summarize our findings and present some recommendations for BLM Corp, both in relation to this replenishment problem and to its supply chain in general.

1.2. Introduction to BLM Corp

In this section, we present an introduction to BLM Corp. A detailed introduction to the Panamanian power system, its origins, privatization and current structure, can be found in Appendix A.

1.2.1. Brief history of BLM Corp

When Panama's IRHE¹ was privatized, one of the generating companies that were created was called EGEMINSA². Later this name was changed to BLM Corp³. The company is named after Las Minas Bay and the port that is located in it, next to the generating units of the company. Fifty one percent of the stock of the company was purchased by Enron while 48.5% remained the property of the Panamanian government. Only 0.5% of the stock was purchased by employees. At the time of the privatization, nine generating units were assigned to BLM:

- 1) Unit 1 (boiler and steam turbine.)
- 2) Unit 2 (boiler and steam turbine.)
- 3) Unit 3 (boiler and steam turbine.)
- 4) Unit 4 (boiler and steam turbine.)
- 5) Unit 5 (gas turbine.)

¹ Institute of Hydraulic Resources and Electrification (Instituto de Recursos Hidráulicos y Electrificación)

² Empresa de Generación Eléctrica Bahía Las Minas, S.A.

³ Bahía Las Minas Corp

- 6) Unit 6 (gas turbine.)
- 7) Unit 7 (engines.)
- 8) Mount Hope (gas turbine.)
- 9) San Francisco (engines.)

Very soon, four of these units were dismantled, because of their obsolescence: Unit 1, Unit 7, Mount Hope and San Francisco. At the same time, BLM Corp inherited from IRHE the plans of a new generating complex. This plan included a new gas turbine (Unit 8) and a new steam turbine (Unit 9): the three gas turbines (Units 5, 6 and 8) would supply their exhaust gas to the new steam turbine (Unit 9) in order to operate the complex as a Combined Cycle. BLM Corp implemented this plan, and the Combined Cycle was operating by 1999. In December 2000, lightning struck Unit 6 and put it out of service for six months.

1.2.2. Brief description of BLM's current generation capacity

Currently, BLM Corp encompasses seven generating units, which we can classify in two groups: the steam turbines and the combined cycle.

Steam turbines

There are three steam turbines (units 2, 3 and 4). Each has an original plate capacity of 40 MW. Currently, their gross capacity is around 35 MW. These units run on Fuel Oil No 6 (Bunker C). Running at full capacity, each steam unit requires around 1,300 barrels of No 6 per day. They burn Light Diesel during start-up.

Combined cycle

The combined cycle includes one steam turbine (unit 9) and three gas turbines (units 5, 6 and 8). Unit 5 has a plate capacity of 33 MW. Unit 6 has a plate capacity of 33 MW. Unit 8 has a plate capacity of 34 MW. Unit 9 has a plate capacity of 58 MW. The combined cycle has a plate capacity of 158 MW, although currently its gross capacity is around 147 MW. The gas turbines run on Marine Diesel No 2. For its use in BLM's turbines, Marine Diesel

includes 80% of Light Diesel and 20% of Fuel Oil No 6. The gas turbines burn Light Diesel during start-up. The complex can operate in six different configurations:

- 1) Unit 5 can operate as an independent gas turbine, with a gross capacity of 32 MW.
- 2) Unit 6 can operate as an independent gas turbine, with a gross capacity of 32 MW.
- 3) Unit 8 can operate as an independent gas turbine, with a gross capacity of 33 MW.
- 4) One gas turbine can operate with Unit 9 as a combined cycle, with a gross capacity of approximately 49 MW.
- 5) Two gas turbines can operate with Unit 9 as a combined cycle, with a gross capacity of approximately 98 MW.
- 6) Three gas turbines can operate with Unit 9 as a combined cycle, with a gross capacity of approximately 147 MW. Running at full capacity, the combined cycle requires around 4,500 barrels of No 2 per day.

Thus, the total plate capacity of BLM is 278 MW, which is 20% of the total 1.38 GW of installed capacity available for the Panamanian system, and 33% of the current maximum demand of the country (0.85 GW).

1.2.3. A word about BLM's future generation capacity

Since the beginning, BLM Corp has considered the use of less expensive fuels. One of the ideas the company considered at the beginning was to run the combined cycle on natural gas. However, since Panama has no wells of natural gas, it would have to be brought from abroad. Liquefied Natural Gas was too expensive. Enron had planned to invest \$300 million in an underwater natural gas pipeline from Cartagena, Colombia to the Bahía Las Minas plant, 3 feet in diameter, almost 600 kilometers long. The pipeline would initially transport 70 million cubic feet (MCF) per day. However, an increase in the projected cost of gas rendered the project unprofitable for the investors, and it was scrapped.

The current plan is to build a new, 120 MW fluidized-bed boiler to burn pulverized coal. This boiler will supply steam to Units 2, 3 and 4. At a cost of 100 million dollars, this boiler should be ready by July 2006. The company estimates that using coal instead of Fuel Oil will cut the variable cost of these units by half, which will increase BLM's ability to compete in the Panamanian generation market.

The coal BLM plans to use, which would be brought from Colombia, has a sulfur content of 0.7%. This represents a gain from an environmental perspective, because the Fuel Oil currently in use has 3.0% content of sulfur. The new boiler will also be furnished with pollution-abatement technology. The terminal to receive coal in the Las Minas bay port was inaugurated in early February 2004. BLM has considered keeping the capacity to burn either coal or fuel oil in this boiler, which would allow the company to choose the fuel that is less expensive at any given time.

1.3. Why optimize BLM's fuel inventory management?

Generally speaking, the goal of inventory management optimization is to reduce the relevant costs, including purchase cost and holding cost, while providing the desired product availability.

For an electricity generating company, the goal of fuel inventory management is to balance the cost of purchasing fuel and holding it in inventory against the risk of not having enough fuel available to satisfy demand in real time. In the case of BLM, most uncertainty is associated with fuel demand, although fuel supply is not necessarily certain. BLM must keep fuel inventory at proper levels to avoid running out of fuel, without building excessive inventory levels.

Also, since fuel prices are not fixed, it is possible to take advantage of price changes. When fuel price is expected to increase substantially, it is desirable to build a stock of fuel

purchased at the present low prices, to avoid the purchase of more expensive fuel in the near future.

Thus, the purpose of optimizing BLM's fuel inventory management is to prevent unnecessary fuel purchases and inventory accumulation, to take advantage of variations in fuel prices, and to ensure the availability of fuel, both stored and ordered, to satisfy availability and reserve expectations, at the lowest possible total cost.

1.4. Literature review

1.4.1. On fuel inventory management for thermal generating plants

Some of the publications that discuss topics similar to that of our work include the following:

"A Utility Fuel Inventory Model", published by Peter A. Morris et al. in *Operations Research*, Vol. 35, No. 2 (Mar.-Apr., 1987), 169-184. This paper discusses an inventory modeling system, called "The Utility Fuel Inventory Model, or UFIM, which was designed to help electric utilities set a long-term fuel inventory strategy that specifies the most cost-effective inventory levels to maintain during normal times (e.g. times when there are no disruptions). Morris's paper describes the model, its use and one of its first applications. He mentions that UFIM was used at that time by more than 50 utilities.

"A System Integration and Optimization Model for Fuel Management and Scheduling of Power Generators", doctoral thesis of Babul Patel, presented in 1981 at the University of Nebraska. Patel examines the problem of scheduling power generators and develops a dynamic programming technique to obtain an optimal dispatch schedule. The final results of his analysis include weekly or monthly summaries of the unit's fuel inventory and fuel costs.

"Studies in fuel supply and air quality planning by electric utilities", doctoral thesis of Yung-Tang Shen, presented in 1996 at Ohio State University. Shen develops a

comprehensive linear programming model for air quality and fuel management, dealing with fuel purchases, consumption, inventory, and air quality control strategies. He considers, among other variables, inventory requirements and electricity demand. His objective was to minimize total fuel purchase, inventory, and air quality control costs over a given planning horizon.

1.4.2. On replenishment and inventory management in general

The literature on replenishment and inventory management is vast and specialized. Our model uses rather basic concepts. For the reader that is not familiar with the general concepts of inventory management and wants to acquire more knowledge on this area, we list the names of two books that will prove helpful.

The first is "Inventory Management and Production Planning and Scheduling" by Edward A. Silver, David F. Pyke and Rein Peterson. This book was repeatedly recommended to us during the MLOG program. It tries to close the gap between the theoretical advancements and the industrial applications of inventory management.

Another excellent book, which deals with Supply Chain Management as a whole, is "Supply Chain Management", by Sunil Chopra and Peter Meindl. The replenishment techniques described in this book for unique products with stochastic demand (Chapter 9) were very helpful for developing our model, since it considers stochastic variables.

1.4.3. On the economics and operation of power systems in general

Although power systems are not the core subject of our thesis, we realize that some of the readers of this document will be interested in learning more about them. There is a vast literature regarding the operation, economics and risk management of power systems. We want to recommend four of them.

As an introduction to the engineering and economic factors involved in operating and controlling power generation systems, our recommendation is "Power Generation,

Operation and Control" by Allen J. Wood and Bruce F. Wollenberg. This book covers several topics that we mention in Appendix A, such as economic dispatch of thermal units, unit commitment, fuel supply contracts, and hydrothermal coordination.

For the reader interested in power market, we recommend two excellent books. The first is "Power System Economics", by Steven Stoft, which introduces key economic, engineering and market design concepts for power systems, and explores in more depth specific areas of interest in power markets. The second book is "Market Operations in Electric Power Systems", by Mohammad Shahidehpour, Hatim Yamin and Zuyi Li, which discusses market structure, operation, forecasting, scheduling, and risk management in electric power systems. It dedicates one whole chapter to one of the topics we mention in Appendix A: short-term load forecasting.

Readers interested in exploring the risk management techniques available to power companies are invited to read "Energy and Power Risk Management", by Alexander Eydeland and Krzysztof Wolyniec. Although not related to the core subject of this thesis, risk management techniques such as modeling, pricing and hedging, can complement our inventory management approach to protect generating companies from fuel prices volatility.

2. Understanding the Problem

2.1. Problem Definition

2.1.1. Objective

We can define the problem as follows: create in Excel an inventory optimization model for BLM that determines, for a single fuel type, the size of next week's fuel order, so that the purchase and holding costs are minimized considering a horizon of four weekly orders into the future, while respecting all constraints and assumptions.

As we stated before, the way we will create this model is by defining in an Excel spreadsheet the relevant variables (including the stochastic forecasts for fuel demand and fuel prices) and the equations that describe the model. We will not write our own optimization code. Instead, we will use Excel Solver to find the order sizes that minimize the relevant costs (storage cost, money cost and fuel cost), while respecting all constraints.

2.1.2. Decision variables

The decision variables are the only variables that can be adjusted by the model to minimize the objective function. These variables are the size of the weekly orders for the next four weeks, which are the relevant orders within the horizon of the model. The sizes of the orders are subject to the constraints of the model.

2.1.3. Constraints

While minimizing the objective function, the model should consider the following constraints:

1. For each week in the horizon of the model, order sizes cannot be negative.
2. For each week in the horizon of the model, order sizes cannot be larger than a maximum size specified by the user. The purpose of this upper cap is to account for the limitations in:

- a. The financial capacity of BLM to purchase fuel.
 - b. The fuel availability of the supplier.
3. For each week in the horizon of the model, the probability of a stockout should be no more than the maximum acceptable probability specified by the user. This assumption is based on the significant penalties of stockout.
 4. For each week in the horizon of the model, the probability of fuel reserve levels falling below the expected reserve level should be no more than the maximum acceptable probability specified by the user. Fuel reserve is not a hard constraint (as proven by the fact that, in times of shortage, the reserve has been lower than the expected reserve levels), but it is considered a best-practice criterion that should be respected whenever possible.
 - a. Fuel reserve is defined as the sum of the fuel stored both in the local and external storages plus the fuel of orders that have already been placed (even if this fuel is not physically “in transit”).
 - b. Expected reserve is defined as the fuel needed for the next 10 days of forecasted generation.

If the model is asked to decide the order sizes so that there is no probability of availability or reserve violations, the model will make its replenishment decision based on the scenario with the greatest demand. This results in the model recommending very high order sizes, which is not desirable. Accepting a probability of violations greater than zero tempers these results. That is why the model will ask the user to provide values for the maximum accepted probabilities for violations of availability and reserve for each week.

Please notice that there are no constraints related to fuel prices.

2.1.4. Assumptions

Our model will use as working assumptions the following statements:

1. The cost of placing an order is negligible. The rationale for this assumption is that the only variable cost of placing an order is the cost of the bank transfer, which is negligible compared to the amount paid for the fuel itself. Other costs, such as those related to staff and fuel transportation, do not depend on the number of orders placed.
2. We assume that one order for each fuel type is placed every week, and that the size of each order is decided on Friday afternoon⁴. Of course, BLM can decide to place an order for 0 barrels (e.g. to place no order) on any week.
3. We assume that the first barrel of an order arrives one week later, and the last barrel two weeks later, counting in both cases from the Friday the size of the order was decided.
4. Once an order is placed, it cannot be modified. This assumption is based on information we received from BLM describing its relationship with the supplier.
5. The shipping capacity of the supplier is not considered a constraint. The receiving capacity of BLM is discussed in section 2.2.
6. The storage capacity of the supplier is not considered a constraint. The storage capacity of BLM is discussed in section 2.2.
7. All orders of both Marine Diesel and Fuel Oil are delivered by truck.
8. There is no correlation between short-term variations in fuel price (e.g. variations in price in the following four weeks) and variations in fuel consumption. This assumption is based on the experience of BLM staff, and is supported by the fact that in Panama, most thermal generating units receive their fuel from the same markets: the variations in fuel prices are felt almost equally by all generating companies. Also, the percent of hydroelectric

⁴ In real life, BLM places the orders any day of the week, but to keep the model simple, we assume that orders are placed only on Friday afternoon.

generation at any given time depends on long-term forecasts of fuel prices, but is not affected by variations of fuel prices in the short run. Therefore, short-term variations in fuel prices have no noticeable effect over unit commitment and required generation.

9. At any given time, BLM knows the current fuel inventory levels. BLM has no real-time reading of the fuel levels in its tanks. However, these levels are measured directly in the tanks once a month, and between the readings the levels in the tanks are estimated based on the actual generation and heat rate of the units, and the heating values of the fuel.
10. BLM decides on Friday the size of the order for each fuel type for next week only. However, the planning horizon for the model includes the demand forecast for the next six weeks and the order sizes for the next four weeks. This planning horizon is a compromise between anticipated planning and forecast reliability:
 - a. On the one hand, using more than four weekly orders as decision variables would require using more than six weeks of demand forecast. This is not desirable, because forecasts for demand so far in the future are not reliable.
 - b. On the other hand, using less than four weekly orders as decision variables would severely impair the model's ability to play with inventory levels for the company's benefit, i.e. to purchase excess fuel inventory at lower prices in the present to avoid purchasing fuel later at higher prices.

2.1.5. Parameters

Our model will consider the following parameters, which are presented here with the values of the current supplier. In the future, if these values change, they can be updated in the model.

1. Number of suppliers: Currently BLM has only one fuel supplier.
2. Demand: A stochastic fuel demand forecast is used as input to our model. By *stochastic demand forecast*, we mean a forecast that has, for each one of the following six weeks, fifty possible demand scenarios, each with a probability of 2%. This demand forecast is calculated by BLM based on a stochastic optimization run by the system operator (National Dispatch Center⁵ or CND), in a computer program called SDDP. In section 2.2, we discuss how this demand forecast is prepared.
3. Fuel prices: A stochastic fuel price forecast is used as input to our model. By *stochastic price forecast*, we mean a forecast that has, for each of the following four weeks, five possible price scenarios, each with a probability of 20%. This price forecast is calculated by BLM's staff. Later in this chapter, in section 2.3, we discuss the structure and variability of fuel prices. Fuel price information of BLM is confidential. As we mention in Chapter 4, the same results are achieved by using either a deterministic or a stochastic price forecast.
4. Local storage capacity: BLM has a local storage capacity of 50 thousand barrels of Fuel Oil and 50 thousand barrels of Marine Diesel. As a reference, consider that when the generating units of BLM are operating at full capacity, they require each day around 4 thousand barrels of Fuel Oil for all the three steam units and 4.5 thousand barrels of Marine Diesel for the Combined Cycle. However, very seldom are all units generating flat out. Local storage capacity is a sunk cost. This cost is not relevant for our calculations.
5. External storage capacity: Theoretically, it is possible for BLM to lease additional storage capacity from its supplier, for a fee. Since BLM has never used this service before, there is no information about the exact cost of this service. For

⁵ National Dispatch Center (Centro Nacional de Despacho)

the sake of our model, we have assumed a cost of \$2/BBL-week, that is, two dollars per barrel stored per week stored, paid on the maximum amount of barrels stored each week (not on the average amount of barrels stored.) As we will explain later, for our model weeks start on Saturday.

6. Cost of money: The holding cost is composed of the cost of external storage capacity and the cost of money. BLM informed us that its cost of money can be expressed as LIBOR⁶ plus a confidential constant. For our model, we estimated that the money cost is 25% per year, or $25\% / (52 \text{ weeks/year}) = 0.481\%$ per week.
7. Order lead time: Order lead time depends on the size of the order. According to BLM, for an order of 25 thousand barrels of one type of fuel, from 10 to 12 days elapse from the moment the size of the order is defined to the moment the tank trucks are received at the plant. This delay can be divided into three periods:
 - a. Three days elapse from the moment the size of the order is defined to the moment the bank emits the purchase order. This delay does not depend on the size of the order.
 - b. Between 3 to 4 days elapse from the moment the purchase order is emitted and the moment the supplier sends the first loaded tank truck. Part of this time is used to get a tax exemption for the fuel.
 - c. The rest of the delay depends on the size of the order. In the past, the average order of BLM has been of 25,000 BBL of each fuel. For an order of this size, 4-5 days elapse from the moment the first truck leaves the supplier's premises and the moment the last truck arrives at BLM's plant.

To account for the fact that orders can be larger than 25,000 barrels, and to compensate for any unforeseen authorization, processing or delivery delay, we will consider the

⁶ London Interbank Offered Rate, a widely used reference rate for short term interest rates.

whole lead time to be two weeks (14 days) from the day the order is decided to the day the last tank truck arrives at BLM's plant.

8. Based on the structure of the delays, we make the following assumptions:
 - a. For the sake of calculating the cost of money, we assume that BLM takes financial responsibility of the fuel (e.g. must have separate funds for them) from a day after the amount of the order has been decided. Since the sizes of the orders are decided on Friday, BLM would pay the cost of money starting the following day, Saturday, which is the first day of the new week.
 - b. For the sake of calculating the cost of storage, we assume that BLM must assume physical responsibility for the fuel starting one week after the day the supplier has cleared this fuel for delivery. Speaking in terms of days after the Friday the order size is determined: the fuel is BLM's responsibility 2 weeks after the order has been placed.
9. Receiving capacity: BLM can receive up to 30 tank trucks per day for each type of fuel, each truck with an average of 202 BBL of fuel. This amounts to a maximum receiving capacity of 6,060 thousand barrels per day for each type of fuel. Since BLM's fuel consumption at full capacity is smaller than its receiving capacity, the receiving capacity is not a constraint for our decisions.
10. We assume that demand is uniform during the week. It is a known fact that generation on Saturdays and Sundays is usually lower than on weekdays, but in order to keep our model reasonably simple we will neglect this difference.

2.2. Fuel demand

Forecasting the fuel consumption of BLM is a complex process, since it depends on a variable that is very hard to forecast: the required generation of BLM.

For our model, we do not attempt to prepare our own fuel consumption forecasts. Instead, we use the forecasts prepared every week by the system operator. Section 2.2.1 explains how the system operator forecasts the required generation of BLM's units. Section 2.2.2 explains how these generation forecasts are translated into fuel consumption forecasts.

2.2.1. Forecasting the required generation

As we mentioned in section 1.2.3., the required generation of BLM depends, generally speaking, on two factors: unit commitment and system-wide energy demand. System-wide energy demand occurs in real-time, and has been found to depend on several factors, including: 1) the amount of installed load in the system (nation-wide), 2) the time of the day, 3) the day of the week, 4) whether a day is a holiday, 5) the week of the year, and 6) weather (fresh, rainy days have lower demand than hot, sunny days). Unit commitment is determined by the system operator (CND), who decides to commit a unit depending on many factors, including but not limited to the following: 1) the availability of the generating unit, 2) the variable cost of all the available generating units in the system, 3) long-term weather forecast (in a rainy year hydroelectric units are used more than in dry years), 4) the long-term fuel prices forecast, 5) the expected demand for the day, the week, the month, and the next two years, 6) the start-up cost and operation restrictions of the generating unit, and 7) the stability and safety of the operation of the system.

Obviously, trying to forecast the required generation based on expected energy demand and optimal unit commitment is a very complex task and is beyond the scope of our model. Instead, we will use the forecasts prepared every week by the system operator, CND, with the help of all the stakeholders of the system. These forecasts are used to determine many things, including the operative planning of the system for the short and long terms. They are highly regarded as accurate forecasts for the short-term operation, because they are prepared with the active collaboration of all parties involved, including representatives of the staff of BLM Corp and all the other generating companies.

Every Thursday, before 10:00 AM, the generating, transmission and local distribution companies give CND large amounts of information regarding the status of their equipment, including scheduled maintenances and other relevant data. For example, LDCs give CND the energy demand forecasts for their networks; thermal generators give CND their generation variable cost for the week, based on the current fuel prices; weather experts forecast inflows to the reservoirs for the next two years; ETESA gives CND the information on scheduled transmission interruptions because of network maintenance.

With this information, and the use of specialized software called Stochastic Dual Dynamic Programming (SDDP), CND performs three analyses. Each one of these analyses provides as output a large amount of information, including a forecast of the generation that will be required from the units of BLM Corp. The three analyses are:

1) **The Stochastic Run.** This run uses 3 years of data. It considers reservoir inflows (the amount of water that enters the reservoirs) stochastically, and prepares fifty hydrologic scenarios based on historic data from several decades. SDDP uses dynamic programming to determine the optimal prices (\$/MWh) that should be assigned to the hydroelectric generating units in order to balance their generation, on an economic dispatch basis, with the generation of thermal units. This run provides us with a stochastic forecast of BLM's required generation for every week of the next three years, in GWh. The forecast is stochastic because it considers 50 possible weather scenarios, which translates into 50 possible generation scenarios for BLM, each with a probability of 2%. History teaches that the forecasts of the first four weeks are very accurate.

2) **The Week-Ahead⁷ Deterministic Run.** CND prepares its operative policy using the optimized hydroelectric generation prices obtained from the first run. Considering both these prices and the thermal prices, CND performs a short-term analysis, hour by hour, to determine unit commitment for the next week (from Saturday to Friday). This unit

⁷ In Panama, this run is commonly called "el Predespacho"

commitment will be used as reference in real-time for the economic dispatch of the system, constrained by safety and stability considerations. This run provides a deterministic forecast of the BLM's required generation for every hour of the next week, in MWh.

3) **The Month-Ahead⁸ Deterministic Run.** CND performs a third analysis, using a horizon of two years, with weekly steps, considering weather as a deterministic variable. In the words of Percy Garrido, BLM's senior analyst in charge of SDDP operation, the purpose of this run is to *"provide a robust forecast of the next four weeks"*. The forecast is not considered accurate beyond the fourth week. This run provides us with a deterministic forecast of BLM's required generation for every week of the next month, in GWh.

Every Friday morning, representatives from all the stakeholders of the power system, including BLM, gather in an office at CND's headquarters and discuss the results of these three runs. As a result of this weekly meeting, any discrepancies and errors that are detected in the runs prepared on Thursday will be corrected by CND on Friday noon. Thus, every Friday afternoon all the companies (including BLM) have access to the final forecasts from the three runs of SDDP performed by CND and verified by the all parties.

2.2.2. Translating generation into fuel consumption

For the sake of inventory management, demand should be expressed in terms of fuel units (e.g. BBL), instead of electrical energy (e.g. MWh). It is possible to translate the forecast of required generation into one of fuel consumption. As a matter of fact, the SDDP software makes this automatically. Let us review the logic behind this transformation. The following figure illustrates the calculation process.




Figure 1 - Transformations to translate generation into fuel consumption.

⁸ Commonly known in Panama as "la Corrida Determinística"

Generation is given in Watts-hour or its multiples, such as MWh or GWh. For the purpose of fuel replenishment, it is useful to express fuel consumption in the units used in the purchase orders. In the case of liquid fuels, such as Fuel Oil and Marine Diesel, these units are either gallons or barrels (BBL). In the future, when coal is used in BLM, the consumption of this fuel will likely be expressed in tons.

Energy is the link between generation and fuel consumption. In the Panamanian power system, energy is commonly expressed in terms of MBTU (millions of BTU).

The variable that allows us to determine how much energy is needed to obtain a certain generation from a specific generating unit is known as Heat Rate, but is also called “Efficiency”, and is expressed in terms of MBTU/MWh. Each generating unit of BLM, even each configuration of the Combined Cycle, has an efficiency curve, which expresses the efficiency of the unit at any given generation level.

The variable that allows us to determine how much energy is contained in a unit, i.e. barrel or ton, of a given fuel is the Heating Value, which can be either gross or net. In Panama, it is commonly expressed in MBTU/BBL.

Every week, CND inputs into SDDP the values of Heat Rate for every generating unit of the system, including the three steam units and all the configurations of the combined cycle of BLM. These values are updated every year or so, after efficiency tests made by third-party engineering firms. CND also inputs into SDDP the Heating Values for the specific fuel that every generating unit is using in a given week. This allows SDDP to translate, through simple calculations, the forecast of required generation into a forecast of fuel consumption. After the Stochastic Run, the SDDP software returns its fuel consumption forecast in a file called **fuelcn.csv**. This file contains, among other values, the two stochastic forecasts of fuel consumption for BLM: one for Fuel Oil and one for Marine Diesel. For each fuel, for each one of the following 52 weeks, SDDP forecasts 50 possible consumption scenarios (divided

in four blocks), each with a probability of 2%. These forecasts are expressed in thousand of gallons of consumption per week (not in barrels).

However, it is necessary to make an adjustment to this forecast before we use it in our model. SDDP calculates the forecast considering that the generating units operate all the time at the maximum efficiency point of their efficiency curves. However, in real life they sometimes operate at lower efficiency points that consume more fuel for the same generation output. BLM's Garrido has told us that by augmenting the forecasted values of consumption by 5% we can compensate for this deviation.

Our model uses this stochastic forecast, with the 5% efficiency adjustment and transformed to barrels, as input for its calculations. *Note: blocks are not used individually.*

2.3. Fuel prices

2.3.1. Fuel prices structure

Confidentiality prevents BLM from sharing with us the fuel prices of the contract with its current supplier.

However, they shared with us the structure of the formula that determines these prices: they are calculated as the monthly average of the prices of the fuel in the US Gulf according to Platt's *Latin Wire*, plus a fixed amount that was not disclosed. Unfortunately, the information contained in Platt's Latin Wires is available only to subscribers of this service and cannot be shared with non-subscribers.

When the model is in use in the future, BLM will provide as an input a stochastic price forecast for each fuel type. By stochastic price forecast we mean a forecast that has, for each of the following four weeks, five possible price scenarios, each with a probability of 20%.

For the purposes of our thesis, we will use arbitrary fuel values between \$23/BBL and \$27/BBL, which, based on my experience in BLM, are typical values.

2.3.2. Fuel prices variability

To develop our model, we do not need to know the exact values of fuel prices. However, to be able to prepare examples to evaluate our model, it helps if we know the magnitude of variation that these fuel prices usually display. Based on our experience in BLM, and using as a reference publicly available WTI oil prices for the last 12 months, we have estimated that in the absence of high-profile world scale events, it is common for BLM to suffer changes of fuel price of 8% from one month to the next. To estimate how extraordinary events, such as the recent wars in Iraq and Afghanistan, can increase the price variations, we used data from February and March 2003 (near the beginning of the latest war in Iraq). Prices in these months show that the variability could reach values of up to 18% under these circumstances.

3. Building the model

In section 3.1 we define the equations that describe our model. In section 3.2 we simplify the constraints of the model by defining aggregate lower limits for the order sizes. In section 3.3 we describe the development of the Microsoft Excel model.

While you read this chapter, please keep in mind that the model we develop works with a single type of fuel, which can be either Fuel Oil or Marine Diesel. Therefore, BLM would have to run two copies of the model, one for each fuel. If the “barrels” labels are replaced with “tons” labels, the model can be used for coal, too.

3.1. Equations, objective function and constraints

3.1.1. Nomenclature

In this section we define the nomenclature that we will use later to write the equations that describe the problem, in three groups: 1) general nomenclature, 2) nomenclature of input variables, and 3) nomenclature of calculated variables. The text inside the parenthesis indicates the units this variable will have in our model.

As you read this section, keep in mind that the demand forecast has 50 stochastic scenarios and the price forecast has 5 stochastic scenarios.

3.1.1.1. Nomenclature – General nomenclature

- **Weeks** in the Panamanian power system start on Saturday and end on Friday. This type of week is called “*planning week*” or “*Titistic week*”, and is used in our model because it is used in the input data we receive from SDDP.
- **n** is an integer representing the number of weeks considered in the horizon of the model. In our model, n is four (4) weeks.
- **i** is an integer that represents every given week, from week 1 to week n.

- j is an integer that represents every given demand scenario, from 1 to 50, of the stochastic demand forecast, each with a 2% probability.
- k is an integer that represents every given price scenario, from 1 to 5, of the stochastic price forecast, each with a 20% probability.
- Q_i is the amount of fuel purchased in the order of week i (BBL). The decision of the exact size of this amount is made on Friday of week $i-1$. According to our assumptions, this amount is part of BLM's fuel reserve since the first day of week i , but is physically in BLM's storage starting in the first day of week $i+2$. The sizes of the past two orders, Q_{i-1} and Q_0 , are input variables that should be provided by BLM. On the other hand, the order sizes from weeks from Q_1 to Q_n are decision variables, and are identified collectively in this thesis as \bar{Q} .

3.1.1.2. Nomenclature – Input variables

- LC is the local fuel storage capacity of BLM (BBL). Storing fuel locally implies no relevant cost for our calculation, since there is no variable cost associated with it.
- D_0 is the fuel demand in week zero. This is a single value from the past provided by BLM as input.
- D_{ij} is the fuel demand in week i in demand scenario j . These values originate in the 50 scenarios of CND's Stochastic Run of SDDP, and are later transformed to barrels and increased by 5% to be used as input values to the model. As we explained at the end of section 2.2.2, the 5% adjustment compensates for the fact that efficiency under real generation conditions is lower than the optimal efficiency used by SDDP for its forecast.
- CF_{ik} is the cost of the fuel in week i in price scenario k (\$/BBL)
- CF_i is the average cost of the fuel in week i for the five price scenarios (\$/BBL)
- CM is the cost of money expressed as a percent per week (%).

- CE is the cost of leasing the external storage space for a week (\$/BBL-week). We assume the cost is charged to the maximum amount of fuel stored each week i .
- MQ_i is the maximum or upper cap to the value of Q_i (BBL).
- MVA_i is the maximum acceptable probability of violation of the availability constraint (e.g. stockout) in week i , specified by the user (%). For our model, we assumed 20% for all weeks.
- MVR_i is the maximum acceptable probability of violation of the expected reserve constraint in week i , specified by the user (%). For our model, we assumed 20% for all weeks.

3.1.1.3. Nomenclature – Calculated variables

- TS_{ij} is the maximum amount of fuel that is stored both locally and externally in week i in demand scenario j (BBL). It is also the maximum requirement for fuel storage, and occurs at the beginning of week i . TS_0 is the maximum amount of fuel stored in week 0.
- LS_{ij} is BLM's maximum requirement for local fuel storage in week i in demand scenario j (BBL). This value occurs at the beginning of week i . The purpose of this variable is to simplify the calculation of ES_{ij} .
- ES_{ij} is BLM's maximum requirement for external fuel storage in week i in demand scenario j (BBL). This value occurs at the beginning of week i .
- A_{ij} is the minimum amount of fuel that is physically available for generation in week i in demand scenario j (BBL). This value occurs at the end of week i . The minimum of fuel that is locally available is calculated accounting for the fact that as soon as 7 days after the order is placed, fuel is arriving at the plant in tank trucks.
- VA_{ij} is a variable created to know if the value of A_{ij} is violating the “no stockout” constraint, e. g. that at any given time the available fuel A_{ij} to the plant should be more

than zero. A violation is indicated by a value of $VA_{ij} = 1$, while no violation is represented by 0.

- R_{ij} is the minimum reserve of fuel inventory in week i in demand scenario j (BBL). This value occurs at the end of week i . Fuel reserve is understood as the sum of the fuel stored both in the local and external storages plus the fuel of orders that have already been placed.
- ER_{ij} is the minimum expected fuel reserve that the regulations expect BLM to have, equal to the forecasted demand of the next ten days (BBL).
- VR_{ij} is a variable created to know if the value R_{ij} is violating the minimum expected reserve constraint ER_{ij} , e. g. that at any given time the fuel reserve R_{ij} should be more than the minimum expected reserve ER_{ij} . A violation is indicated by a value of $VR_{ij} = 1$, while no violation is represented by 0.
- FC is the expected cost of fuel over the entire horizon of the model (\$).
- SC is the expected cost of external storage over the entire horizon of the model (\$).
- MC is the expected cost of money over the entire horizon of the model (\$).
- TC is the sum of the expected relevant costs: simply the sum of FC , SC and MC (\$).

3.1.2. Inputs

Let us define week 1 as the week for which we need to determine the order size. Our model is designed to be run on the last day of week 0 (Friday), to decide the size of order Q_1 . In this section we list the information that our model uses as input, which should be provided by BLM before running the model.

- LC : The local storage capacity, (BBL).
- CM : The cost of money per week, (%-week).
- CE : The cost of leasing external storage space, (\$/BBL-week).

- Q_{-1}, Q_0 : The amounts of the last two orders (BBL). Each of this is a single value.
- TS_0 : The maximum amount of fuel that BLM had stored both locally and externally in week zero (BBL). In other words, it is how many barrels of fuel BLM had stored both locally and externally last Saturday (the first day of week 0).
- CF_{ik} for $i=\{1,4\}$ and $k=\{1,5\}$: the five stochastic forecasts of prices for the following four weeks, starting in week 1 (\$/BBL).
- D_0 : The demand of week zero (BBL). On Friday afternoon, this demand is known for the last 6 days, and only the demand of the last day has to be estimated. D_0 is a single value.
- D_{ij} for $i=\{1,6\}$ and $j=\{1,50\}$: the demand forecast for the 50 stochastic scenarios of the next six (e.g. $n+2$) weeks, starting in week 1 (BBL).
- MQ_i for $i=\{1,4\}$: the upper cap to the value of Q_i for $i=\{1,4\}$ (BBL).
- MVA_i for $i=\{1,4\}$: the maximum acceptable probability of availability violation (%).
- MVR_i for $i=\{1,4\}$: the maximum acceptable probability of reserve violation (%).

3.1.3. General equations

The following equations define the variables of the model that are not given as input.

$$TS_{ij} = TS_{i-1,j} + Q_{i-2} - D_{i-1,j}$$

$$LS_{ij} = \text{Min}(TS_{ij}, LC)$$

$$ES_{ij} = TS_{ij} - LS_{ij}$$

$$A_{ij} = LS_{ij} - D_{ij}$$

$$VA_{ij} = \begin{cases} 1, & A_{ij} < 0 \\ 0, & \text{otherwise} \end{cases}$$

$$R_{ij} = (TS_{ij} + Q_{i-1} + Q_i) - D_{ij}$$

$$ER_{ij} = D_{i+1,j} + \frac{3}{7}D_{i+2,j}$$

$$VR_{ij} = \begin{cases} 1, & R_{ij} < ER_{ij} \\ 0, & otherwise \end{cases}$$

$$CF_i = \frac{1}{5} \sum_{k=1}^5 CF_{ik}$$

3.1.4. Weeks of relevance of the general equations

The order amounts from Q_1 to Q_n impact different variables in different time frames. This means, for example, that variables such as LS_{ij} are not relevant, while $TS_{n+2,j}$ is relevant. The following is a list of the weeks for which calculating the variables is relevant, assuming we want the model to decide the sizes of orders Q_i for $i=\{1,4\}$.

- TS_{ij} is relevant for $i=\{1,6\}$. Notice that TS_{1j} is the same for every j .
- LS_{ij} is relevant for $i=\{3,6\}$.
- ES_{ij} is relevant for $i=\{3,6\}$.
- A_{ij} is relevant for $i=\{3,6\}$.
- VA_{ij} is relevant for $i=\{3,6\}$.
- R_{ij} is relevant for $i=\{1,4\}$.
- ER_{ij} is relevant for $i=\{1,4\}$.
- VR_{ij} is relevant for $i=\{1,4\}$.
- CF_i is relevant for $i=\{1,4\}$.

3.1.5. Cost equations

The following formulas describe the cost equations that will be calculated.

3.1.5.1. Cost of fuel (FC)

The following equation defines the average cost of fuel for the five price scenarios over the entire horizon of the model:

$$FC = \frac{1}{5} \cdot \sum_{i=1}^4 \sum_{k=1}^5 Q_i \cdot CF_{ik}$$

It is possible to simplify this expression by using the average cost of fuel for week i, CF_i:

$$FC = \sum_{i=1}^4 Q_i \cdot CF_i$$

3.1.5.2. Cost of the money (MC)

The following equation defines the cost of money over the entire horizon of the model:

$$MC = \frac{1}{5} \cdot \sum_{i=1}^4 \sum_{k=1}^5 \left(CM \cdot CF_{ik} \cdot \left(\frac{1}{50} \cdot \sum_{j=1}^{50} R_{ij} \right) \right)$$

It is possible to simplify this expression by using the average cost of fuel for week i, CF_i:

$$MC = CM \cdot \sum_{i=1}^4 \left(CF_i \cdot \left(\frac{1}{50} \cdot \sum_{j=1}^{50} R_{ij} \right) \right)$$

3.1.5.3. Cost of external storage (SC)

The following equation defines the cost of external storage over the entire horizon:

$$SC = \frac{1}{50} CE \cdot \sum_{i=1}^4 \sum_{j=1}^{50} ES_{i+2,j}$$

The reason this equation uses i+2 instead of i is that the need to store the fuel purchased in week i will materialize two weeks later, in week i+2, when BLM receives the material responsibility for the fuel.

3.1.6. Objective function

The objective function that our model will minimize is the total relevant cost, TC, by changing the order quantities, \bar{Q} :

$$\textit{Minimize } TC(\bar{Q}) = FC + SC + MC$$

s.t. constraints

We do not state the constraints here directly, because their complexity makes them worthy of a separate discussion. The original statement of the constraints is presented in section 3.1.7, while a simplified statement is presented in section 3.1.8.

3.1.7. Original constraints

Here is the original statement of the constraints, based on the constraints presented verbally in section 2.1.3:

$$\textit{s.t. } 0 \leq Q_i \leq MQ_i \quad \forall i$$

$$MVA_i \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VA_{ij} \quad \forall i$$

$$MVR_i \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VR_{ij} \quad \forall i$$

The first expression states that the order size for each week must be between 0 and the upper cap specified by the user for that week.

The second expression states that the probability of availability violation for each week should not be more than the maximum probability specified by the user for that week.

The third expression states that the probability of reserve violation for each week should not be more than the maximum probability specified by the user for that week.

Since i has 4 values and j has 50 values, there are 200 (i,j) pairs. Therefore, those three statements imply the existence of 8 individual constraints related to order size, 200 individual constraints related to availability and 200 individual constraints related to reserve, for a grand total of 408 constraints.

If the model is asked to consider each one of the 408 individual constraints required by the problem statement, the problem becomes too complex for Excel Solver. Tests with an early prototype of the model demonstrated that not even academic-strength Excel plug-ins for optimization could solve the problem stated like this.

Therefore, it was necessary to try a new approach to this issue, redefining the constraints in simpler terms. The simplification process and the new constraint statements are presented in the next section.

3.1.8. Simplified constraints

Analysis of the relationship between the demand scenarios and constraint violations demonstrated that the first violation of the availability or reserve constraints occurs in the scenario with the highest demand.

Subsequent violations of the constraints follow the same logic: the second violation to the constraints occurs in the demand scenario with the second highest demand, and so on.

Further analysis demonstrated that it is possible to find lower limits for the aggregate of the order sizes that will guarantee that the number of availability and reserve violations will not exceed the maximum allowed by the user. We claim that:

1) It is possible to find the minimum value of Q_1 for which the number of availability and reserve violations in week 1 does not exceed the maximum allowed by the user. For any value of Q_1 below this value, the violations in week 1 will be more than the acceptable number. Let us call this minimum value of Q_1 the Aggregate Lower limit 1, or ALC_1 . So:

$$\exists ALC_1 \left[Q_1 \geq ALC_1 \Leftrightarrow \left(MVA_1 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VA_{1j} \quad \text{and} \quad MVR_1 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VR_{1j} \right) \right]$$

Our research has shown that the demand scenarios that include availability and reserve violations when $Q_1 = ALC_1$ are those demand scenarios j with the highest sum of $D_{1j}+D_{2j}+D_{3j}$. For practical purposes, the value of ALC_1 is calculated by an Excel Macro that we wrote, which is shown in section 3.2. This macro starts with a very high Q_1 value and decreases it until it finds the minimum value of Q_1 that will not exceed the maximum allowed violations in either availability or reserve.

2) It is possible to find the minimum value of the sum of Q_1+Q_2 for which the number of availability and reserve violations in week 2 does not exceed the maximum allowed by the user. Let us call this minimum value ALC_2 . So:

$$\exists ALC_2 \left[\sum_{i=1}^2 Q_i \geq ALC_2 \Leftrightarrow \left(MVA_2 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VA_{2j} \quad \text{and} \quad MVR_2 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VR_{2j} \right) \right]$$

For ALC_2 , violations occur in those demand scenarios j with the highest sum of $D_{1j}+D_{2j}+D_{3j}+D_{4j}$. The value of ALC_2 is calculated by the same Macro mentioned above.

3) It is possible to find the minimum value of the sum of $Q_1+Q_2+Q_3$ for which the number of availability and reserve violations in week 3 does not exceed the maximum allowed by the user. Let us call this minimum value ALC_3 . So:

$$\exists ALC_3 \left[\sum_{i=1}^3 Q_i \geq ALC_3 \Leftrightarrow \left(MVA_3 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VA_{3j} \quad \text{and} \quad MVR_3 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VR_{3j} \right) \right]$$

For ALC_3 , violations occur in those demand scenarios j with the highest sum of $D_{1j}+D_{2j}+D_{3j}+D_{4j}+D_{5j}$. The value of ALC_3 is also calculated by the Macro.

4) It is possible to find the minimum value of the sum of $Q_1+Q_2+Q_3+Q_4$ for which the number of availability and reserve violations in week 4 does not exceed the maximum allowed by the user. Let us call this minimum value ALC_4 . So:

$$\exists ALC_4 \left[\sum_{i=1}^4 Q_i \geq ALC_4 \Leftrightarrow \left(MVA_4 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VA_{4j} \quad \text{and} \quad MVR_4 \geq \frac{1}{50} \cdot \sum_{j=1}^{50} VR_{4j} \right) \right]$$

For ALC_4 , violations occur in those demand scenarios j with the highest sum of $D_{1j}+D_{2j}+D_{3j}+D_{4j}+D_{5j}+D_{6j}$. The value of ALC_4 is calculated by the Macro.

Our analysis indicates that, by using these four aggregate lower limits as constraints for the model, we indirectly guarantee that all the 400 availability and reserve constraints are respected by the model. Exhaustive tests corroborated the validity of this claim.

The simplified statement of the constraints is:

$$s.t. \quad 0 \leq Q_i \leq MQ_i \quad \forall i$$

$$Q_1 \geq ALC_1$$

$$\sum_{i=1}^2 Q_i \geq ALC_2$$

$$\sum_{i=1}^3 Q_i \geq ALC_3$$

$$\sum_{i=1}^4 Q_i \geq ALC_4$$

Since i has 4 values, these statements imply 18 individual constraints related to order size, and nothing else. This approach simplifies the problem to the point that Excel Solver solves it optimally in a few seconds.

3.2. Implementing the model in Excel

The implementation in Excel of the mathematical model we developed in section 3.1 is composed of three major parts: a spreadsheet for calculations, a macro to find the limits, and

a set of specific parameters and options for Solver. In this section we discuss all parts in detail.

3.2.1. The spreadsheet

The spreadsheet is the most visible part of the model. Its purpose is to allocate the variables in an orderly fashion, and to perform most calculations. The spreadsheet has four distinct areas: 1) the Input-Output Area, 2) the Price-Related Area, 3) the Demand-Related Area, and 4) the Useful Metrics Area. We recommend the reader to see the model directly in Excel while reading this section. Figure 2 is a screenshot of the spreadsheet and its areas.

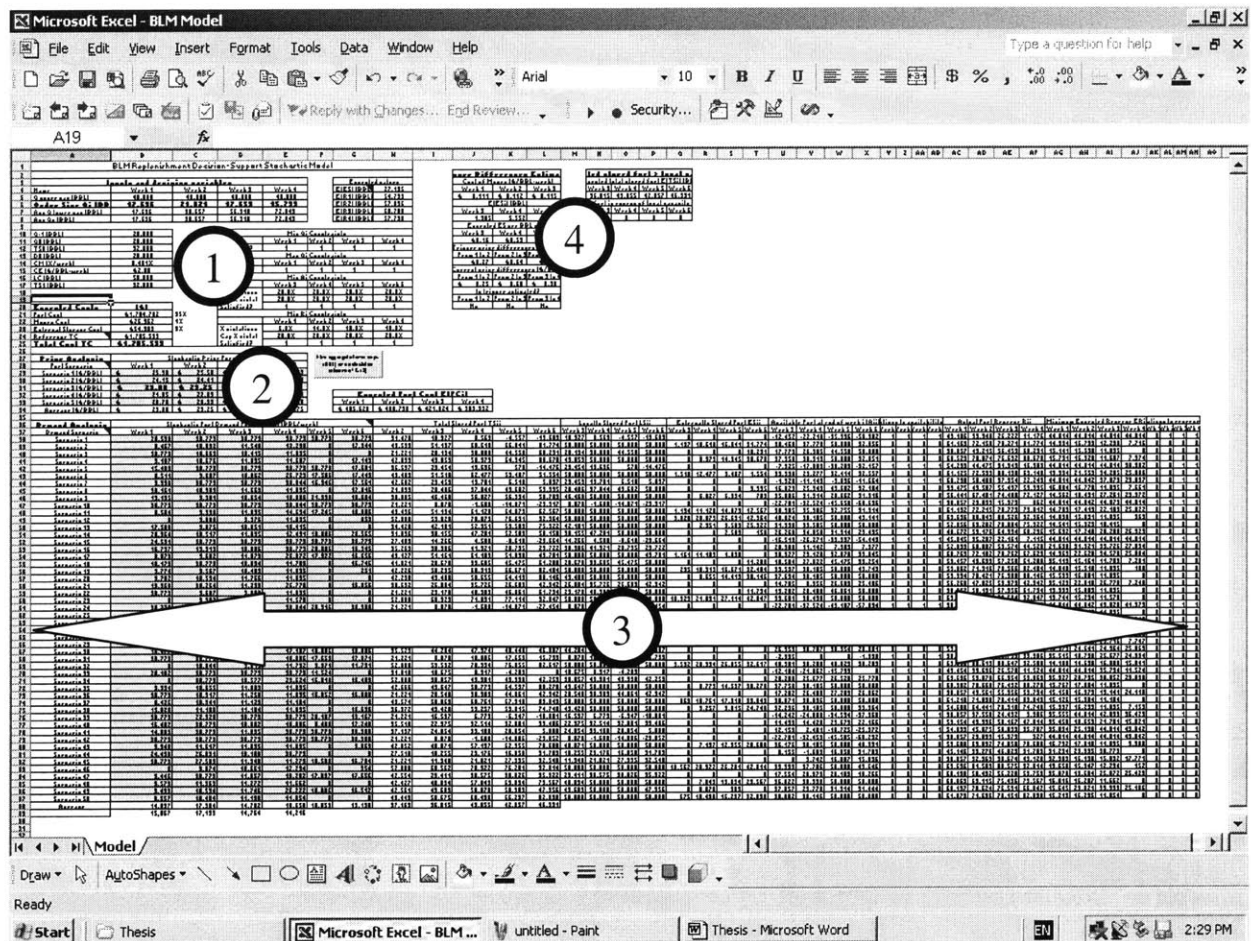


Figure 2 – Screenshot of the whole spreadsheet, with labels over the four different areas

The rest of this section discusses the purpose of the four different areas.

3.2.1.1. The Input-Output Area

The Input-Output Area is labeled with number 1 in Figure 2. As you can see in the close-up screenshot shown in Figure 3, this area includes cells from A1 to H25.

	A	B	C	D	E	F	G	H
1	BLM Replenishment Decision-Support Stochastic Model							
2								
3	Inputs and decision variables					Expected values		
4	Name	Week 1	Week 2	Week 3	Week 4		E[ES] (BBL)	27,195
5	Q upper cap (BBL)	40,000	40,000	40,000	40,000		E[R1] (BBL)	54,799
6	Order Size Qi (BBL)	17,636	21,021	17,653	15,733		E[R2] (BBL)	57,836
7	Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043		E[R3] (BBL)	60,708
8	Agg Qs (BBL)	17,636	38,657	56,310	72,043		E[R4] (BBL)	57,790
9								
10	Q-1 (BBL)	20,000	Min Qi Constraints					
11	Q0 (BBL)	20,000		Week 1	Week 2	Week 3	Week 4	
12	TS0 (BBL)	32,000	Satisfied?	1	1	1	1	
13	D0 (BBL)	20,000	Max Qi Constraints					
14	CM (%/week)	0.481%		Week 1	Week 2	Week 3	Week 4	
15	CE (\$/BBL-week)	\$2.00	Satisfied?	1	1	1	1	
16	LC (BBL)	50,000	Min Ai Constraints					
17	TS1 (BBL)	32,000		Week 3	Week 4	Week 5	Week 6	
18			% violations	20.0%	20.0%	20.0%	20.0%	
19			Cap % violat	20.0%	20.0%	20.0%	20.0%	
20	Expected Costs	(\$)	Satisfied?	1	1	1	1	
21	Fuel Cost	\$1,704,782	95%	Min Ri Constraints				
22	Money Cost	\$26,362	1%	Week 1	Week 2	Week 3	Week 4	
23	External Storage Cost	\$54,389	3%	% violations	6.0%	14.0%	18.0%	18.0%
24	Reference TC	\$1,785,533		Cap % violat	20.0%	20.0%	20.0%	20.0%
25	Total Cost TC	\$1,785,533	Satisfied?	1	1	1	1	

Figure 3 – Screenshot of the Input-Output Area of the model

Input variables - All the input variables that are provided by the user, with the exception of the two stochastic forecasts, are introduced in this area. The user introduces the upper limits for the Q_1 in cell B5, for Q_2 in C5, for Q_3 in D5 and for Q_4 in E5. The user introduces the maximum acceptable probabilities of availability violation in cells E19:H19, and the maximum probabilities of reserve violations in cells E24:H24. The user introduces the sizes of past orders Q_{-1} and Q_0 in cells B10 and B11 respectively, the maximum amount of fuel in storage of week zero TS_0 in cell B12, the demand of week zero D_0 in cell B13, the cost of money (%/week) in cell B14, the cost of external storage (\$/BBL-week) in cell B15, and the local storage capacity in cell B16.

Decision variables – The decision variables, e.g. the variables that Excel Solver will manipulate to minimize the total cost, are in cells B6:E6. The answer to our problem, which is the optimal size of Q_1 , is in cell B6 after the solver has found a solution.

Expected costs – The expected costs are listed in cells A20:B25. The Reference TC (or reference total cost) is the total cost of ordering for every week the minimum order sizes that respect all the aggregate lower limits. As we will see in the next chapter, this reference cost is used to evaluate the savings that are produced by anticipated fuel purchases in the wake of forecasted price increases. The percents shown in cells C21:C23 indicate what percent of the expected TC represent the expected fuel, money and storage costs.

Other expected values – Cells G3:H8 show other expected values. Cell H4 shows the average of the sum of the external storage variables for all weeks of all demand scenarios. This value is used by the model to calculate the expected cost of external storage. Cells H5:H8 show the expected values of fuel reserve for each week.

Status of constraints – Cells D10:H25 are dedicated to the different constraints. In this area, a 1 in a cell means that the constraint is satisfied, and a 0 means that constraint is violated. Cells E12:H12 notify us whether the non-negative order constraints are satisfied each week. Cells E15:H15 do the same for the upper cap constraints. Cells E18:H18 give us the expected probabilities of violation of the availability constraints for each week, while cells E20:H20 verify if these expected probabilities satisfy the user-defined constraints. The same verification is performed for the reserve constraints by cells E23:H23 and E25:H25.

3.2.1.2. The Price-Related Area

The Price-Related Area is labeled with number 2 in Figure 2. As you can see in the close-up screenshot shown in Figure 4, this area includes cells from A27 to J34.

	A	B	C	D	E	F	G	H	I	J
26										
27	Price Analysis	Stochastic Price Forecast i					Find aggregate lower caps (BBL) and calculated reference TC (\$)			
28	Fuel Scenario	Week 1	Week 2	Week 3	Week 4					
29	Scenario 1 (\$/BBL)	\$ 25.30	\$25.58	\$ 26.24	\$27.23					
30	Scenario 2 (\$/BBL)	\$ 24.15	\$24.41	\$ 25.04	\$25.99					
31	Scenario 3 (\$/BBL)	\$ 23.00	\$ 23.25	\$ 23.85	\$ 24.75					
32	Scenario 4 (\$/BBL)	\$ 21.85	\$22.09	\$ 22.66	\$23.51	Expected Fuel Cost E[FCi]				
33	Scenario 5 (\$/BBL)	\$ 20.70	\$20.93	\$ 21.47	\$22.28	Week 1	Week 2	Week 3	Week 4	
34	Average (\$/BBL)	\$ 23.00	\$23.25	\$ 23.85	\$24.75	\$405,628	\$488,738	\$421,024	\$389,392	

Figure 4 – Screenshot of the Price-Related Area of the model

Price forecast – The user inputs the stochastic forecast of fuel prices in cells B29:E33.

Average forecast – The model calculates an average of the fuel prices in cells B34:E34.

Expected fuel costs – Based on the forecasted fuel prices and the sizes of the orders, the model calculates the expected fuel cost for each week, in cells G34:J34.

The model was originally designed to consider individually each price scenario. Analysis confirmed that the results obtained by using the stochastic forecast are identical to the results obtained by using the average forecast of prices.

3.2.1.3. The Demand-Related Area

The Demand-Related Area is labeled with number 3 in Figure 2. As you can see in the close-up partial screenshots shown in Figures 5, 6, 7 and 8, this area includes cells from A36 to AN89 (rows from 48 to 89 do not appear in the figures, for the sake of space).

	A	B	C	D	E	F	G	H	I	J	K	L	
36	Demand Analysis	Stochastic Fuel Demand Forecast Dij (BBL/week)						Total Stored Fuel TSij					
37	Demand Scenario	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 2	Week 3	Week 4	Week 5	Week 6	
38	Scenario 1	20,530	30,779	30,779	30,779	30,779	30,779	31,470	18,327	8,569	-4,557	-19,603	
39	Scenario 2	8,467	10,032	11,540	12,230	0	17,944	43,533	51,137	60,618	66,041	81,774	
40	Scenario 3	30,779	10,663	10,415	11,895	0	0	21,221	28,194	38,800	44,558	60,291	
41	Scenario 4	9,107	10,676	11,895	11,687	0	17,149	42,893	49,853	58,979	64,945	80,678	
42	Scenario 5	15,403	30,779	30,779	30,779	30,779	17,681	36,597	23,454	13,696	570	-14,476	
43	Scenario 6	8,511	9,607	10,062	26,723	17,586	17,443	43,489	51,518	62,477	53,407	51,554	
44	Scenario 7	9,398	30,779	30,779	30,844	16,346	17,561	42,602	29,459	19,701	6,510	5,897	
45	Scenario 8	30,161	10,989	11,663	11,895	0	17,816	21,839	28,486	37,844	43,602	59,335	
46	Scenario 9	13,195	9,981	10,654	18,086	21,338	18,684	38,805	46,460	56,827	56,394	50,789	
47	Scenario 10	30,779	30,779	30,779	30,844	30,779	30,779	21,221	8,078	-1,680	-14,871	-29,917	

Figure 5 – Screenshot of the extreme-left of the Demand-Related Area of the model

About the labels of the scenarios – The label that identifies the 50 stochastic scenarios of the demand forecast is in cells A38:A87. These labels apply across all the Demand-Related Area, because all the variables in this area, not only the demand forecast, are scenario-specific. The model will perform individual storage, availability and reserve calculations for each scenario. Therefore, the label of column A refers to all other calculations in a same row.

Demand forecast – The user inputs the stochastic demand forecast for weeks 1-6 in cells B38:G87. It is important that this forecast be introduced in barrels, with the adjustment of 5% and unsorted. Sorting the values of each column individually would destroy the sequential nature of the forecast, so it should be avoided.

Total stored fuel – The model calculates the amount of fuel stored both externally and locally, for each demand scenario of weeks 2-6, in cells H38:L87. These values are in barrels.

	A	M	N	O	P	Q	R	S	T
36	Demand Analysis	Locally Stored Fuel LSij				Externally Stored Fuel ESij			
37	Demand Scenario	Week 3	Week 4	Week 5	Week 6	Week 3	Week 4	Week 5	Week 6
38	Scenario 1	18,327	8,569	-4,557	-19,603	0	0	0	0
39	Scenario 2	50,000	50,000	50,000	50,000	1,137	10,618	16,041	31,774
40	Scenario 3	28,194	38,800	44,558	50,000	0	0	0	10,291
41	Scenario 4	49,853	50,000	50,000	50,000	0	8,979	14,945	30,678
42	Scenario 5	23,454	13,696	570	-14,476	0	0	0	0
43	Scenario 6	50,000	50,000	50,000	50,000	1,518	12,477	3,407	1,554
44	Scenario 7	29,459	19,701	6,510	5,897	0	0	0	0
45	Scenario 8	28,486	37,844	43,602	50,000	0	0	0	9,335
46	Scenario 9	46,460	50,000	50,000	50,000	0	6,827	6,394	789
47	Scenario 10	8,078	-1,680	-14,871	-29,917	0	0	0	0

Figure 6 – Screenshot of the center-left of the Demand-Related Area of the model

Fuel stored locally – The model calculates the amount of fuel stored locally, for each demand scenario of weeks 3-6, in cells M38:P87. These values are in barrels.

Fuel stored externally – The model calculates the amount of fuel stored externally, for each demand scenario of weeks 3-6, in cells Q38:T87. These values are in barrels.

	A	U	V	W	X	Y	Z	AA	AB
36	Demand Analysis	Available Fuel at end of week i (Aij)				Violations to availability VAij			
37	Demand Scenario	Week 3	Week 4	Week 5	Week 6	Week 3	Week 4	Week 5	Week 6
38	Scenario 1	-12,452	-22,210	-35,336	-50,382	1	1	1	1
39	Scenario 2	38,460	37,770	50,000	32,056	0	0	0	0
40	Scenario 3	17,779	26,905	44,558	50,000	0	0	0	0
41	Scenario 4	37,958	38,313	50,000	32,851	0	0	0	0
42	Scenario 5	-7,325	-17,083	-30,209	-32,157	1	1	1	1
43	Scenario 6	39,938	23,277	32,414	32,557	0	0	0	0
44	Scenario 7	-1,320	-11,143	-9,836	-11,664	1	1	1	1
45	Scenario 8	16,823	25,949	43,602	32,184	0	0	0	0
46	Scenario 9	35,806	31,914	28,662	31,316	0	0	0	0
47	Scenario 10	-22,701	-32,524	-45,650	-60,696	1	1	1	1

Figure 7 – Screenshot of the center-right of the Demand-Related Area of the model

Minimum available fuel – The model calculates the minimum level of available fuel, for each demand scenario of weeks 3-6, in cells U38:X87. These values are in barrels.

Availability violations – The model evaluates whether there are availability violations, for each demand scenario of weeks 3-6, in cells Y38:AB87. A violation is indicated by 1.

	A	AC	AD	AE	AF	AG	AH	AI	AJ	AK	AL	AM	AN
36	Demand Analysis	Actual Fuel Reserve Rij				Minimum Expected Reserve ERij				Violations to reserve VRij			
37	Demand Scenario	Week 1	Week 2	Week 3	Week 4	Week 1	Week 2	Week 3	Week 4	Wk 1	Wk 2	Wk 3	Wk 4
38	Scenario 1	49,106	39,348	26,222	11,176	44,014	44,014	44,014	44,014	0	1	1	1
39	Scenario 2	61,169	72,158	78,271	81,774	14,994	16,799	12,230	7,716	0	0	0	0
40	Scenario 3	38,857	49,215	56,453	60,291	15,141	15,530	11,895	0	0	0	0	0
41	Scenario 4	60,529	70,874	76,632	80,678	15,791	16,920	11,687	7,374	0	0	0	0
42	Scenario 5	54,233	44,475	31,349	16,303	44,014	44,014	44,014	38,382	0	0	1	1
43	Scenario 6	61,125	72,539	80,130	69,140	13,934	21,553	34,285	25,086	0	0	0	0
44	Scenario 7	60,238	50,480	37,354	22,243	44,014	44,042	37,873	23,897	0	0	1	1
45	Scenario 8	39,475	49,507	55,497	59,335	16,004	16,778	11,895	7,661	0	0	0	0
46	Scenario 9	56,441	67,481	74,480	72,127	14,562	18,431	27,261	29,372	0	0	0	0
47	Scenario 10	38,857	29,099	15,973	862	44,014	44,042	44,079	44,014	1	1	1	1

Figure 8 – Screenshot of the extreme-right of the Demand-Related Area of the model

Minimum actual reserve – The model calculates the minimum level of fuel reserve, for each demand scenario of weeks 1-4, in cells AC38:AF87. These values are in barrels.

Expected reserve – The model calculates the expected level of fuel reserve, for each demand scenario of weeks 1-4, in cells AG38:AJ87. These values are in barrels.

Reserve violations – The model evaluates whether there are reserve violations, for each demand scenario of weeks 1-4, in cells AK38:AN87. A violation is indicated by 1.

3.2.1.4. The Useful Metrics Area

The Useful Metrics Area is labeled with number 4 in Figure 2. As you can see in the close-up screenshot shown in Figure 9, this area includes cells from J2 to Q20. The purpose of this area is to present some useful variables that did not fit in any of the previous areas.

Trigger Differences – Trigger difference is the holding cost per barrel between two weeks. If the difference in fuel prices between these two weeks is greater than the trigger difference, purchasing fuel in anticipation of the price rise could be the best option.

Expected stored fuel – To know the expected stored fuel for each week is useful, because if it exceeds the local storage capacity, we can expect the need of external storage.

	J	K	L	M	N	O	P	Q
2	Trigger Differences Estimation				Is expected stored fuel > local capacity?			
3	Cost of Money (\$/BBL-week)				Expected total stored fuel E[TSi] (BBL)			
4	Week 1	Week 2	Week 3		Week 3	Week 4	Week 5	Week 6
5	\$ 0.111	\$ 0.112	\$ 0.115		36,815	43,055	42,057	46,931
6	E[ESi] (BBL)				Stored fuel in excess of local capacity (BBL)			
7	Week 3	Week 4	Week 5		Week 3	Week 4	Week 5	Week 6
8	1,385	5,552	7,276		0	0	0	0
9	Expected ES per BBL of Q							
10	Week 3	Week 4	Week 5					
11	\$0.16	\$0.53	\$0.82					
12	Trigger price differences (\$/BBL)							
13	From 1 to 2	From 2 to 3	From 3 to 4					
14	\$0.27	\$0.64	\$0.94					
15	Current price differences (\$/BBL)							
16	From 1 to 2	From 2 to 3	From 3 to 4					
17	\$ 0.25	\$ 0.60	\$ 0.90					
18	Is trigger activated?							
19	From 1 to 2	From 2 to 3	From 3 to 4					
20	No	No	No					

Figure 9 – Screenshot of the Useful Metrics Area of the model

3.2.2. The Macro

The following is the code of the macro that finds the values of aggregate lower limits (ALC). The reader should use the Input-Output Area screenshot, shown in Figure 3, as reference to understand the names of cells mentioned in this Macro. The macro begins:

```
Sub FindCaps()
' FindCaps Macro
' Macro recorded 3/20/2004 by Roberto Perez-Franco
```

The macro clears the cells where the cap values will be introduced later on:

```
Range("B7").Value = ""
Range("C7").Value = ""
Range("D7").Value = ""
Range("E7").Value = ""
```

The macro sets Q1, Q2, Q3 and Q4 to their maximum acceptable values:

```
Range("B6").Value = Range("B5").Value
Range("C6").Value = Range("C5").Value
Range("D6").Value = Range("D5").Value
Range("E6").Value = Range("E5").Value
```

The macro finds ALC1, by decreasing Q1 until it reaches the maximum violations:

```
MaxQ = Range("B5").Value
StepQ = MaxQ / -100
For TryQ = MaxQ To 0 Step StepQ
Range("B6").Value = TryQ
If Range("E20").Value * Range("E25").Value = 1 Then
LatestGoodQ = TryQ
End If
Next TryQ
```

The macro refines the value of ALC1 by repeating the process with smaller steps:

```
Quanto = -1 * StepQ
MaxQ = LatestGoodQ + Quanto
If (StepQ / -1000) > 1 Then
StepQ = LatestGoodQ / -1000
Else
StepQ = -1
End If
For TryQ = MaxQ To (MaxQ - 2 * Quanto) Step StepQ
Range("B6").Value = TryQ
If Range("E20").Value * Range("E25").Value = 1 Then
LatestGoodQ = TryQ
End If
Next TryQ
```

Having found and refined the value of ALC1, the macro copies it in its place:

```
Range("B6").Value = LatestGoodQ
Range("B7").Value = Range("B6").Value
```

That is all for ALC1. The macro repeats the same procedure for ALC2, ALC3 and ALC4.

The following code finds ALC2.

```
MaxQ = Range("C5").Value
StepQ = MaxQ / -100
```

```

For TryQ = MaxQ To 0 Step StepQ
Range("C6").Value = TryQ
If Range("F20").Value * Range("F25").Value = 1 Then
LatestGoodQ = TryQ
End If
Next TryQ
Quanto = -1 * StepQ
MaxQ = LatestGoodQ + Quanto
If (StepQ / -1000) > 1 Then
StepQ = LatestGoodQ / -1000
Else
StepQ = -1
End If
For TryQ = MaxQ To (MaxQ - 2 * Quanto) Step StepQ
Range("C6").Value = TryQ
If Range("F20").Value * Range("F25").Value = 1 Then
LatestGoodQ = TryQ
End If
Next TryQ
Range("C6").Value = LatestGoodQ
Range("C7").Value = Range("B6").Value + Range("C6").Value

```

The following code finds ALC3.

```

MaxQ = Range("D5").Value
StepQ = MaxQ / -100
For TryQ = MaxQ To 0 Step StepQ
Range("D6").Value = TryQ
If Range("G20").Value * Range("G25").Value = 1 Then
LatestGoodQ = TryQ
End If
Next TryQ
Quanto = -1 * StepQ
MaxQ = LatestGoodQ + Quanto
If (StepQ / -1000) > 1 Then
StepQ = LatestGoodQ / -1000
Else
StepQ = -1
End If
For TryQ = MaxQ To (MaxQ - 2 * Quanto) Step StepQ
Range("D6").Value = TryQ
If Range("G20").Value * Range("G25").Value = 1 Then
LatestGoodQ = TryQ
End If
Next TryQ
Range("D6").Value = LatestGoodQ
Range("D7").Value = Range("B6").Value + Range("C6").Value + Range("D6").Value

```

The following code finds ALC4.

```

MaxQ = Range("E5").Value
StepQ = MaxQ / -100
For TryQ = MaxQ To 0 Step StepQ
Range("E6").Value = TryQ
If Range("H20").Value * Range("H25").Value = 1 Then

```

```

LatestGoodQ = TryQ
End If
Next TryQ
Quanto = -1 * StepQ
MaxQ = LatestGoodQ + Quanto
If (StepQ / -1000) > 1 Then
StepQ = LatestGoodQ / -1000
Else
StepQ = -1
End If
For TryQ = MaxQ To (MaxQ - 2 * Quanto) Step StepQ
Range("E6").Value = TryQ
If Range("H20").Value * Range("H25").Value = 1 Then
LatestGoodQ = TryQ
End If
Next TryQ
Range("E6").Value = LatestGoodQ
Range("E7").Value = Range("B6").Value + Range("C6").Value + Range("D6").Value
+ Range("E6").Value

```

Once all the limits have been found, the macro preserves the reference total cost TC:

```
Range("b24").Value = Range("B25").Value
```

To finish, the macro sets all Q to their maximum values, so that the model is ready to be solved by the user immediately after the macro is finished:

```

Range("B6").Value = Range("B5").Value
Range("c6").Value = Range("c5").Value
Range("d6").Value = Range("d5").Value
Range("e6").Value = Range("e5").Value
End Sub

```

This macro has proven to be effective and fast. In all the tests we ran, it quickly found after a few seconds the four aggregate lower limits with high accuracy.

3.2.3. Solver parameters and options

In this section, we describe how we set the parameters and options of the Excel Solver to run the model. The reader should use Figure 3 as a reference to understand what variables belong to the cells mentioned in this section.

3.2.3.1. Solver Parameters

Based on the positions of the relevant variables in the spreadsheet we prepared in section 3.2.1, the Solver Parameters should be set to the following values:

- Set Target Cell: $\$B\25 (the cell with the total cost)
- Equal to: Min (minimize)
- By Changing Cells: $\$B\$6:\$E\6 (the cells with Q_1 , Q_2 , Q_3 and Q_4)
- Subject to the Constraints:
 - $\$B\$6 \leq \$B\5 (order Q_1 should be smaller than its upper cap)
 - $\$B\$6 \geq \$B\7 (order Q_1 cannot be smaller than ALC_1)
 - $\$C\$6 \leq \$C\5 (order Q_2 should be smaller than its upper cap)
 - $\$C\$6 \geq 0$ (order Q_2 cannot be negative)
 - $\$C\$8 \geq \$C\7 (the sum of Q_1+Q_2 cannot be smaller than ALC_2)
 - $\$D\$6 \leq \$D\5 (order Q_3 should be smaller than its upper cap)
 - $\$D\$6 \geq 0$ (order Q_3 cannot be negative)
 - $\$D\$8 \geq \$D\7 (the sum of $Q_1+Q_2+Q_3$ cannot be smaller than ALC_3)
 - $\$E\$6 \leq \$E\5 (order Q_4 should be smaller than its upper cap)
 - $\$E\$6 \geq 0$ (order Q_4 cannot be negative)
 - $\$E\$8 \geq \$E\7 (the sum of $Q_1+Q_2+Q_3+Q_4$ cannot be smaller than ALC_4)

Notice that there is no need to specify $\$B\$6 \geq 0$ (order Q_1 cannot be negative) as a constraint, because there is already another constraint that specifies that Q_1 will be larger than ALC_1 , and the value of ALC_1 is never negative (the macro takes care of that).

Figures 9 and 10 show screenshots of Excel's "Solver Parameters" dialog box. The reason we show two screenshots is that there are eleven constraints but only six fit in the dialog box window at a time. Figure 10 shows the first constraints, and Figure 11 shows the rest.

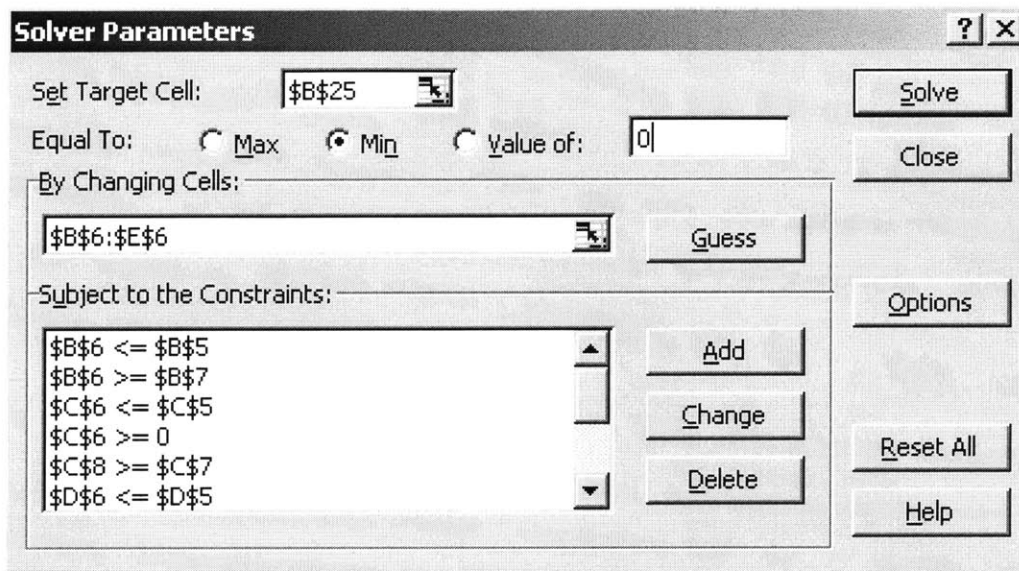


Figure 10 – Excel Solver Parameters for the model – First six constraints are shown

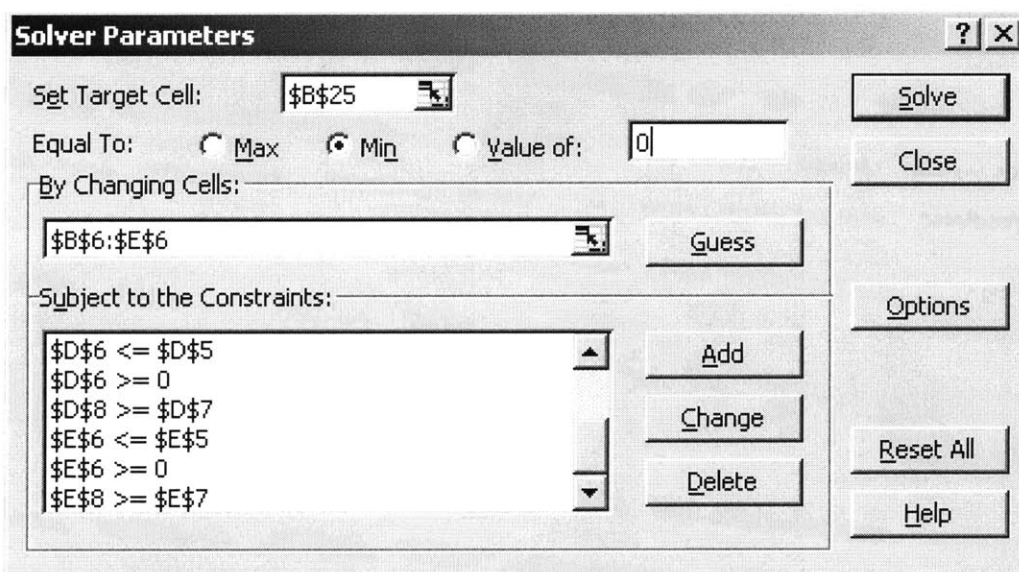


Figure 11 – Excel Solver Parameters for the model – Remaining constraints are shown

3.2.3.2. Solver Options

Although other settings might work, the model was tested using one specific group of settings for “Solver Options” and always returned satisfactory results. These settings are:

- Max Time: 1000 seconds
- Iterations: 20000
- Precision: 0.000000001
- Tolerance: 0.001
- Convergence: 0.00000001
- Assume Linear Model: No
- Assume Non-Negative: Yes
- Use Automatic Scaling: Yes
- Show Iteration Results: No
- Estimates: Tangent
- Derivatives: Central
- Search: Conjugate

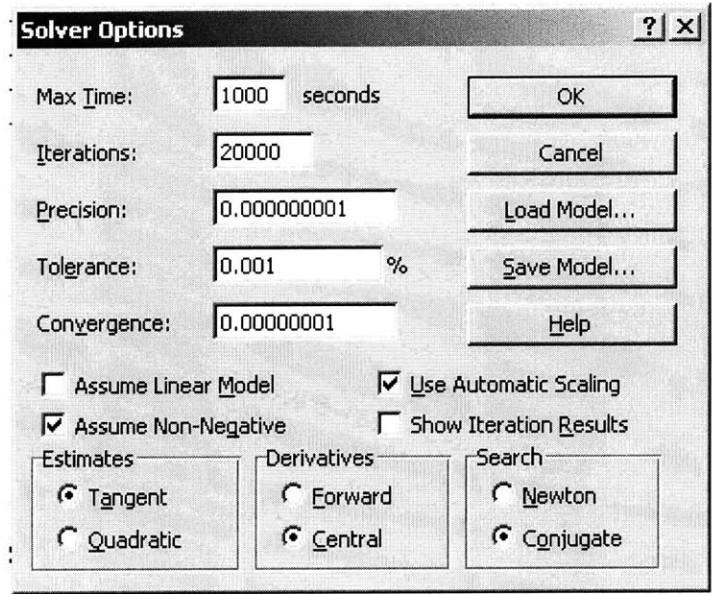


Figure 12 – Excel Solver Options settings for the model

After setting the parameters and options of the Solver to these values, the model is ready to be used. The next chapter explains how to use it and discusses some examples.

4. Analysis of Results

4.1. Using the model

Using the model is simple. Let's say that today is Friday April 2, 2004, and the user wants to calculate the size of the order that will be placed next week. Based on today's date, the weeks' notation is as follows:

- Week -1 goes from Saturday March 20 to Friday March 26.
- Week 0 goes from Saturday March 27 to Friday April 2nd (e.g. today).
- Week 1 goes from Saturday April 3rd (e.g. tomorrow) to Friday April 9.
- Week 2 goes from Saturday April 10 to Friday April 16.
- Week 3 goes from Saturday April 17 to Friday April 23.
- Week 4 goes from Saturday April 24 to Friday April 30.
- Week 5 goes from Saturday May 1st to Friday May 7.
- Week 6 goes from Saturday May 8 to Friday May 14.

To determine the best size for next week's order (Q_1), just follow these steps:

1) Open the model in Excel. The file is called "BLM Model.xls". If asked by Excel when it is opening the file, the user should accept the use of Macros.

2) The user should input the following information:

- The upper limits for Q_1 in cell B5, for Q_2 in C5, for Q_3 in D5 and for Q_4 in E5.
- The maximum acceptable probabilities of availability violation in cells from E19 to H19, and the maximum probabilities of reserve violations in cells E24:H24.
- The sizes of past orders Q_{-1} and Q_0 in cells B10 and B11 respectively.
- The maximum amount of fuel that was in storage in week zero TS_0 , in cell B12.
- The demand of week zero D_0 in cell B13. Notice that today's demand is not known yet, but can easily be estimated by BLM's staff.

- The cost of money per week (%/week), in cell B14. This cost of money is calculated as the cost of money per year (%), divided by 52 weeks.
- The cost of external storage (\$/BBL-week) in cell B15.
- The local storage capacity in cell B16.
- The stochastic forecast of fuel prices. Our analysis shows that using a stochastic price forecast yields the same results as using a deterministic price forecast. Therefore, the user is free to use either one. If the user wants to input a stochastic forecast, it should be introduced in cells B29:E33. If the user prefers to use a deterministic value, it should be introduced in cells B34:E34.
- The stochastic demand forecast in cells B38:G87. It is important that this forecast be introduced in barrels, with the adjustment of 5% and unsorted. The stochastic demand forecast we use for our examples is based on real data from a January run of SDDP, provided to us by BLM.

3) With all input data in place, click the macro button. This is the only button in our model. It has the following label: “Find aggregate lower caps (BBL) and calculated reference TC (\$).” When we click this button, the model will run the macro, which will find the aggregate lower limits and the reference total cost.

4) Run the solver. This is done by clicking on Tools / Solver, in Excel’s menu. (If you do not find Solver in your Excel menu, do not panic. As of Excel 2002, the Solver is not installed by default: it is an Add-in that should be added once. You might need to install it with the Microsoft Office installation CD.) The Solver Parameters dialog box will appear. Click on the “Solve” button. A couple of seconds later, you will see a dialog box called “Solver Results”. Click the “OK” button to keep the solution found by Solver.

After these four steps, the answer (e.g. the recommended order size Q_1) is in cell B6.

Let us analyze nine examples which illustrate the performance of the model.

4.2. Example #1

Consider that fuel prices remain stable during the four weeks.

Example #1 - Stable fuel prices				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$25.00	\$25.00	\$25.00	\$25.00
Order Size (BBL)	17,636	21,021	17,653	15,733
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	17,636	38,657	56,310	72,043
% violations of Ai	20%	20%	20%	20%
Expected Costs	(\$)		E[TSI] - LC (BBL)	
Fuel Cost	\$1,801,075		Week 3	0
Money Cost	\$27,780		Week 4	0
External Storage Cost	\$54,389		Week 5	0
Reference TC	\$1,883,245		Week 6	0
Total Cost TC	\$1,883,245			
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.28	\$ 0.65	\$ 0.94	

Figure 13 – Summary of the most important variables of Example #1

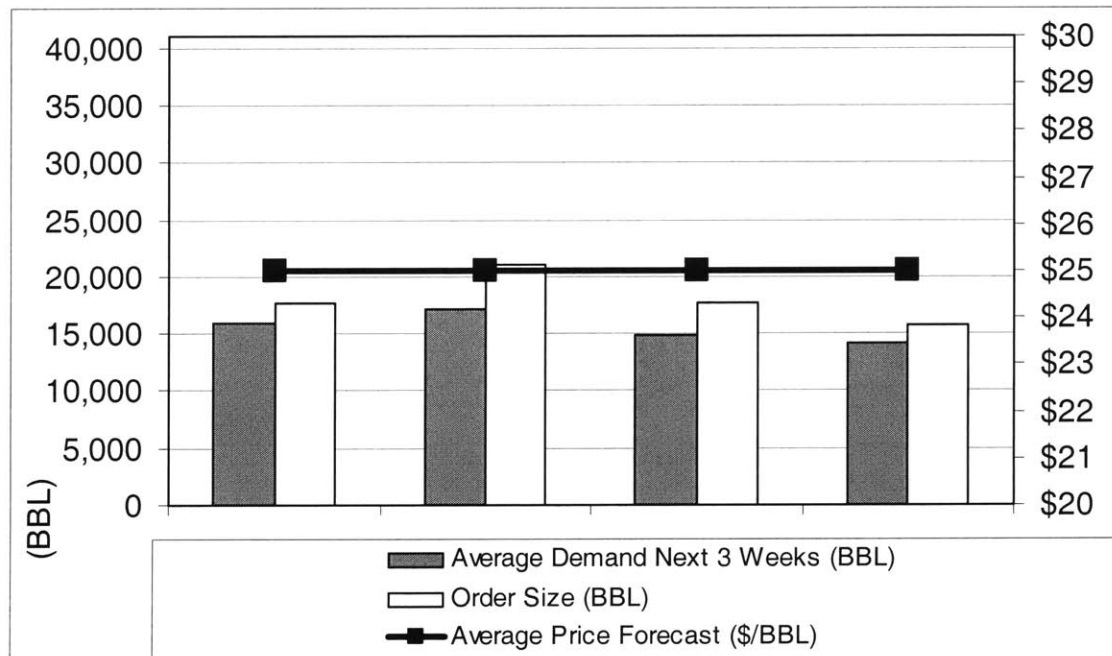


Figure 14 – Graphical representation of demand, price and orders of Example #1

Analysis of Example #1

With constant prices, we would expect the model to buy each week the minimum fuel that satisfies the constraints, because of money and storage costs. In other words, since there are no price increases in the future – as far as the model knows - there is no need to build up inventory that would accrue holding and money costs. That is exactly the result we get. In cases like this, the model recommends for every week the minimum amounts that respect the aggregate lower limits.

Notice that availability violations have a probability of 20% in all weeks, which is exactly the value that was specified as maximum acceptable by the user. The reason we do not show the probability of reserve violations is that our analysis showed they are not constraining. For this example, the reserve violation probabilities for the four weeks are 6.0%, 14.0%, 18.0% and 18.0%, respectively.

The reason the order sizes in the graph of Figure 14 are larger than the average demand for every week, is that order sizes are chosen to satisfy demand in 80% of the scenarios, while the demand shown in the graph is the average value of the next three weeks for all scenarios (not only for the top 20%).

Please notice that the Reference Total Cost is equal to the Total Cost. The Reference total cost is the relevant cost associated to purchasing the minimum amount of fuel each week that will not violate any of the aggregate lower limits. As the reader can verify, in this example the aggregate order sizes are identical to the aggregate lower limits. This will be the case in every example where there is no fuel purchase in advance of a price increase.

4.3. Example #2

Consider that the fuel prices continually decrease from each week to the next.

Example #2 - Decreasing fuel prices				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$26.00	\$25.00	\$24.00	\$23.00
Order Size (BBL)	17,636	21,021	17,653	15,733
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	17,636	38,657	56,310	72,043
% violations of Ai	20%	20%	20%	20%
Expected Costs	(\$)		E[TSi] - LC (BBL)	
Fuel Cost	\$1,769,592		Week 3	0
Money Cost	\$27,196		Week 4	0
External Storage Cost	\$54,389		Week 5	0
Reference TC	\$1,851,178		Week 6	0
Total Cost TC	\$1,851,178			
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.28	\$ 0.65	\$ 0.94	

Figure 15 – Summary of the most important variables of Example #2




Figure 16 – Graphical representation of demand, price and orders of Example #2

Analysis of Example #2

With prices that decrease constantly, the decision we would expect from the model is similar to that of Example #1: to buy for every week the minimum amounts that respect the aggregate lower limits. That is, in fact, the result we obtain.

As opposed to the cases with significant price increases, where the model could save money by buying in advance of the price rise, when prices are going down there is nothing the model can do to save money beyond buying the minimum that will respect the constraints. That is why it is important to correctly assess the risk that the company is willing to take with regard to stockout probability: the higher the risk, the lower the order sizes. We recommend a probability of stockout between 10% and 33%. We used 20% in our examples.

In the results of this example, availability violations have a probability of 20% in all weeks, exactly the value specified as maximum by the user. Again, the Reference Total Cost is equal to the Total Cost, signaling that there is no speculative fuel purchase.

4.4. Example #3

Consider that fuel prices decrease only once, remaining constant the other two weeks.

Example #3 - Decrease from one level to another level				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$25.00	\$25.00	\$23.00	\$23.00
Order Size (BBL)	17,636	21,021	17,653	15,733
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	17,636	38,657	56,310	72,043
% violations of Ai	20%	20%	20%	20%
Expected Costs	(\$)		E[TSI] - LC (BBL)	
Fuel Cost	\$1,734,303		Week 3	0
Money Cost	\$26,641		Week 4	0
External Storage Cost	\$54,389		Week 5	0
Reference TC	\$1,815,333		Week 6	0
Total Cost TC	\$1,815,333			
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.28	\$ 0.65	\$ 0.93	

Figure 17 – Summary of the most important variables of Example #3




Figure 18 – Graphical representation of demand, price and orders of Example #3

Analysis of Example #3

The rationale for this example is that, with its current supplier, BLM experiences no more than one change in price every four weeks. As we explained in section 2.3.1, the current supplier charges BLM the same price for any order placed in one given month. For example, all orders placed in April will have the same price, all orders placed in May will have the same price (although May's price could be different from April's price), and so on.

In this example, where prices decrease once in four weeks and remain constant the rest of the time, the decision the model makes is the same it made for Example #2: to buy for every week the minimum amounts that respect the aggregate lower limits.

4.5. Example #4

Consider that fuel prices increase slightly from every week to the next.

Example #4 - Fuel prices increasing less than trigger difference				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$23.00	\$23.25	\$23.85	\$24.75
Order Size (BBL)	17,636	21,021	17,653	15,733
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	17,636	38,657	56,310	72,043
% violations of Ai	20%	20%	20%	20%
Expected Costs	(\$)			
Fuel Cost	\$1,704,782	E[TSi] - LC (BBL)		
Money Cost	\$26,362	Week 3	0	
External Storage Cost	\$54,389	Week 4	0	
Reference TC	\$1,785,533	Week 5	0	
Total Cost TC	\$1,785,533	Week 6	0	
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.27	\$ 0.64	\$ 0.94	

Figure 19 – Summary of the most important variables of Example #4




Figure 20 – Graphical representation of demand, price and orders of Example #4

Analysis of Example #4

In this example, prices increase slightly from every week to the next. We would expect the model to buy fuel in advance of the price rises if, and only if, the savings of anticipated purchases exceed the cost of holding the excess fuel for the future.

Holding costs, in our model, are composed of the costs of storage and money. The holding costs per barrel from one week to the next are what we call the “trigger price differences”, on the basis that if the difference in prices from one week to the next is larger than this holding cost (the trigger difference), considering the purchase of fuel in anticipation of the price rise is worth considering.

For example, the difference in fuel prices from week 1 to week 2 is $\$23.25/\text{BBL} - \$23.00/\text{BBL} = \$0.25/\text{BBL}$. The trigger price difference for the same period is $\$0.268/\text{BBL}$, calculated as the sum of money cost ($\$0.111/\text{BBL}$) and the storage cost spread across all the stored barrels ($\$0.157/\text{BBL}$.) Since the price difference is smaller than the trigger (e.g. $\$0.25/\text{BBL} < \$0.268/\text{BBL}$), the user knows there is no incentive to build additional fuel inventory in week 1 for use in week 2. By the same token, the user knows there is no incentive for extra inventory in the other weeks: a) the price difference between weeks 2 and 3 ($\$0.60/\text{BBL}$) is smaller than the respective trigger ($\$0.640/\text{BBL}$), and b) the price difference between weeks 3 and 4 ($\$0.90/\text{BBL}$) is smaller than its trigger ($\$0.939/\text{BBL}$.)

The calculation of the trigger difference prices is just a reference for the user. The model itself does not use the triggers for anything: Solver makes the decision based on cost.

For this example, the model decides to buy the minimum amounts every week, just as it did for the three previous examples. The result indicates that the model does not find it attractive to purchase fuel in anticipation of the price rise. The trigger price differences are consistent with this conclusion: the actual price differences are smaller than the triggers, and therefore there is no reason to consider anticipated fuel purchases. Please notice that trigger price differences are not fixed, as the next example shows.

4.6. Example #5

Consider that fuel prices increase a little more than the trigger differences.

Example #5 - Fuel prices increasing a bit more than trigger difference				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$23.00	\$24.00	\$24.00	\$25.35
Order Size (BBL)	24,216	14,441	24,051	9,335
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	24,216	38,657	62,708	72,043
% violations of Ai	16%	20%	20%	22%
Expected Costs	(\$)			
Fuel Cost	\$1,717,418	E[TSI] - LC		(BBL)
Money Cost	\$28,247	Week 3	0	
External Storage Cost	\$66,103	Week 4	0	
Reference TC	\$1,813,806	Week 5	0	
Total Cost TC	\$1,811,768	Week 6	0	
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.44	\$ 0.88	\$ 0.99	

Figure 21 – Summary of the most important variables of Example #5




Figure 22a – Graphical representation of demand, price and orders of Example #5

Analysis of Example #5

In this example, the prices in the second and fourth week are significantly higher than the prices in the first and third weeks, respectively. The model responds to this situation, by building additional fuel inventory in the first and third weeks, anticipating the price rises.

It is interesting to notice that the total amount of the four orders together is the same as the previous four examples: 72,043 BBL. However, it is the way the barrels are distributed in the individual orders that changes. Now, Q_1 is larger than in the previous examples, while Q_2 is smaller than in the previous examples. The response of the model is what we would have expected: it is building an inventory of cheap fuel to avoid buying when it is expensive.

To quantify the expected savings of the anticipated purchases, we can use the reference total cost. The reference total cost, as shown in Figure 21, is \$1,813,806, which represents the cost of buying in each week just the minimum fuel needed to satisfy the constraints. The actual total cost, as shown in the same figure, is \$1,811,768. This indicates that the expected savings of purchasing some fuel in anticipation of the price rise is around \$2,038.

But, why did not the model just buy all the fuel for the first two weeks in Q_1 and make Q_2 equal to zero? Would not that provide us with even larger savings? The answer is no. If we buy nothing in week 2 and instead purchase all the fuel for the first two weeks in week one, the total cost is \$2,204 higher: \$1,813,972. The reason the cost is higher is that, although we save $\$1,717,418 - \$1,702,977 = \$14,441$ in fuel cost, we incur in an additional storage cost of $\$81,152 - \$66,103 = \$15,049$, and an additional money cost of $\$29,844 - \$28,247 = \$1,597$. The reason we pay **less fuel cost** is that we are buying all fuel in the first order, avoiding the purchase of more expensive fuel in the second week. The reason we pay **more storage cost** is that, since we placed a much larger order the first week, we will have to lease more external storage to store this extra fuel, which will accrue extra storage cost. The reason we pay **more money cost** is that money cost is charged on the fuel that remains

unused from one week to the next: the excess fuel purchased the first week remains as reserve, and therefore represents an extra money cost charge for the company.

The sizes of Q_1 and Q_2 that the model selected are the best, as can be seen in Figure 22b, which plots the total cost for the whole horizon versus the size of Q_2 , from 0 BBL to 21,000 BBL. This plot assumes that the sum of Q_1+Q_2 remains constant at 38,657 BBL. It confirms the decision of the model: the lowest total cost is obtained when Q_2 equals 14,441 BBL and therefore Q_1 equals $38,657 \text{ BBL} - 14,441 \text{ BBL} = 24,216 \text{ BBL}$




Figure 22b – Total cost (\$) versus the order size of Q_2 (BBL), keeping Q_1+Q_2 constant

The reason total storage costs vary, even if we keep the sum of the orders constant, is that locally fuel stored does not pay storage cost. Since only the stored fuel in excess of the local storage capacity is stored externally, the amount of fuel stored externally depends directly on the order sizes. Larger order sizes will create the need to store larger amounts of

fuel, increasing the amount in excess of local storage capacity that must be stored externally, therefore increasing the storage cost.

An admittedly oversimplified, deterministic example, considering only two order sizes, can help us understand why the overall cost of storage can vary by the size of one order, even if the overall total of the orders remains the same. In this example, let us assume that we only have two order sizes to decide: Q_1 and Q_2 . Pay attention to how changes in the distribution of the total amount ordered between these two orders changes the total storage cost for the whole horizon.

Case 1 – Let us assume that both orders were of the same size and that everything is identical for both weeks (demand, inventory, etc.). Let us consider the total stored fuel as:

- In week 1: 52,000 barrels-week
- In week 2: 52,000 barrels-week
- In both weeks: 104,000 barrels-week

In this first case we pay external storage for:

- 2,000 barrels-week in week 1
- 2,000 barrels-week in week 2
- 4,000 barrels-week in both weeks

Case 2 - Now consider a second case, where the first order was larger, while the second order was smaller by the same amount. Let us consider the total stored fuel as:

- In week 1: 72,000 barrels-week
- In week 2: 32,000 barrels-week
- In both weeks: 102,000 barrels-week

Notice the overall total is the same for both cases. The only difference is that in this second case, both orders are not the same. In this second case we pay external storage for:

- 22,000 barrels-week in week 1
- Zero barrels-week in week 2

- 22,000 barrels-week in both weeks

Let us assume that the cost of storing a barrel externally is 1\$/barrel-week. We pay:

- In the first case, \$4,000 of storage
- In the second case, \$22,000 of storage

If we divide this storage cost across the total number of barrels that were stored in the whole horizon, we find that the cost of our inventory per barrel stored everywhere is:

- Case 1: $\$4,000/104,000$ barrels = $\$0.0385/\text{barrel-week}$
- Case 2: $\$22,000/104,000$ barrels = $\$0.2115/\text{barrel-week}$

The only difference between the two cases was the distribution of inventory in each week, not the total amount of barrels stored. Since the levels of fuel that have to be stored each week depend on the orders, then we can conclude that the way we distribute barrels between orders Q_1 and Q_2 has an impact on the overall cost of storage.

The same reasoning that explains the increase in storage costs can be applied to money costs, because the latter is paid based on the levels of reserve, which – just as the levels of external storage – depend on the sizes of individual orders.

4.7. Example #6

Consider that fuel prices increase much more than the trigger differences.

Example #6 - Fuel prices increasing a lot more than trigger difference				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$23.00	\$25.00	\$24.00	\$27.00
Order Size (BBL)	38,657	0	33,386	0
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	38,657	38,657	72,043	72,043
% violations of Ai	6%	20%	14%	20%
Expected Costs	(\$)			
Fuel Cost	\$1,690,375	E[TSi] - LC (BBL)		
Money Cost	\$31,657	Week 3	7,836	
External Storage Cost	\$93,241	Week 4	0	
Reference TC	\$1,861,523	Week 5	7,790	
Total Cost TC	\$1,815,274	Week 6	0	
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.70	\$ 0.12	\$ 1.11	

Figure 23 – Summary of the most important variables of Example #6




Figure 24 – Graphical representation of demand, price and orders of Example #6

Analysis of Example #6

In this example, the differences between the prices in the second and fourth week and the prices in the first and third weeks, respectively, are higher than in the previous example. The model responds adequately to this situation, by building sufficient fuel inventory in the first and third weeks to avoid buying any fuel in the second and fourth weeks. In this case, a moderate response like that of Example #5 would not be the cheapest option: a radical response yields the greatest savings, because the savings in fuel cost compensate the increases in storage and money costs.

The total amount of the four orders together is the same as with all the previous examples: 72,043 BBL. However, only two orders are placed: Q_1 and Q_3 . The expected savings of the anticipated purchases, easily quantified with the help of the reference total cost, is \$46,249.

Remember that, for the sake of the model, the total lead time of an order is two weeks. BLM must assume physical responsibility for the fuel two weeks after the order has been placed. This means that the large amounts of fuel purchased in Q_1 and Q_3 result in large amounts of fuel that need to be stored when these orders arrive two weeks after they were placed.

In Example #6, the expected total amount of fuel that BLM will have to store ($E[TS_i]$) is larger than the local storage capacity (LC) for week 3 (as result of Q_1) and week 5 (as result of Q_3). We can calculate the expected need for external storage as $E[TS_i] - LC$. We should not calculate it as $E[ES_i]$, because ES_{ij} cannot take negative values. As you can see in Figure 23, the expected need of external storage capacity ($E[TS_i] - LC$) is 7,836 barrels for week 3 and 7,790 barrels for week 5. This is the first example to have an expected total stored fuel exceeding the local storage capacity. Every time BLM decides to take advantage of low fuel prices, it should be aware that large purchases could require leasing external storage capacity to store the excess fuel.

4.8. Example #7

Consider that fuel prices show two consecutive significant increases, starting in week 3.

Example #7 - Two consecutive significant increases				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$23.00	\$23.00	\$25.00	\$27.00
Order Size (BBL)	17,636	40,000	14,407	0
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	17,636	57,636	72,043	72,043
% violations of Ai	20%	14%	14%	20%
Expected Costs	(\$)			
Fuel Cost	\$1,685,803	E[TSi] - LC (BBL)		
Money Cost	\$31,243	Week 3	0	
External Storage Cost	\$95,888	Week 4	12,034	
Reference TC	\$1,836,869	Week 5	7,790	
Total Cost TC	\$1,812,934	Week 6	0	
All Q upper caps = 40,000 BBL	Is trigger diff value exceeded?			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	No	No	No	

Figure 25 – Summary of the most important variables of Example #7




Figure 26 – Graphical representation of demand, price and orders of Example #7

Analysis of Example #7

In this example, fuel prices show two consecutive significant increases, starting in the third week. We would expect the model to build up as much inventory as it could, to avoid purchasing expensive fuel when the price rises. This is exactly the result: the model purchases as much fuel in Q_2 as it is allowed to: 40,000 BBL. No previous example had shown an order size being constrained by the upper cap. The upper cap is necessary, and should be selected to reflect the reality of BLM's financial capacity and the supplier's fuel availability. The expected savings of the anticipated fuel purchase are \$23,935.

Notice that the model still purchases some fuel in Q_3 . This decision is correct and yields the lowest expected total cost: \$1,812,934. The reason the model buys this amount in Q_3 instead of buying it in Q_4 is that the price in week 3 is enough lower than the price in week 4 to justify the holding cost.

But why buy in week 3 at all? Consider, for example, that we decide to purchase no fuel in week 3, and instead we will buy that amount in week 1, that is: $Q_1 = 32,043$ BBL, $Q_2 = 40,000$ BBL, $Q_3 = 0$ BBL, $Q_4 = 0$ BBL. This way, we save in fuel cost, because we are buying cheap fuel in week 1 instead of buying it more expensively in week 3. However, the increase in holding cost (storage and money costs), caused by the huge inventory and reserves that result from Q_1 being so large, exceeds the savings of fuel cost. This alternative has an expected total cost of \$1,823,040, which is \$10,106 higher than the solution recommended by the model. The solution recommended by the model is the best.

As any example with large purchases, this one requires external storage in some weeks.

4.9. Example #8

Consider that fuel prices show only one significant increase in week 3.

Example #8 - Significant increase from one level to another level				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$23.00	\$23.00	\$26.00	\$26.00
Order Size (BBL)	17,636	40,000	0	14,407
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	17,636	57,636	57,636	72,043
% violations of Ai	20%	14%	20%	20%
Expected Costs	(\$)			
Fuel Cost	\$1,700,210	E[TSi] - LC		(BBL)
Money Cost	\$29,532	Week 3	0	
External Storage Cost	\$78,499	Week 4	12,034	
Reference TC	\$1,838,803	Week 5	0	
Total Cost TC	\$1,808,240	Week 6	0	
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.27	\$ 0.96	\$ 0.13	

Figure 27 – Summary of the most important variables of Example #8

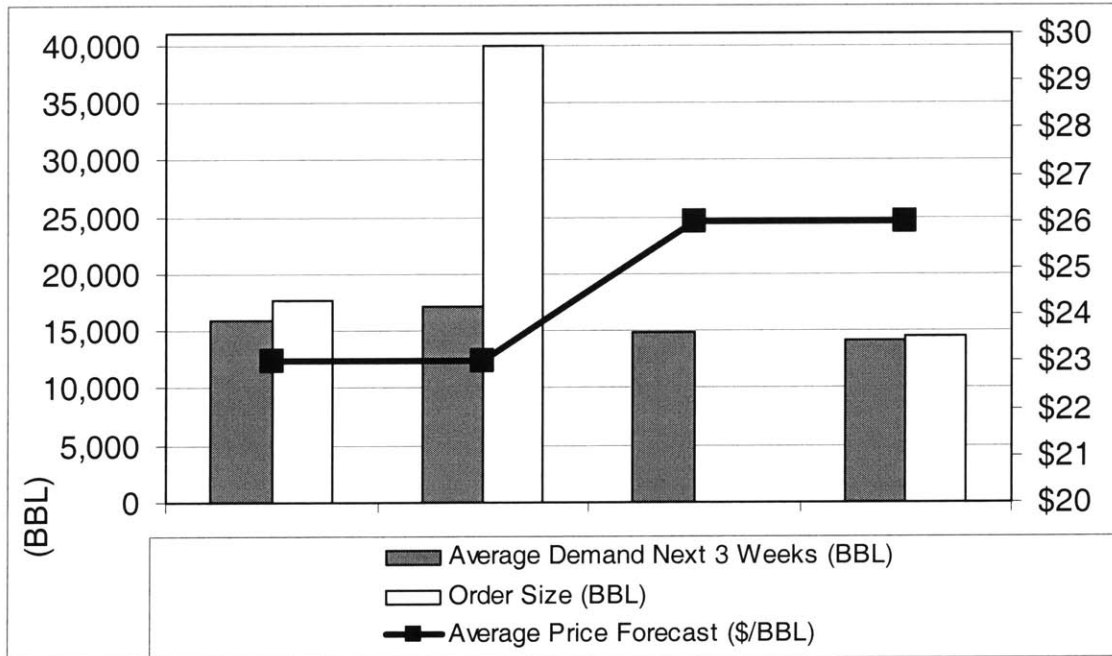


Figure 28 – Graphical representation of demand, price and orders of Example #8

Analysis of Example #8

Just as for Example #3, the rationale for this example is that, with its current supplier, BLM experiences no more than one change in price every four weeks.

In this example, fuel prices show a single significant increase in the third week. We would expect the model to build up as much inventory as it could in the second week, to avoid purchasing expensive fuel in the third week. This is exactly the result: the model purchases as much fuel in Q_2 as it is allowed to: 40,000 BBL, and buys nothing in Q_3 . The expected savings of the anticipated fuel purchase are \$30,563. Remember that calculating this expected savings is easy: just subtract the total cost from the reference total cost.

The reason the model buys no fuel in Q_3 , and buys in Q_4 , although both have the same price, is simple: holding costs make the difference. Why buy in week 3 when you can buy in week 4 at the same price, saving the holding costs? The reserve for week 3 was already purchased in week 2. In the previous example, the model decided to buy in week 3 instead of week 4 because there was a price increase in week 4 that justified the anticipated purchase.

4.10. Example #9

Consider that fuel prices significantly increase in week 2 and then decrease in week 4.

Example #9 - Significant increase, steady and then decrease				
Variable:	Week 1	Week 2	Week 3	Week 4
Average Demand Next 3 Weeks (BBL)	15,867	17,139	14,764	14,216
Average Price Forecast (\$/BBL)	\$23.00	\$26.00	\$26.00	\$23.00
Order Size (BBL)	40,000	0	16,310	15,733
Agg Q lower cap (BBL)	17,636	38,657	56,310	72,043
Agg Qs (BBL)	40,000	40,000	56,310	72,043
% violations of Ai	6%	20%	20%	20%
Expected Costs	(\$)		E[TSI] - LC (BBL)	
Fuel Cost	\$1,705,919		Week 3	9,179
Money Cost	\$29,909		Week 4	0
External Storage Cost	\$77,690		Week 5	0
Reference TC	\$1,854,668		Week 6	0
Total Cost TC	\$1,813,518			
All Q upper caps = 40,000 BBL	Trigger price differences (\$/BBL)			
All A max violations = 20%	From 1 to 2	From 2 to 3	From 3 to 4	
All R max violations = 20%	\$ 0.73	\$ 0.13	\$ 1.02	

Figure 29 – Summary of the most important variables of Example #9

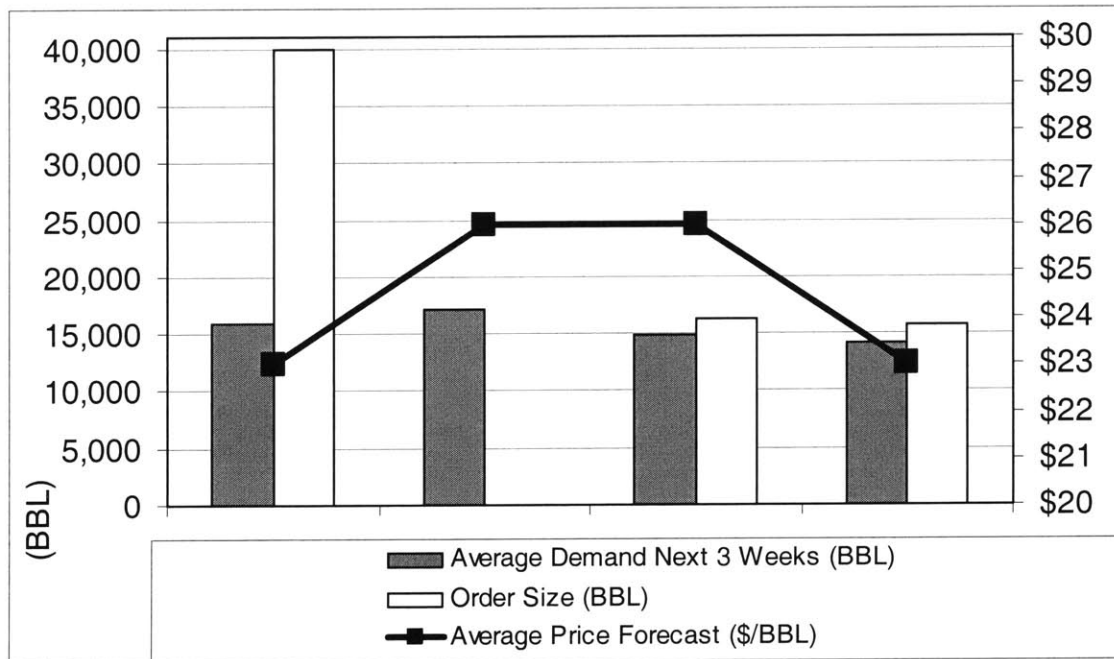


Figure 30 – Graphical representation of demand, price and orders of Example #9

Analysis of Example #9

This example depicts a case that BLM will not face with the current pricing structure of its supplier. However, it is offered for its theoretical value.

In this example, fuel prices show a single significant increase in the second week, and then a significant decrease in the fourth week. We would expect the model to build up as much inventory as it could in the first week, to avoid purchasing expensive fuel in the second week. The model recommends exactly that: to purchase as much fuel in Q_1 as possible: 40,000 BBL, and buy nothing in Q_2 . The expected savings of the anticipated fuel purchase are \$41,150.

Since there is no opportunity for speculative purchases in weeks 3 and 4, the model orders the same amount of fuel in Q_3 and Q_4 that it ordered in the early examples, just enough to satisfy the constraints and not a barrel more.

4.11. Unexpected results

As the model was being developed, many cases were run to test it. The nine examples that we have presented here are samples of the most typical and meaningful cases that the user will encounter, but are only a small fraction of the total of cases that we ran to test and understand the responses of the model.

As these cases were run in the model, two unexpected results called our attention, which we describe in this section.

4.11.1. No added value in price forecast being stochastic

We found that using a stochastic forecast of fuel prices adds no value to the model over using a single deterministic forecast. This was verified by comparing the results of many examples using a 5 scenario stochastic price forecast versus using a single deterministic price forecast (calculated as the weighted average of the stochastic values).

The results were found to be identical for the stochastic and deterministic fuel forecasts. This is explained by two facts: 1) there is no correlation between demand and fuel prices in the short run, and 2) there are no price-related constraints. Notice that the same cannot be said of the stochastic demand forecast, whose different scenarios play an important role.

This finding is relevant because preparing deterministic forecasts is much easier than preparing stochastic forecasts. Knowing that stochastic forecasts add no value will save BLM's staff some precious time.

4.11.2. Reserve violations are not constraining

In all examples we have run so far (around fifty in total), violations to the availability constraints have outnumbered the violations to the reserve constraints. Even manipulating every variable within their reasonable values, we did not find a single example where reserve violations exceed availability violations.

This finding is important because it suggests that, as long as the maximum acceptable probability for reserve violations is set to a value equal to or larger than that for availability violations, it is possible to make order size decisions based solely on the availability constraints, because the reserve constraints will also be satisfied. Obviously, if the maximum acceptable probability for reserve violations is set to a value smaller than that for availability violations, this finding is not relevant and both constraints will have to be considered.

For example, if the maximum acceptable probability for availability violations is set to 20% and that for reserve violations is set to 20% or more, then the user can expect that a solution that satisfies the availability constraints will also satisfy the reserve constraints. However, if the maximum acceptable probability for reserve violations is set to a much lower value, such as 5%, then it will probably become the binding constraint and must be considered separately.

However, currently there is no rational scenario where the user would have to set a maximum acceptable probability for reserve violations to a value lower than that for availability violations, because reserve constraints are soft and have no penalty, as opposed to availability constraints whose violation implies a stockout with dramatic economic consequences, measured in hundreds of thousands of dollars.

5. Conclusions and recommendations

In this chapter we present our conclusions, drawn from the experience of developing and using the model. But most importantly, we offer some recommendations to BLM related to possible opportunities for improvement through better supply chain management, which deserve further exploration.

5.1. Conclusions

After researching BLM's fuel management, in the frame of the Panamanian system and its regulations, we successfully developed a replenishment decision-support model using only Microsoft Excel. In this section we share some of the lessons we learned while developing this model.

BLM must have fuel available at all times, and a reserve equivalent to ten days of forecasted demand. Since we considered a horizon of four orders, and our demand forecasts include 50 scenarios, the availability and reserve requisites translate into four hundred constraints that our model has to satisfy. This large number of constraints is beyond the capability of Excel's Solver. To avoid putting BLM in the position of buying new software to run the model, we developed a workaround for this issue. It consists of four lower limits for the aggregate of the order sizes. A simple macro finds these lower limits, and passes them as constraints to the Solver. This way, the total number of constraints was decreased from 408 to only 18. Excel Solver satisfactorily solves this problem in a few seconds.

The model was designed to consider the forecast of fuel prices as stochastic. However, our analysis demonstrated that, since there are no price-related constraints and there is no correlation between demand and fuel price forecasts, a stochastic forecast does not add any value, and a deterministic forecast yields exactly the same results.

We found that fuel availability constraints are binding and are often the main constraint to the size of the orders selected by the model. On the other hand, reserve constraints are

not binding, because they are satisfied even when availability constraints have reached their limit. This means that it is necessary to select very carefully the maximum probability of stockout in each week, because it will usually be the determining factor in sizing the orders.

The performance of the model is satisfactory, because it responds as we would expect under a wide variety of scenarios. When the price forecast does not indicate a sharp increase in fuel prices, the model will buy conservatively, just enough fuel to satisfy the constraints. On the other hand, when there is the opportunity of saving money by purchasing excess fuel in anticipation of a price rise large enough to cover the holding costs, the model will buy the optimal amount of fuel in advance, taking advantage of the lower prices.

Although we are satisfied with this model, and know it will help the management of BLM to make their weekly replenishment decisions, we would like to make some recommendations related to important issues that, in our opinion, BLM should address to improve the overall performance of its supply chain and reduce costs.

5.2. Recommendations

The benefits that BLM can receive from the practice of good supply-chain management go well beyond the modest benefits of our model.

At this time, the lead time for a fuel order, from the moment the size of the order is decided to the moment the last truck arrives at the plant, is from 10 to 12 days, depending on the size of the order. The first truck leaves the supplier's facilities seven days after the order size is determined. Of these seven days of delay, three are caused by the internal process of BLM of getting authorization for the order and the bank issuing the purchase order. The remaining four days of delay originate outside BLM. Part of this time is used by the supplier to test the properties of the fuel, issue a report about these properties, send it to the certifying authority in Panama and obtaining the clearance for a tax exemption. It is our opinion that BLM should explore alternatives to reduce this lead time. Internally, BLM

should explore whether processes could be expedited to reduce the time used to approve the order and issue the purchase order. Externally, BLM could ask the supplier to find a mechanism to perform the testing of the fuel properties in advance and keep it ready for the time BLM places the order. This could reduce the delay in one or two days. It makes sense to streamline the fuel purchase process especially for summertime, when BLM is generating constantly for months.

The number of fuel suppliers is another issue that we think BLM should address. It has occurred in the past that BLM's single supplier has run out of fuel. We recommend that BLM keep a portfolio of suppliers. Although one of them could be the main supplier, BLM should always have more than one source of fuel, to reduce the probability of not having a supplier with available fuel to sell.

In the past, the stockouts of its supplier took BLM by surprise, further compounding the problem. This is a symptom of poor communication between BLM and its supplier, and a total lack of coordination of the channel. We recommend strengthening the communication of BLM and its supplier and the coordination of the channel. BLM should analyze the advantages that could result from sharing its demand forecasts with the supplier, and of receiving from the supplier the projected levels of fuel available locally for sale in the mid-term, e.g. two months into the future at any given time. Having an employee of the supplier visit BLM, and having an employee of BLM visit the supplier, or even working at each other's facilities, should be evaluated as an option.

Currently, the only form of hedging that BLM has is the dampening effect of buying fuel not at the daily price, but at the monthly average price. Although this avoids the daily peaks, it does not protect against sustained increase trends in fuel prices. BLM should evaluate the advantage of using a stronger hedging mechanism for its fuel procurement.

BLM has no formal fuel price forecasting system. We recommend that BLM evaluate the benefits of either purchasing fuel price forecasts from firms dedicated to this activity or

developing these forecasts internally. Purchasing millions of dollars worth of fuel based on eyeball estimates might not be the best approach for BLM.

Finally, BLM has no person in its staff dedicated to the coordination of its supply chain. It is not known whether the size of the company calls for a person dedicated to logistics, but it could be worth exploring the advantages that such a position could represent to BLM.

Appendix A - The Panamanian power system

With a population of 3 million people and a surface of 29 thousand square miles, Panama is a small country by most standards. Its power system and energy market are also small. Here we discuss the origin, privatization, and present state of the Panamanian power system.

A.1. Origin of the Panamanian power system

In 1954, the governments of Panama and the United States of America created an organization called SCIFE⁹, whose purpose included to research and evaluate the hydroelectric potential of Panama. From 1955 to 1960, some consulting firms were hired to recommend ways to develop the hydroelectric resources and study the problems of rural electrification in the central provinces¹⁰ of the country, whose power came from three private companies: Panamá Eléctrica S.A., Eléctrica del Interior S.A., and Santiago Eléctrica S.A.

In January 31, 1961, the government created the Institute of Hydraulic Resources and Electrification (IRHE¹¹), with 300 employees, to coordinate and expand the utilities services. It was also in charge of gas and telecommunications. In September 16 of the same year, IRHE took charge of the operation and maintenance of the utilities in the central provinces and other regions, including Panamá Eléctrica S.A. and Eléctrica del Interior S.A.

In 1969, less than a year after the military *coup d'etat* of 1968 that started 21 years of military dictatorship in the country, the government transformed IRHE into an autonomous institution, with capacity to promote the electrification of the Republic. In 1972, the

⁹ Servicio Cooperativo Interamericano de Fomento Económico

¹⁰ Panama encompasses nine provinces and five Indian reserves. The western provinces are Chiriquí and Bocas del Toro. The central provinces are Veraguas, Los Santos, Herrera and Coclé. The eastern provinces are Colón, Panamá and Darién. Panama City, the largest load center of the country, is found in Panama province. The largest hydroelectric resources are found in Chiriquí and Bocas del Toro.

¹¹ Instituto de Recursos Hidráulicos y Electrificación

company that supplied energy to the provinces of Panamá and Colón, called Compañía Panameña de Fuerza y Luz, was nationalized by the government and became part of IRHE. Its 84 MW of installed capacity represented 37% of the total capacity of the country; its 1,300 employees joined IRHE's workforce. The same was done in 1973 to Compañía Santiago Eléctrica, which supplied Veraguas, and Empresas Eléctricas de Chiriquí, which supplied Chiriquí. This added 25 MW and 721 new workers to IRHE. In 1974, the responsibility for telecommunications was transferred to a new institution, called INTEL¹².

In 1976, Bayano, a new hydroelectric plant in Panama province, with a capacity of 150 MW and a cost of \$104M, started operations. In 1979, two additional hydroelectric plants located in Chiriquí, with combined capacities of 90 MW and a cost of \$90M, started operations. The same year, the western part of the country, where most hydroelectric potential is located, was connected through 943 transmission towers and 230-KV transmission lines to Panama City in the eastern part of the country, where most loads are located. Several thermal generating units were added to the system, including three steam turbines with a total capacity of 120 MW, installed in Colon, and five gas turbines: two were installed in Panama in 1983 (43 MW, \$9M) and three in Colon in 1988 (60 MW, \$33M).

Led by the populist policies of the military government, the utility started to expand its generation capacity in increments disproportionate for the country's load of the time. The best example of this is Fortuna, a 300 MW hydroelectric project finished in 1984, when demand in the country was only 386 MW. Inefficient investment resulted in high capital requirements for installed capacity: Fortuna cost \$532M, twice its originally expected cost.

In 1985, Panama and Costa Rica connected its networks through a 230 KV transmission line. Up to this date, there is no connection between Panama and Colombia. In 1997, after decades of diverting utility revenues to finance other areas, the Panamanian government

¹² Instituto Nacional de Telecomunicaciones (National Institute for Telecommunications)

decided that IRHE's future as a government-owned utility was not bright and decided to do what many other countries were doing at that time: to privatize the utility.

A.2. Privatization of the Panamanian power system

The law that started the process of privatization of the electric utility, known as Law 6, was approved on February 3, 1997. This law defined the frame for the activities of generation, transmission, distribution and marketing of electric energy. A Restructuring Unit was created inside IRHE to create the necessary markets and at least 6 new companies from the pieces of the former utility, at the least cost and with the least impact on service and workforce. The restructuring process took 22 months. The consulting firm Arthur Andersen performed an operative and financial analysis and decided to divide IRHE in 8 companies, including 3 local distribution companies (LDCs), 1 transmission company (Transco) and 4 generation companies (Gencos). The public bid for the stock of these companies was held in May 28, 1998 for the Discos and in November 18, 1998 for the Gencos. Table 1 shows the names of the companies, the buyers, the amount paid, and the percent of stock sold. The government kept the rest of the stock, including 100% of the stock of ETESA¹³, the only transmission company in the country. The government received \$590M from these sales of stock.

Name of the Company	Description	Buyer	Price (Million \$)	% Sold
EDEMET ¹⁴ , EDECHI ¹⁵	2 Discos	Unión Fenosa	\$212M	51%
Elektra NE ¹⁶	1 Disco	Constellation	\$90M	51%
EGE Bahía Las Minas	1 Genco	Enron	\$78M ¹⁷	51%
EGE Fortuna	1 Genco	Hydro Quebec, Costal	\$118M	49%
EGE Bayano, EGE Chiriquí	2 Gencos	AES	\$92M	49%

Table 1 - Results of the privatization of the Panamanian electric utility

¹³ Empresa de Transmisión Eléctrica, S.A.

¹⁴ Empresa de Distribución Eléctrica Metro-Oeste, S.A.

¹⁵ Empresa de Distribución Eléctrica Chiriquí, S.A.

¹⁶ Empresa de Distribución Noreste, S.A.

¹⁷ Enron paid \$92M, but later received \$14M from the government as compensation for overcharge

A.3. Current structure of the Panamanian power system

A.3.1 Generation

In January 2004, 70% of generation was hydro; the rest was powered by fossil fuels.

There is no natural gas in Panama: fossil plants use fuel oil, marine diesel and light diesel. The average marginal cost¹⁸ in the energy spot market, from 1998 to 2003, is \$51.00/MWh.

The system has a total installed capacity of 1.5 GW. EGE Bahía Las Minas, now called BLM Corp, has a capacity of 280 MW. EGE Fortuna has 300 MW. EGE Chiriquí and EGE Bayano were merged in one company, now called AES Panamá; it built several new hydroelectric projects, for a total of 550 MW. Since 1997, four privately-owned Gencos were installed: PEP (60 MW), PanAm (96 MW), Copesa (44 MW), and Pedregal (53 MW). The rest of the generation comes from the generation units of the Panama Canal Authority.

A.3.2 Distribution

The maximum demand registered in the Panamanian system to the date is 883 MW. Distribution and retail sales are in charge of local distribution companies (LDC). LDCs buy power in bulk through contracts and the spot market, deliver it, and sell it to consumers in their area. There are three LDCs in Panama: 1) EDEMET, which distributes power to the central provinces and half of Panama City, has 50% of the energy demand. 2) Elektra NE, which distributes power to the eastern provinces and the other half of Panama City, has 40% of the energy demand. 3) EDECHI, which distributes power to the western provinces, has 8% of the energy demand. The remaining 2% of energy is sold retail to large customers (see A.3.3).

¹⁸ In Panama, marginal cost includes only variable generation costs: fuel, operation and management costs.

A.3.3 Retail Sales

There is no retail wheeling in the Panamanian market for customers with demand smaller than a certain limit. Customers above this limit are called large customers, while those below the limit are called regulated customers. Originally the limit was 500 kW. Currently, it is 200 kW, and in 2005 it will be 100 kW. Regulated customers have no choice of their supplier: they must buy from their LDC. Large customers can buy energy either from the LDC, or from other agents, including Gencos. When a large customer buys energy from a source other than its LDC, it pays the LDC and the Transco for the use of their networks.

A.3.4 Transmission

ETESA, Panama's only Transco, is fully owned by the government. Its purpose is to move power in bulk quantities from the generation sites to the delivery sites. Distribution is not performed by ETESA: this is the job of the LDCs. ETESA owns and maintains the transmission facilities, and performs many management and engineering functions to ensure that the system continues to work. ETESA is paid for the use of the transmission lines and other equipment: Gencos and LDCs pay half of the transmission expenses each. The system and the market operators are structurally part of ETESA, but they function as independent entities.

A.3.5 System operation

The system operator is called the National Dispatch Center (CND¹⁹). It is in charge of forecasting, long-term planning, short-term scheduling, unit commitment, real-time operation with safety and stability considerations, and ex post evaluation of the operation of the system.

¹⁹ Centro Nacional de Despacho

A.3.6 Market operation

The market operator, which is itself a unit inside CND, is called the Direction of Electrical Wholesale Market (MME²⁰). It is in charge of authorizing supply contracts and managing the energy and capacity spot markets. MME notifies the agents how much energy was sold, from whom, and to whom. It also allocates the payments for ancillary services.

A.3.7 Regulatory Institution

The regulatory institution is the Public Utilities Regulating Entity (ERSP²¹), created in 1996 by the government to regulate all public utilities, including electricity and telecommunications.

²⁰ Dirección del Mercado Mayorista de Electricidad

²¹ Ente Regulador de los Servicios Públicos

Biographical Reference

Roberto Perez-Franco is an MIT Master of Engineering in Logistics (MLOG) Candidate. Prior to moving to Cambridge, Roberto worked in Panama as Contracts Administrator for the Marketing Department of BLM Corp, the largest fossil fuel generating company in Central America. He was involved in forecasting and calculating the sale of energy and power of BLM Corp to the contract and spot markets in Panama and Central America.

Roberto holds a B.Sc. in Electromechanical Engineering from Panama Technological University. He was awarded the Fulbright and Barsa scholarships to study Logistics at MIT.

Born in Chitre, Panama, in 1976, he has published three short fiction collections, and appears in several national and international short story anthologies. Fluent in Esperanto, he is a member of the Universal Esperanto Association, Mensa, ISPE and CLM. He lives in La Heroica Villa de Los Santos, Panama, which named him Meritorious Son in 1998.

Web-site: <http://perez-franco.com>

E-mail: roberto@perez-franco.com

Bibliographical Reference

1. Chopra, Sunil; Meindl, Peter. (2000) "Supply Chain Management", 1st edition, 544 pages, Prentice Hall College Div, ISBN 0130264652.
2. Eydeland, Alexander; Wolyniec, Krzysztof. "Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging", 1st edition, 504 pages, John Wiley & Sons, ISBN 0471104000.
3. Morris, Peter A. ; James Sandling, M.; Fancher, Richard B.; Kohn, Michael A.; Chao, Hung-Po; Chapel, Stephen W. (1987) "A Utility Fuel Inventory Model". Operations Research, Vol. 35, No. 2 (Mar.-Apr., 1987), 169-184.
4. Oilnergy - WTI Crude Oil Posted Price: <http://www.oilnergy.com/1opost.htm>
5. Patel, Babul, Ph.D. (1981) "A System Integration and Optimization Model for Fuel Management and Scheduling of Power Generators." The University of Nebraska – Lincoln (Source: DAI, 42, no. 04B, (1981): 1563)
6. Shahidehpour, Mohammad; Yamin, Hatim; Li, Zuyi. (2002) "Market Operations in Electric Power Systems", 1st edition, 450 pages, John Wiley & Sons, ISBN 0471443379.
7. Shen, Yung-Tang, Ph.D. (1996) "Studies in fuel supply and air quality planning by electric utilities." The Ohio State University (Source: DAI, 57, no. 02A (1996): p. 891)
8. Silver, Edward A.; Pyke, David F.; Peterson, Rein. (1998) "Inventory Management and Production Planning and Scheduling", 3rd edition, 784 pages, Wiley Text Books, ISBN 0471119474.
9. Stoft, Steven. (2002) "Power System Economics", 1st edition, 496 pages, John Wiley & Sons, ISBN 0471150401.
10. Wood, Allen J.; Wollenberg, Bruce F. (1996) "Power Generation, Operation and Control", 2nd edition, 592 pages, Wiley-Interscience, ISBN 0471586994.