Flexibility in Transportation Procurement: A Real Options Approach

By

Georgi Zhelev Zhelev

B.A. Economics
Williams College, 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

June 2004

Signature of Author..
Engineering Systems Division
May 7, 2004

Certified by..............................
Christopher Caplice
Executive Director, Master of Engineering in Logistics Program
Engineering Systems Division
Thesis Supervisor

Accepted by ....
Yosef Sheffi
Professor, Engineering Systems Division
Professor, Civil and Environmental Engineering Department
Director, MIT Center for Transportation and Logistics

Flexibility in Transportation Procurement: A Real Options Approach

by

Georgi Z. Zhelev

Master of Engineering
in Logistics

Submitted to the Center for Transportation and Logistics on May 7, 2003 in partial fulfillment of the requirements for the degree of

Master of Engineering
in Logistics

Abstract

Shippers who use for-hire carriers need a careful estimation of their need for trucking capacity in order to be able to haul products in a timely and cost-effective manner. This thesis outlines a flexible approach to transportation procurement. According to the proposed method, a shipper makes a firm commitment to a carrier to use a guaranteed level of capacity and agrees to pay a penalty for any unutilized equipment. In exchange, the carrier provides the shipper with an option to fulfill the shipper's request for additional capacity above the guaranteed level and up to a predermined limit if needed. The analysis focuses on actual data from a shipper and explores a "what if" scenario – whether it would have been possible to achieve savings and generate value by having such a flexible contract.

Thesis Supervisor: Christopher Caplice
Title: Executive Director, Master of Engineering in Logistics Program
Engineering Systems Division

To Mitko Fitkov

Acknowledgements

I would like to thank:

My thesis advisor, Chris Caplice, for his great patience with me and the advice and insights he provided in helping me complete this thesis.

Akihiro Kaga and Ee Learn Tay, who helped me understand better the intricacies of real options analysis.

Gary Wicker of JB Hunt, Kevin McCarthy of C.H. Robinson, and Matt Stalter of Pepsico/Quaker for their professional insights and contributions to this work.

My classmates from MLOG '04, who made this year fun and less painful. A special mention to grape historian Antoine.

My parents and my little brother for their love and support during hard times.

Table of Contents

List of Tables and Figures

Chapter One

1.1 Introduction

One of the constant problems facing businesses is how to deal with uncertainty. No matter how accurate forecasts are, ultimately they are wrong. For example, a fashion apparel manufacturer can rarely if ever predict the demand of its latest collection, let alone for the individual items in it. In order to deal with uncertainty, it is essential to build systems that are flexible and can change dynamically reflecting the latest available information and allow for a wide range of actions to be taken depending on current developments.

This thesis focuses on the uncertainty faced by shippers when procuring trucks needed for sending their products to customers. For many firms, transportation costs represent a significant part of their overall expenditures. As a result, even small reductions in those costs can have a large impact on profitability. If a shipper knows with certainty the future demand for his products, then he can accurately estimate his needs for trucking capacity and procure it at the lowest possible price. If, however, demand pattern changes, suddenly the shipper finds itself in a situation where there is a mismatch between the capacity that he needs and the actual capacity he will get. In that case, what initially seemed to be an optimal assignment that minimizes costs may in fact turn out to be inefficient.

In order for the shipper to acquire more flexibility that will enable him to respond quickly and adjust his needs to constantly changing demand levels, we propose to use option contracts in transportation procurement. In the field of finance, financial option contracts are widely used to hedge against risk and uncertainty. Although not as common, real option contracts are becoming increasingly popular in the non-finance related sphere. Unlike their financial counterparts that are centered around financial, or intangible, assets (stocks, currencies, etc.), real options deal with tangible assets (thus the name "real") such as production capacity, production technology, etc.

1.2 Thesis outline

The thesis is organized as follows. Chapter 2 provides an overview of the commercial transportation market in the U.S., the latest trends and developments. Chapter 3 introduces the Real Options methodology, provides examples of industries in which it has been used and models and tools that are used to perform the analysis. Chapter 4 presents analysis of historical shipping data from a consumer packaged goods company, determines appropriate parameters for a flexible option contract for transportation procurement and uses simulation to verify the cost advantages of such a contract for a future unknown demand. Chapter 5 summarizes the findings, presents conclusions and suggests areas for future research.

Chapter Two

2.1 U.S. Commercial Freight Transportation Market

The U.S. commercial freight transportation market is valued at about $790.6 billion in 2002 (S & P Industry Surveys, Dec 2003). The trucking business represents more than 80% of this market with $585.3 billion in 2002. Trucking itself is divided into three main groups: Private, Truckload (TL), and Less-than-truckload (LTL). Each is discussed in turn.

2.1.1 Private carriers

Private carriers are truck fleets that are operated by companies who want to haul their own raw materials and products without resorting to for-hire carriers (TL or LTL). They account for the highest percentage of trucking operations, valued at $276.7 billion in 2002. Private carriers are employed by companies from various industries – wholesale and retailing, food distribution, manufacturing. Firms that use their own private fleets opt for vertical integration with their own transportation divisions to guarantee themselves a high level of service and reliable shipments of their raw materials and finished goods.

2.1.2 Truckload carriers

TL carriers hold the second largest share of trucking revenues after private carriers ($250.2 billion in 2002). These carriers transport large shipments of a single shipper from origin to destination. This key feature leads to a very high fragmentation in the TL sector. Without the need for expensive investments in terminals and following the deregulation of the industry, anyone who possesses a truck can become a carrier. According to a report published in 2000, this is one of the three most fragmented business sectors (CSFB, 2000) with over 30000 companies of different sizes. While most of these are relatively small private operations with annual revenues not exceeding $1 million, bigger publicly traded companies such as J.B. Hunt Transport Services., the top TL carrier, had revenues of $2.31 billion in 2002 (S & P Industry Surveys, 2003).

2.1.3 Less-than-truckload carriers

This sector is the smallest of the three with annual revenues of $58.4 billion in 2002. LTL carriers have an extensive terminal network through which they consolidate numerous small shipments and then transport them to individual customers. Unlike the TL sector, LTL has serious barriers to entry because of the high infrastructure costs. Another factor is the high degree of unionization compared to the TL carriers which translates into higher labor costs and thinner profit margins.

2.2 Problem formulation

According to a recent survey (Bear Stearns 2004), TL capacity has tightened over the past year, while LTL capacity has remained relatively balanced. Expectations for the TL sector are that capacity will become even tighter. Kevin McCarthy of C.H. Robinson echoes these findings by observing that for the first time in the past 6-7 years shippers find themselves in a "sellers" market and that trend will only grow stronger as there has been no capacity addition by the carriers (McCarthy, 2004).

Tightening capacity poses a problem for shippers. One of the key problems faced by shippers using truckload carriers is the accurate estimation of their need for trucks because if they underestimate it, it will be expensive to procure the needed capacity later when their true needs become known. On the one hand, shippers do not want to order too much capacity from the carriers. At the same time, shippers do not want to order too little and run out of trucks. Here are two typical situations that might arise.

Scenario 1

The shipper has underestimated demand and faces a shortage of trucks. In that case, he would have to request additional capacity from the carriers. They may not have it since capacity is fairly tight. In that case, the shipper will be forced to turn to the spot market. The rates, however, are going to be much higher than his usual contract rate with a carrier and will require greater effort.

Scenario 2

The shipper has overestimated demand and faces a surplus of trucks. This time it is the carrier who is unhappy because he has committed his trucks only to find out that they are not needed. The profit margins for the TL carriers are typically very low (2-3%) and every idle unit is foregoing valuable revenue, so from the carrier's perspective equipment utilization has a very high priority. Without high equipment utilization rates, the carriers face extinction for not being able to achieve a return on assets that allows them to keep their operations going.

Is it possible to avoid such situations and ensure that both shippers and carriers use the limited capacity more efficiently? One approach that a shipper could use would be to diversify among many carriers. When he faces an unexpected surge in demand, he is more likely to find the additional trucks if his carrier base is wider. This, however, has implications for the service level provided by the carriers. If a shipper uses many carriers and allocates small volumes to each of them, the carriers will probably dedicate more efforts to serving their bigger customers than a shipper that represents a very small share of their business.

When a shipper wants to ensure high service levels and chooses to allocate most of its volume to a few dedicated carriers, the capacity problem cannot be easily solved by requesting additional trucks because the carriers might not have any available. If, however, the shipper has a contract with the carriers according to which he has the opportunity to request additional capacity if needed and the carriers are prepared to fulfill such a request, the shipper obtains a greater flexibility in handling unexpected demand spikes without incurring prohibitive transportation costs.

Such an approach was employed by Reynolds Metals Company in the late 1980s (Gorban, Moore, Warmke, 1991). The company, which at that time was a Fortune 75 metals producer, used to have about 200 van and flatbed carriers in order to serve its interstate truckload freight operations. This presented a few disadvantages. The large number of carriers decreased the leverage that the company could have on any single carrier to lower rates. Service levels varied greatly among carriers. A serious issue came as a result of decentralization. Each of the company's 200 plants, warehouses, and suppliers chose their respective carriers and made their dispatches without coordinating their efforts with the other players in the company's network.

This prompted the company's management to move to a centralized system in an attempt to raise the quality and reliability of the transportation network and at the same time stabilize and reduce the costs. As a result, Reynolds Metals decreased its carrier base to 14 premium-quality carriers, improved service levels, and reduced annual freight expenditures by $7 million. Each of the carriers was guaranteed to serve a large portion of Reynolds' freight and this provided incentives for the carriers to follow stringent rules and meet high service standards while dedicating a large number of trucks to fulfill the company's requirements.

Because Reynolds Metals reduced significantly the number of carriers, it exasperated the problem of dealing with demand uncertainty. It could no longer rely on 200 carriers that provided more flexibility in obtaining extra capacity if needed. In order to deal with peak demand instances, the company devised two truck commitment strategies.

- fixed-truck commitment

For high-volume, high-variability locations, the company required a carrier to provide a fixed number of trucks daily. If for some reason, demand was lower than expected, the company paid a penalty for the unused trucks.

- variable-truck commitment

This strategy was established for <u>all</u> locations, including those covered by fixed-truck commitment. It required the carriers to supply a location with a variable number of trucks daily up to a weekly maximum that had been established. The company was not required to pay a penalty for any of the unused "variable" trucks.

The experience of Reynolds Metals Company provides an insight to the question posed in the beginning of the chapter, i.e. how can the shippers and carriers better contract so that they can use the limited truck capacity more efficiently. Suppose that the shipper wants to have flexibility without sacrificing his service level. If he can design a contract according to which he can request additional trucks from his carriers if he needs them, he could still enjoy the benefits of lower rates and have the desired flexibility. From the carrier's perspective, it is important to ensure that he utilizes fully his total capacity by either using his trucks or being compensated for any unused trucks.

In effect, such an arrangement represents an options contract. The shipper has the option to request additional trucks only if he needs to handle peak demand. At the same time, the carrier has the option to sell his unused trucks to the shipper if demand turns out to be less than expected and the shipper does not need them. The two parties have seemingly conflicting interests – the shipper is more worried about lacking capacity, whereas the carrier is apprehensive about providing too many trucks and finding some of them unused. In order to reconcile their difference and provide a mutually beneficial solution, we propose to use a real options approach that will help determine the optimal parameters of the contract between shipper and carrier.

Chapter Three

3.1 Real Options

A real option is the right but not the obligation to take an action (for example, contract or expand a plant, switch to alternative fuel in power generation, abandon a research project, etc.) at a predetermined cost (exercise price), during a specified period of time (the option life), (Brach 2002). Real options provide managers with flexibility by allowing them to take an action only if it will be beneficial. If not, they simply choose not to exercise the option and in this way can limit their losses.

Suppose we bought a license to make a video game based on the latest episode of a movie. Let's assume that our revenues from selling the game are R while it costs us C to produce. The revenue R depends on the success of the movie. If it proves to be a great hit, R will be sufficiently high so that $R - C > 0$, i.e. our payoff is positive and we will make the game. If the movie is disappointing, then demand and revenue will be low and $R - C < 0$. In this case, however, we can choose not to make use of the license and simply abandon the game project, thus saving further losses.

The value of flexibility that Real Options provide cannot be underestimated. Seemingly unprofitable projects that are considered in a static framework where all the key variables are pre-defined and fixed can in fact prove to be worth undertaking when we add the possibility to vary one's actions and decisions based on constantly updated information that reduces uncertainty. In this respect, the real options framework is a much closer approximation to the real world since it allows for active management, interaction with the business environment, and the opportunity to react to the latest developments instead of being a constrained and passive observer.

3.2 Types of Real Options

There are several main types of real options. The following examples are adapted from (Mun, 2002) and (Brach, 2003).

Table 3-1 Real Option Types

Option Type	Gives Owner the Right to
Deferral Option	Postpone decision in order to gather more information
Expansion Option	Increase capacity if conditions are favorable
Contraction Option	Decrease capacity if conditions are unfavorable
Abandonment Option	Scrap an unprofitable project
Switching Option	Change input/output parameters, production technology
Option to Choose	Adjust capacity based on market conditions
Compound Option	Make a decision based on subsequent, conditional steps

3.2.1 Deferral Option

The option to defer reduces uncertainty by delaying the decision to take action until more relevant information is gathered. For example, a chemical producer might delay its decision to build a plant for a new product until it develops a better understanding of the compound performance and whether it will be well received by the market. A drug manufacturer might choose to wait before committing itself to building an expensive plant for a new product if it is not certain whether the new drug will get government approval. In these cases, the uncertainty can be partially or completely resolved by waiting (the chemical producers might get an initial idea of how the market will react to the new product whereas the drug manufacturer will know whether its drug has passed government approval).

3.2.2 Expansion Option

This option allows its owner to expand his operations if market conditions turn out to be more favorable than previously expected. For example, a computer disk drive manufacturer faces uncertain demand for its latest product. If it builds a facility that has a capacity much bigger than even the most optimistic scenario, it can be certain that it will meet demand and never stock out and lose sales. However, if demand is low, the facility will be underutilized and, considering the short product life cycle for computer components, the company might end up with piles of obsolete inventory. Instead of building a huge plant, the company can do better by investing in a facility that will be able to cover the minimum expected demand with the option to expand it and increase capacity if demand increases.

3.2.3 Contraction Option

The contraction option is the opposite of the expansion option. It enables a company to scale down if market conditions are unfavorable. For example, an airline company might have experienced a period of rapid growth during which it considerably increased its fleet and terminal network in an attempt to gain market share from its competitors. If suddenly demand for the company's services declines it can choose to close down some of the unprofitable destinations and decrease the number of planes it operates in an attempt to avoid incurring losses and improve profitability. Without the option, the airline will have to continue operating on routes that are diminishing its value instead of increasing it.

3.2.4 Abandonment Option

It represents the right to dispose of an asset and recover salvage value once market conditions deteriorate enough so as to make keeping the asset unprofitable. This option represents a hedge against market downturns. Suppose a company is considering a costly investment in the development of a new product. It is estimated that the research will take three years and if successful, the product will be launched, if not – the company may choose to resell any patents that it has registered for the product to another firm. More specifically, the company has the right, but not the obligation to resell its research at any point during the next three year to another interested party, thus it is guaranteed some salvage value for its efforts and the value of the project would be greater compared to a base case without such an option.

3.2.5 Switching Option

This is a very important option that offers considerable value in industries where inputs to the production process can be substituted. If the price of an input relative to its substitute rises, then the substitute becomes more attractive and the company may begin to use it instead of the original input in the production process. Such an option requires flexible facilities to be built, i.e. equipment which can use multiple inputs. Alternatively, it could be the outputs that change depending on their relative prices while keeping inputs the same. Having such facilities provides the company with more cost-effective production capabilities. For example, a power plant that can use both oil and coal to generate electricity is inherently more flexible than a plant that can use only one of the two inputs. Even though the dual-use equipment of the flexible plant is likely to be more expensive, it provides a hedge against future uncertain price fluctuations of the production process inputs (oil and coal in this example) and justifies the investment in more expensive technology in the long run.

3.2.6 Option to Choose

This option allows management to alter the scale of operations in response to changing market conditions. Suppose a company has successfully accomplished a research phase and introduced a new product. In the beginning it may not have enough information about the expected demand so it may choose to set up an option that allows it to acquire additional capacity (either by building it or purchasing it from another company) should demand prove very strong and exceed its current capacity. If demand remains the same or even goes down, then the firm needs not exercise the option and can limit its spending on overcapacity. Suppose the company was instead very optimistic about the new product launch and decided to build a greater capacity in anticipation for future growth. However, as a hedge against declining demand, it can set up a contraction option which allows it to sell a pre-determined share of its total capacity to another firm in case demand forecasts were not met. In this case, the company manages to limit its losses if the market environment is unfavorable. Note that the option to abandon is an extreme case of the option to contract (100% contraction).

3.2.7 Compound Option

Many projects usually pass through a series of steps, each of which relies on the successful completion of the previous one. After each step, the management reviews the outcome and then decides whether to proceed or not. Such an option is compound: it is an option on an option. A good example is the development process of a new drug. It starts with pre-clinical research aimed at producing a molecule for the new drug. If successful, the company has the option to take the chemical compound to the next stage, phase I trials. Upon successful completion of this stage, the company may proceed to phase II trials and so on, until the drug is filed for approval by the FDA. This kind of flexibility allows management to consider the investments as a series of decision that build upon each other without necessarily committing investment upfront for stages that might never become viable.

3.3 Possible Real Option Types in Transportation Procurement

The expansion/contraction option seems to be the most obvious RO because it allows for responding and adjusting capacity to market conditions. If the shipper experiences surges in demand, he can request additional trucks from his carriers while during downturns he can choose to decrease the number of trucks used.

A switching option could also be valuable if it allows the owner of the option to switch to equipment which best suites his current product needs. For example, during extreme weather conditions a beverage manufacturer will require specialized trucks that protect his products from very low or high temperatures. Having such an option would allow the shipper to obtain the necessary trucks before other competitors.

3.4 Real Options Valuation

Obviously real options provide a certain value to their owners but the question is how much exactly is an option worth and what are the main determinants of the option value. The value of the option is defined as the difference between the return from a flexible investment program versus the return from an inflexible program (Brach 2002). Real options are valued similarly to their counterparts – the financial options. There are five main determinants of the option price.

1. Value of the underlying asset
2. Volatility
3. Time to expiration
4. Risk-free rate
5. Exercise price

3.4.1 Real Options vs. Financial Options

The underlying asset is the basis for option valuation. For financial options it is simply the price of the stock on which the option is based. In the case of real options, the underlying asset is more difficult to define. Since it is difficult to find a publicly traded stock which to use as an underlying asset, some authors assume that the present value of the project (without the flexibility provided by the option) on which the option is based to be the underlying asset (Copeland and Antikarov, 2001). There is a fundamental difference in this respect between financial and real options. Whereas the management of the company can potentially influence present value of the project it has undertaken, the owners of financial options have no control over the stock prices.

Volatility represents the speed at which the value of the underlying asset tends to diverge away from today's value or how hard it is to predict the price of the underlying asset into the future

(Howell et al, 2001). Higher volatility means larger divergence from today's value of the underlying asset. It may seem counterintuitive, but the higher the volatility of the underlying asset, the more valuable the option becomes. This is because real options introduce flexibility into managers' decisions and a higher volatility implies that there is a greater chance of the project reaching such a value that exercising the option is justified.

Time to expiration is the period over which the option may be exercised. The longer it is, the more valuable the option. Again, this follows from basic intuition about uncertainty – the farther we go into the future, the more difficult it is to predict all possible outcomes, the uncertainty increases and with it the chance of the option becoming more valuable increases as well.

The exercise price is the amount of money invested to exercise the option. In the expansion option example, it will be the money spent to increase capacity. If the exercise price goes up, the value of the option decreases because it becomes more expensive to be exercised.

The risk-free rate indicates the rate that the market is willing to pay for an asset whose payoffs are completely predictable (i.e. there is no risk associated with them). The higher it is, the greater the value of the option.

3.4.2 Option Pricing models

There are several approaches that are used to solve real options problems: closed form equations (Black-Scholes model), binomial method, and Monte Carlo simulation (Mun 2002). The Black-Scholes formula, although widely used to price financial options, has certain assumptions that make it impossible to be used to price real options. For example, it assumes constant volatility over time for the project that is valued, a pre-determined expiration date for the option, normal distribution of asset returns (Brach 2003). These conditions, though satisfied by financial options, rarely hold true for real options.

3.4.3 Binomial Pricing Model

This model was developed in 1979 by Cox, Ross and Rubinstein (Cox, Ross, Rubinstein, 1979). The strength of the model lies in its simplicity. It does not require any knowledge of complex mathematics and it is understandable to a wider audience, not intimidating potential users (Mun, 2002). The binomial option model divides the project life into discrete time steps (weeks, months, years). It uses binomial trees to approximate the behavior of the underlying asset. The model assumes that at each step the value of the asset can either go up or down, with a probability of p and 1-p, respectively. The magnitude of the down movement is inversely proportional to the up movement, i.e. $u = 1/d$. Let's assume that the value of the asset today is S_0 and it can go up to uS_0 or down to dS_0.

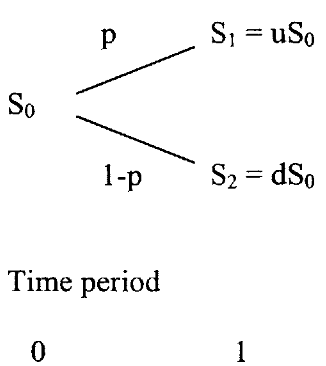

The option price is then calculated backwards starting at the end of the tree and discounting the value of money at the risk free rate.

3.4.4 Monte Carlo Simulation

Deriving its name from the casino haven in Monaco, this approach is a stochastic technique based on the generation of random numbers to simulate behavior of uncertain inputs (Mun 2002). If our underlying asset follows a path that cannot be replicated using a binomial lattice, we can use a simulation to generate a probability distribution for the value of the option. After defining the model that simulates option value through a set of equations, we can assign probability distributions to the key variables. Then perform a series of trials and draw random values from the previously defined probability distribution generating a range of option values (Mittal, 2004).

3.4.5 Example of Real Option Calculation

We need the following inputs in order to calculate the value of an option

S – the value of the underlying asset

X – present value of the implementation cost of the option

σ – volatility of the underlying asset

T – time to expiration in years

rf – risk-free interest rate

We can then calculate the size of the up and down movements and the probability associated with them.

Their derivation is based on the risk-neutral valuation principle in the financial world. Risk neutrality assumes that the expected return from traded securities is the risk-free interest rate and present value of expected future cash flows can be obtained by discounting at the risk-free interest rate (Hull 1991).

The expected return of an asset is the risk-free interest rate, r. Therefore, after a time interval dt, the asset is expected to have a value of S_0e^{rdt}. At the same time, the value of the asset is expected to be $puS_0 + (1-p)dS_0$ as depicted in 3.3.3.

We have

$S_0e^{rdt} = puS_0 + (1-p)dS_0$

$e^{rdt} = pu + (1-p)d$ (1)

Therefore, $p = \dfrac{e^{(rdt)} - d}{u - d}$ (follows from (1) after rearranging and solving for p)

The up and down movements in the asset price represent proportional changes. The standard deviation of the proportional change in asset price in a short time interval δt is defined as $\sigma\sqrt{dt}$.

Then the variance of the actual change in dt will be $S^2 \sigma^2 dt$.

Taking into account that $Var(X) = E(X^2) - E(X)^2$

$$S^2 \sigma^2 dt = pS^2 u^2 + (1-p) S^2 d^2 - S^2 (pu + (1-p)d)^2$$

$$\sigma^2 dt = pu^2 + (1-p)d^2 - (pu + (1-p)d)^2 \qquad (2)$$

It follows from (1) and (2) (Hull 1991), provided that dt is small, that

$$u = e^{\sigma\sqrt{dt}}$$

$$d = e^{-\sigma\sqrt{dt}} = 1/u$$

We first need to find the value and volatility of the underlying asset. Finding the value of the underlying asset for a real option is tricky. Unlike financial options for which the underlying asset is simply the value of the stock, there are no publicly traded securities that can be tied to the real option. An approach suggested by Copeland and Antikarov (2001) relates the value of the underlying asset to the present values of the cash flow of the project without the managerial flexibility.

Finding the volatility is relatively easy if we can estimate the future cash flows that the project is expected to generate (Mun 2002).

Table 3-2 Volatility Estimation

Time Period	Cash Flow	Relative Returns	Ln(Relative Returns) = X
0	C_o	-	-
1	C_1	C_1/ C_o	$Ln(C_1/ C_o)$
2	C_2	C_2/C_1	$Ln(C_2/C_1)$
…..	….	….	….
N	C_n	C_n/C_{n-1}	$Ln(C_n/C_{n-1})$

Volatility = $\sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(x_i - \overline{x})2}$

Note: if the cash flows were not based on annual date, we need to convert the calculated volatility to an annual base.

Annual volatility = Volatility ($\sigma\sqrt{T}$)

where T is the number of periods per year. E.g. if our cash flows were monthly data, T = 12, weekly – 50, etc.

If we have historical data for past cash flows, we can estimate the volatility and use the value to simulate future cash flows. The problem for real options is that very often we need to estimate the volatility of cash flows associated with projects for which no historical data exists. In that case, we can estimate volatility using simulation software that runs different scenarios about possible future outcomes.

3.4.6 Real Option Numerical Example

As an illustration, we provide a simple example for a deferment option. Suppose a company considers a project that has a net present value of $100 and implementation cost of $100. What will be the value for the company to wait one more year before investing in the project (i.e. have a real option to defer investment in order to gather more information and resolve some of the uncertainty of the project)?

Figure 3-1 Real Option Example

Expiration in Years	**1**	**Number of Time Steps**	**5**
Volatility	**0.45**	**Time Step Size (dt)**	**0.2000**
PV Asset Value	**100**	**Up Jump Size (u)**	**1.2229**
Risk-free Rate	**0.05**	**Down Jump Size (d)**	**0.8177**
Dividend rate	**0**	**Risk-Neutral Probability (p)**	**0.4747**
Strike Cost	**100**		

	Step 0	Step 1	Step 2	Step 3	Step 4	Step 5
	100.00	122.29	149.55	182.89	223.67	273.53 A
		81.77	100.00	122.29	149.55	182.89
			66.87	81.77	100.00	122.29
Lattice 1				54.68	66.87	81.77
					44.71	54.68
						36.56

	Step 0	Step 1	Step 2	Step 3	Step 4	Step 5
	20.75	34.35	55.07	84.87	124.66	173.53 B
		8.86	16.28	29.20	50.55	82.89
			2.31	4.92	10.48	22.29
Lattice 2				0.00	0.00	0.00
					0.00	0.00
						0.00

Lattice 1 traces the possible outcomes of the project value based on the up and down movements from the beginning value. The last column in Lattice 2 represents the corresponding values from Lattice 1 – $100 (the implementation cost). For example, if the underlying asset followed five up movements, it will reach a value of $100(1.2229)^5 = 273.53$ (A). The value of the option at this stage (step 5) will be 273.53 – 100 = $173.53 (B). Obviously, if the value in Lattice 1 is < 100, it doesn't make sense to implement the project because the payoff will be smaller than the implementation costs, thus the last three values in the last column of Lattice 2 are 0. Working backwards from Step 5, we derive the values in Step 4, 3, 2, 1, and 0 for Lattice by discounting at the risk-free rate and using the probability of up and down movements. For example, the top value in step 4, 124.66, is derived from [pu + (1-p)d](1 + rfdt) = $[0.4747(173.53) + (1-0.4747)(82.89)][1+0.05(0.2)]^{-1}$. Similarly, the value below it, 50.55 = from $[0.4747(82.89) + (1-0.4747)(22.29)][1+0.05(0.2)]^{-1}$. Continuing step by step backwards until the beginning time period (step 0), we obtain the value of the option to wait: $20.75.

Chapter Four

4.1 Data Analysis

The data analyzed represents the annual transactions costs of a large consumer packaged goods (CPG) manufacturer. Total annual TL shipment costs were approximately $60 million. The analysis focuses on the outbound TL shipments to customers which represent about 45% of total costs. The network consists of 19 origins and more than 1200 destinations. Table 4.1 shows total shipment costs and number of loads for the highest volume outbound origins. These eight locations represent over 80% of all outbound traffic to customers (costs rounded to the nearest $100,000).

Table 4-1: Total Outbound Shipment Costs (Shipment Costs > $1,000,000)

Location	Total Shipment Costs	Total Shipments
A	$5,200,000	10,600
B	$4,500,000	10,900
C	$3,800,000	6,600
D	$2,400,000	4,800
E	$2,100,000	3,900
F	$2,100,000	2,500
G	$1,600,000	2,600
H	$1,500,000	2,000
Total	$23,200,000	43,900

4.1.1 Cost Estimation and Assumptions

The actual transportation costs are taken as given. Total costs are the actual number of shipments (L_a) multiplied by the contract rate (R_c). The analysis focuses on the savings, if any, that can be achieved by using a flexible contract that allows the shipper to adjust its truck requirements based on the demand it faces. The value of the option, as mentioned previously, is the difference between the costs without flexibility contract and the cost with flexibility contract. If total transportation costs with flexibility exceed total actual transportation costs, then the option is worthless. Therefore,

Option Value = Max [(Total costs without flexibility – Total costs with flexibility), 0]

The contract is structured as follows:

1. The shipper commits to the carrier a guaranteed number of loads (L_g) at rate R_g, even if actual number of loads, L_a, is less than L_g. He pays a penalty cost C_p per unused truck if the actual number of used trucks, L_a, is smaller than L_g.
2. The carrier additionally supplies, if requested by the shipper, a variable number (L_v) of trucks at rate R_v up to a certain number. $L_v = e\, L_g$ where e is an expansion factor, $0 \leq e \leq 0.2$. This additional capacity is optional – the shipper can request any number between 0 and e Lg and if he does not use some or all of the variable trucks, he is not required to pay a penalty.
3. If the guaranteed and variable number of trucks are not sufficient to meet the shipper's requirements, he procures the remaining trucks at the spot market (L_{sp} at rate R_{sp}).

Total costs for the carrier =

(1) $L_aR_g + (L_g - L_a)C_p$, if $L_a \leq L_g$

(2) $L_gR_g + L_vR_v$, if $L_g < L_a \leq (1+e)L_g$

(3) $L_gR_g + L_vR_v + L_{sp}R_{sp}$ if $(1+e)L_g < L_a$

We attempt to determine the appropriate level of L_g based on historical number of loads and expected R_d, C_p, R_v, R_{sp}, e.

4.1.2 Locations

We begin by aggregating all outbound lanes and the number of shipments at each location into 53 weekly periods. Although this assumption is unrealistic, because different lanes might have different demand patterns, it allows us to smooth out the fluctuations. We achieve a risk-pooling effect through the aggregation of loads at both weekly period and facility level.

4.1.3 Penalty cost

Based on discussions with TL carriers, in order to make a profit, a carrier needs to have a daily revenue of $600 per truck. If a carrier needs to reposition an unused truck, he would lose at least half a day, hence a rough estimate for the unused truck penalty is $300. (Wicker, 2004).

4.1.4 Fixed (guaranteed) load rate

Carriers are willing to offer a rate lower than the contract rate if they are guaranteed capacity. If they do not have such a guarantee, they build an insurance in the contract rate that covers the risk in case the shipper does not utilize all of the committed trucks. Having guaranteed capacity and a penalty cost would allow the carriers to subtract the insurance premium and offer a lower rate (Wicker, 2004). In this analysis, we assume the guaranteed rate to be 5% lower than contract rate.

4.1.5 Spot rate

Spot rates can be much higher than the contract rate, especially in a tight capacity market. Depending on the type of goods and the urgency for their delivery, it is not impossible to have spot rates that are 50% or even 100% higher (McCarthy, Wicker 2004). For our analysis, we assume that the spot rate is 20% higher than the contract rate.

4.1.6 Contract Rate

Since we have total annual costs and number of truckloads, we can estimate the average weighted contract rate and the average spot rate for our data sample. In order to do this, we need to make an assumption for the percentage of loads that went to the spot market. Evidence from actual data suggests that the percentage of turndowns, i.e. the number of loads that were not accepted by a carrier can exceed 25% (Harding, 2004). For our analysis, we will adopt a conservative number, 5%.

Given our assumptions,

$$Spot\ Cost = (Loads_{spot})Rate_{spot} = 0.05(Total\ Loads)1.2(Contract\ Rate)$$

$$\frac{SpotCost}{TotalCost} = \frac{0.05(\text{Total Loads})1.2(\text{Contract Rate})}{\{(0.05(\text{Total Loads})1.2(\text{Contract Rate}) + (1 - 0.05)\text{Total Loads}(\text{Contract Rate})\}}$$

Canceling out *Total Loads* and *Contract Rate*, $\frac{SpotCost}{TotalCost} = \frac{0.05(1.2)}{0.05(1.2) + (1 - 0.05)}$

Therefore,

$$Spot\ Cost = \frac{0.05(1.2)}{0.05(1.2) + (1 - 0.05)}\,Total\ Cost$$

We can now estimate the average weighted spot and contract rate.

4.1.7 Variable load rate and expansion factor

We assume the variable rate to be equal to the contract rate without a flexible contract. The contract rate itself is set to the average weighted rate at facility level for the whole year, e.g. for

Location A the contract rate is equal to Location A's total shipment costs divided by Location A's total number of shipments. Intuitively under such arrangement the shipper would like to have an expansion factor as large as possible in order to accommodate unexpected surges in demand while enjoying a rate lower than the spot. The carrier, on the other hand, would like to limit variable capacity in order to avoid maintaining or rerouting considerable capacity in case the shipper experiences a surge and also to provide incentives for the shipper to order more guaranteed capacity. Therefore, we limit the expansion factor to 20% of the guaranteed number of trucks.

Note that these estimates are initial numbers. After the basic run we conduct a sensitivity analysis to observe the changes in option value as a result of varying inputs.

4.2 Problem Definition

Our objective is to minimize total costs under the flexible contract and maximize the savings it has compared to the base case when no flexibility exists.

$$\text{Max } [TC_{\text{without flexibility}} - TC_{\text{with flexibility}}] \text{ or}$$

$$\text{Max} \sum_{i=1}^{n} \{La_iRc - (LgRg + Lv_iRv + Lsp_iRsp)\}$$

Subject to:

$Lg_i \geq 0$

$Lv_i = \text{Max}\{\text{Min}[La_i - Lg, (1+e)Lg - Lg], 0\}$

$Lsp_i = \text{Max}[La_i - (1+e)Lg, 0]$

$Rg = 0.95R_c$

$Rv = R_c$

$Rsp = 1.2R_c$

$0 \leq e \leq 0.2$

where:

Indices:

i: weeks in the period, $i = 1, 2, \ldots..n$

Decision Variable:

Lg: number of guaranteed (fixed) loads per week. Constant for every week within each period.

Data:

La_i: actual number of loads observed in week i.
Lv_i: number of loads that were covered by the variable capacity commitment, range is [0, (1+e)Lg]
Lsp_i: number of loads that went to the spot market, range is $[0, La_i - (1+e)\,Lg]$
e: expansion factor, determines the percentage of variable capacity above the guaranteed commitment. $0 \leq e \leq 0.2$.
Rc: average weighted contract rate from historical data.
Rg: rate charged for guaranteed capacity.
Rv: rage charged for variable capacity.
Rsp: rate charged for loads at the spot market.

Based on seasonality, it might be necessary to identify more than one period throughout the year. Ideally, we would like to have a separate period for each week in order to achieve maximum accuracy of our model. This, however, is not a good strategy because we cannot expect the demand pattern for a single week to repeat itself in the future. It is more likely that we will face a probability distribution for the expected number of loads and therefore it is more important to select such value for Lg_i that will provide consistent savings given a large number of possible scenarios. In the analysis of the data, we limit the number of periods to two in order to avoid falling into the trap of using hindsight to make the model fit the existing data perfectly.

4.3 Analysis

We describe the analysis for locations A, B, and F and provide a summary analysis for the remaining facilities. The complete results for all locations are found in the Appendix.

4.3.1 Location A

The average weekly number of shipments in Location A was 200. Based on that, we were able to identify two periods. During the first 28 weeks, the number of shipments was always below 200 per week, during the next 25 weeks in only two cases shipments were below 200 per week (see Appendix for chart of actual loads per week for the entire year). Consequently, we divided the year into "Low" and "High" season and introduced different guaranteed levels for each season. If we had a single level without taking into account seasonality, this would lead to assigning fixed level that was a compromise between the two seasons and result in high penalty costs during the low season and more frequent resorting to spot contracts during the high season. Since our objective is to maximize the difference between total actual shipment cost and total shipment costs with a flexibility contract subject to the different rates, penalty cost, and expansion factor, intuitively we want to make sure that most shipments are handled by either fixed or variable capacity and avoid penalty costs and spot rates. Optimizing total transportation costs with flexibility contract gives us the following results for the recommended number of guaranteed loads:

Table 4-2 Location A: Optimal Number of Guaranteed Loads

	Location A
Penalty Cost	300
Fixed Number of Trucks	
High Season	227
Low Season	121
Expansion Factor	0.2
Total Cost Without Option	$5,216,982
Total Cost With Option	$5,062,045
Savings	$154,936
%Savings	2.97%

The results show a total savings of almost 3% for the flexible contract over the contract without flexibility. This, however, is a result based on historical data for the actual transportation costs and volume outbound from location A. It is unrealistic to expect that in the future we will have exactly the same outcome in terms of weekly demand for trucking capacity. Following the observations made in the problem definition section, we decide to conduct a simulation in order to assess how good the optimal result will be under a different data set. In order to estimate the range of values that savings could take, we keep the already estimated fixed number commitments, then simulate future weekly demand over 1000 trials, and generate a range of values for the savings. We use Crystal Ball, an Excel plug-in, to conduct Monte Carlo simulation for weekly demand levels. Based on historical volumes, we generate a probability distribution for weekly demand for each season (see Appendix: Simulations for distribution graphs).

After assigning weekly demand to the two relevant probability distributions, we run a simulation to calculate the range of values that cost savings could take. In each trial, the transportation costs without a flexibility option are assumed to be equal to the total number of **simulated** loads times the weighted average of cost per load estimated from **historical** data on total loads and costs.

Table 4-3 Location A: Guaranteed Loads = 227 High Season, 121 Low Season, Simulated Savings

Statistics	Savings	%Savings
Trials	1000	1000
Mean	$131,120	2.51%
Median	$135,215	2.59%
Mode	---	---
Standard Deviation	$36,394	0.70%
Coeff. of Variability	0.28	0.28
Range Minimum	$3,037	0.06%
Range Maximum	$220,944	4.24%
Range Width	$217,907	4.18%
Mean Std. Error	$1,150.88	0.02%

As witnessed by simulation results, even though we used an optimal level of guaranteed loads that suited the data from a single year, there are still significant savings achieved provided that the future distribution of loads follows historical data. The average savings amount to $131,120, or 2.51% of total costs without a flexible contract.

Next, we run a series of optimizations to estimate how sensitive savings are to various inputs into total costs.

Table 4-4 Location A, Sensitivity Analysis: Modifying Spot Rate

Spot Rate as % Contract rate	120%	130%	150%	170%	200%
Savings	$154,936	$155,006	$169,282	$191,384	$226,908
% Savings	2.97%	2.97%	3.24%	3.67%	4.35%
% Variable Capacity	20.00%	20.00%	20.00%	20.00%	20.00%
Fixed Low	121	126	134	136	141
Fixed High	227	231	236	237	241
Penalty	$300	$300	$300	$300	$300
Fixed rate as % Contract rate	95%	95%	95%	95%	95%
Spot Loads					
With Option	583	466	320	284	226
Without Option	528	528	528	528	528

Savings increase as the spot rate rises. Although this relationship depends on the actual distribution of loads and may not hold always, in general a flexible contract should do better under higher spot rates because a flexible contract is designed to avoid the occurrences of using the spot market to procure truck capacity. As we can see from the results, under a higher spot rate the flexible contract tries to minimize the number of spot loads by increasing the guaranteed truck commitment which in turn increases the range covered by variable capacity and ultimately leaves relatively few loads that are not covered by either guaranteed or variable commitment.

Table 4-5 Location A, Sensitivity Analysis: Modifying Penalty Cost

Penalty	$300	$350	$400	$450	$500
Savings	$154,936	$145,922	$137,872	$131,382	$125,682
% Savings	2.97%	2.80%	2.64%	2.52%	2.41%
% Variable Capacity	20.00%	20.00%	20.00%	20.00%	20.00%
Fixed Low	121	118	118	116	116
Fixed High	227	226	221	216	216
Spot Rate as % Contract rate	120%	120%	120%	120%	120%
Fixed rate % Contract rate	95%	95%	95%	95%	95%

Overall, savings are not very sensitive to higher penalty rates. Even if the shipper pays $500 penalty (a 67% increase) per unused truck, total savings decrease approximately by only 25% and are still relatively high at almost 2.5% of total costs without flexibility contract. If we think in terms of elasticity, savings drop by only 18.86% following a 67% increase in penalty costs, which translates into an elasticity of - 0.28. Therefore the shipper can offer a much higher compensation to a carrier without compromising his savings significantly. In fact, he might be able to increase savings by offering a higher penalty price in exchange for even further discounts on the guaranteed rate. For example, increasing the penalty cost to $600 per truck in exchange for lowering the guaranteed rate to 0.93 of the contract rate actually increases savings to $185,000 or 3.55% even if we keep the original assignment of guaranteed loads without further optimization.

Table 4-6 Location A, Sensitivity Analysis: Modifying Variable Capacity

% Variable Capacity	20.00%	15.00%	10.00%	5.00%	0.00%
Savings	$154,936	$136,427	$113,588	$86,399	$53,642
% Savings	2.97%	2.62%	2.18%	1.66%	1.03%
Penalty	$300	$300	$300	$300	$300
Fixed Low	121	122	126	126	128
Fixed High	227	227	231	234	240
Spot Rate as % Contract rate	120%	120%	120%	120%	120%
Fixed rate as % Contract rate	95%	95%	95%	95%	95%

The different levels of expansion have a clearly negative effect on the size of savings. As a result, the shipper is forced to increase his weekly fixed capacity commitment in order to compensate the decreased variable capacity As discussed before, the carriers will prefer to keep the size of variable (optional) capacity as low as possible in order to induce shippers to increase their commitment to guaranteed capacity thus decreasing the likelihood that carriers will have to provide a lot more capacity to cover the shippers' needs in case of a surge.

4.3.2 Location B

The loads distribution at location B followed a similar pattern as at location A. The average weekly number of shipments was 206. During the first 28 weeks, the average number of shipments was 162 per week, during the next 25 weeks it rose to 255 per week. We found the optimal number of guaranteed trucks to be 131 for the low season and 216 for the high season. This resulted in savings of $129,610 or 2.91% of the total costs without a flexible contract.

Table 4-7 Location B: Optimal Number of Guaranteed Loads

	Location B
Penalty Cost	300
Fixed Number of Trucks	
High Season	216
Low Season	131
Expansion Factor	0.2
Total Cost Without Option	$4,450,014
Total Cost With Option	$4,320,404
Savings	$129,610
%Savings	2.91%

Using the above number for guaranteed loads and historical data for the weekly distribution of loads, we simulated future weekly loads and generated a distribution of the potential savings that could be realized (see Appendix for distribution of loads).

Table 4-8 Location B: Guaranteed Loads = 216 High Season, 131 Low Season, Simulated Savings

Statistics	Savings	% Savings
Trials	1000	1000
Mean	$119,418	2.68%
Median	$120,540	2.71%
Mode	$109,434	2.46%
Standard Deviation	$20,305	0.46%
Coeff. of Variability	0.17	0.17
Range Minimum	$49,114	1.10%
Range Maximum	$184,969	4.16%
Range Width	$135,855	3.05%
Mean Std. Error	$642.10	0.01%

On average, the contract with flexibility results in savings of $119,418 or 2.68% of total costs without flexibility.

Similar to our analysis for Location A, we run a series of optimizations to estimate the sensitivity of savings to rising spot rates and penalty costs and decreasing expansion factor.

Table 4-9 Location B, Sensitivity Analysis: Modifying Spot Rate

Spot Rate as % Contract rate	120%	130%	150%	170%	200%
Savings	$129,610	$121,131	$114,280	$114,964	$120,914
% Savings	2.91%	2.72%	2.57%	2.58%	2.72%
% Variable Capacity	20.00%	20.00%	20.00%	20.00%	20.00%
Fixed Low	131	136	141	146	151
Fixed High	216	218	221	226	231
Penalty	$300	$300	$300	$300	$300
Fixed rate as % Contract rate	95%	95%	95%	95%	95%
Spot Loads					
With Option	771	682	596	506	436
Without Option	544	544	544	544	544

Unlike Location A, Location B fails to increase its savings under higher spot rates. We observe that the number of loads going to spot under the flexible contract is relatively high and exceeds the assumed 5% spot loads. However, as the spot rate increases, the savings gradually begin to increase and even though they fail to match the initial value, there is clear rising trend. In fact, if the spot rate were to increase to 220%, the savings will be higher than they were at 120%.

Table 4-10 Location B, Sensitivity Analysis: Modifying Penalty Cost

Penalty	$300	$350	$400	$450	$500
Savings	$129,610	$124,079	$119,529	$114,979	$110,429
% Savings	2.91%	2.79%	2.69%	2.58%	2.48%
% Variable Capacity	20.00%	20.00%	20.00%	20.00%	20.00%
Fixed Low	131	131	131	131	131
Fixed High	216	211	211	211	211
Spot Rate as % Contract rate	120%	120%	120%	120%	120%
Fixed rate % Contract rate	95%	95%	95%	95%	95%

Again we witness that savings seems to be little affected by higher penalty rates. The increase of the penalty rate by 67% reduces savings by 14.7%, elasticity = - 0.22.

Table 4-11 Location B, Sensitivity Analysis: Modifying Variable Capacity

% Variable Capacity	20.00%	15.00%	10.00%	5.00%	0.00%
Savings	$129,610	$112,858	$91,833	$69,637	$41,125
% Savings	2.91%	2.54%	2.06%	1.56%	0.92%
Penalty	$300	$300	$300	$300	$300
Fixed Low	131	135	138	141	145
Fixed High	216	216	217	217	220
Spot Rate as % Contract rate	120%	120%	120%	120%	120%
Fixed rate as % Contract rate	95%	95%	95%	95%	95%

Confirming our intuition, if the ability to request additional capacity is restricted, savings are adversely affected and the commitment for guaranteed capacity increases.

4.3.3 Location F

Although weekly number of loads for facility was quite spiky, there were no clearly identifiable seasons. The average number of loads per week throughout the year was 48, with the lowest being 27, highest -78. We found the optimal number of guaranteed loads to be 42 which resulted in annual savings of $66,512, 3.18% of total costs without a flexible contract.

Table 4-12 Location F: Optimal Number of Guaranteed Loads

	Location F
Penalty Cost	300
Fixed Number of Trucks	
All Year	42
Expansion Factor	0.2
Total Cost Without Option	$2,090,036
Total Cost With Option	$2,023,524
Savings	$66,512
%Savings	3.18%

Table 4-13 Location F: Guaranteed Loads = 42 All Year, Simulated Savings

Statistics	Savings	%Savings
Trials	1000	1000
Mean	$66,449	3.18%
Median	$67,612	3.23%
Mode	---	---
Standard Deviation	$10,444	0.50%
Coeff. of Variability	0.16	0.16
Range Minimum	$30,960	1.48%
Range Maximum	$94,283	4.51%
Range Width	$63,322	3.03%
Mean Std. Error	$330.27	0.02%

Using the optimal estimate for guaranteed loads, we simulate savings and find that on average they amount to $66,449, or 3.18% of total cost without flexibility contract. Next, we estimate the sensitivity of savings to varying spot rates, penalty rates and variable capacity.

Table 4-14 Location F, Sensitivity Analysis: Modifying Spot Rate

Spot Rate as % Contract rate	120%	130%	150%	170%	200%
Savings	$66,512	$67,254	$76,756	$88,184	$105,213
% Savings	3.18%	3.22%	3.67%	4.22%	5.03%
% Variable Capacity	20.00%	20.00%	20.00%	20.00%	20.00%
Fixed All Year	42	42	136	136	141
Penalty	$300	$300	$300	$300	$300
Fixed rate as % Contract rate	95%	95%	95%	95%	95%
Spot Loads					
With Option	111	111	47	47	39
Without Option	125	125	125	125	125

Location F exhibits a clear pattern of increasing savings as spot rate increases. The number of spot loads under the flexibility contract decreases rapidly and leads to considerable savings compared to the base case.

Table 4-15 Location F, Sensitivity Analysis: Modifying Penalty Cost

Penalty	$300	$350	$400	$450	$500
Savings	$66,512	$63,013	$59,713	$56,413	$53,113
% Savings	3.18%	3.01%	2.86%	2.70%	2.54%
% Variable Capacity	20.00%	20.00%	20.00%	20.00%	20.00%
Fixed All Year	42	41	41	41	41
Spot Rate as % Contract rate	120%	120%	120%	120%	120%
Fixed rate % Contract rate	95%	95%	95%	95%	95%

An increase of penalty cost by 67% decreases savings only by 20%, elasticity = 0.30. The increase in penalty cost has limited effect on savings as it was the case at the previous two facilities.

Table 4-16 Location F, Sensitivity Analysis: Modifying Variable Capacity

% Variable Capacity	20.00%	15.00%	10.00%	5.00%	0.00%
Savings	$66,512	$59,916	$52,002	$43,252	$28,849
% Savings	3.18%	2.87%	2.49%	2.07%	1.38%
Penalty	$300	$300	$300	$300	$300
Fixed All Year	42	42	43	44	46
Spot Rate as % Contract rate	120%	120%	120%	120%	120%
Fixed rate as % Contract rate	95%	95%	95%	95%	95%

Although percentage savings are slightly higher compared to the previous two locations, limiting variable capacity ultimately erodes savings because with limited ability to expand the shipper is forced to increase his guaranteed capacity commitment in order to avoid higher spot rates but in this way he increases the frequency of incurring penalty costs. The tradeoff of having a low level of guaranteed capacity to avoid penalty costs versus high level of guaranteed capacity to avoid spot contracts depends heavily on the expansion factor and eliminating it completely decreases the attractiveness of a flexible contract . In reality, the shipper is likely to have additional options before going to the spot market. Having more than one carrier can alleviate the problem

4.3.4 Summary Results

Table 4-17 All Locations: Simulated Savings

Location	A	B	C	D	E	F	G	H
Trials	1,000	1,000	1,000	1,000	1,000	1,000	1,000	1,000
Mean Savings	$131,120	$119,418	$19,942	$5,150	$56,335	$66,449	-$29,634	$49,814
Range Minimum	$3,037	$49,114	-$77,459	-$63,766	$22,300	$30,960	-$81,634	$30,236
Range Maximum	$220,944	$184,969	$88,062	$49,147	$83,325	$94,283	$17,635	$65,716
Mean % Savings	2.51%	2.68%	0.52%	0.21%	2.66%	3.18%	-1.87%	3.26%
Range Minimum	0.06%	1.10%	-2.04%	-2.61%	1.05%	1.48%	-5.16%	1.98%
Range Maximum	4.24%	4.16%	2.31%	2.01%	3.93%	4.51%	1.12%	4.30%

Table 4.17 shows that five of the eight locations have significant savings (> 2.5%), two locations provide almost no savings, and one location is actually doing worse with the flexible contract. The underlying factor that drives savings is the distribution of loads throughout the year.

Table 4-18 All Locations: Distributions of Loads Statistics

Location	A	B	C	D	E	F	G	H
Weeks	53	53	53	53	53	53	53	53
All Year								
Mean	199.26	205.40	125.28	91.23	69	47.36	48.64	37.92
Standard Deviation	66.42	62.28	66.52	38.72	11.58	9.34	30.18	9.66
Range	269	292	219	155	43	51	134	37
Minimum	86	106	36	41	45	27	17	21
Maximum	355	398	255	196	88	78	151	58
Sum	10561	10886	6640	4835	1587	2510	2578	2010
Coefficient of Variation	0.33	0.30	0.53	0.42	0.17	0.20	0.62	0.25

In general, the locations that showed the greatest savings have a coefficient of variation for their distribution of loads at or below 0.3. The locations with little or negative savings have higher coefficients of variation. Figure 4.1 shows a negative relationship between percentage of savings and the coefficient of variability for the actual truckloads. This implies that in order to be effective, a flexible option contract needs to be applied to situations where the variability of truckload capacity needed does not exceed a certain threshold beyond which the gains from avoiding spot contracts are offset by frequent penalty costs. This seems counterintuitive because normally we would expect an option to be more valuable if the underlying asset (the number of weekly loads in our case) is more volatile. The reason for this discrepancy lies in the assumptions of our model. When estimating contract and spot rates, we assumed that only 5% of all loads in the historical data were covered by spot contracts. If we look at the distribution of actual weekly loads in the appendix, we observe that some locations are much more volatile than others. In particular, Location G has spikes that are often 2 or 3 times higher than the volume in the preceding week. It is likely that there will be more instances of spot loads at such a high volatility location, yet we assumed a universal ratio of $\frac{SpotLoads}{TotalLoads}$ across all facilities. Based on our assumption, Location G had only 130 loads that went to the spot market. Since the level of guaranteed loads is set as a compromise between having a high guaranteed commitment and incur penalty cost and having low guaranteed commitment and incur high spot rate, under the option contract we have more than 500 loads going to the spot market. We do not know the exact contract that the shipper had with its carriers but we would expect a flexibility contract by design to tend to avoid the higher spot rates and result in fewer loads going to the spot market, therefore it is possible that we have underestimated the percentage of spot loads for the higher volatility locations. For example, if we assume the upper limit for the percentage of turndowns, a flexible contract at Location G provides savings of almost 2.5% compared to the base case.

Figure 4-1 %Savings vs. Coefficient Of Variability For Weekly Loads

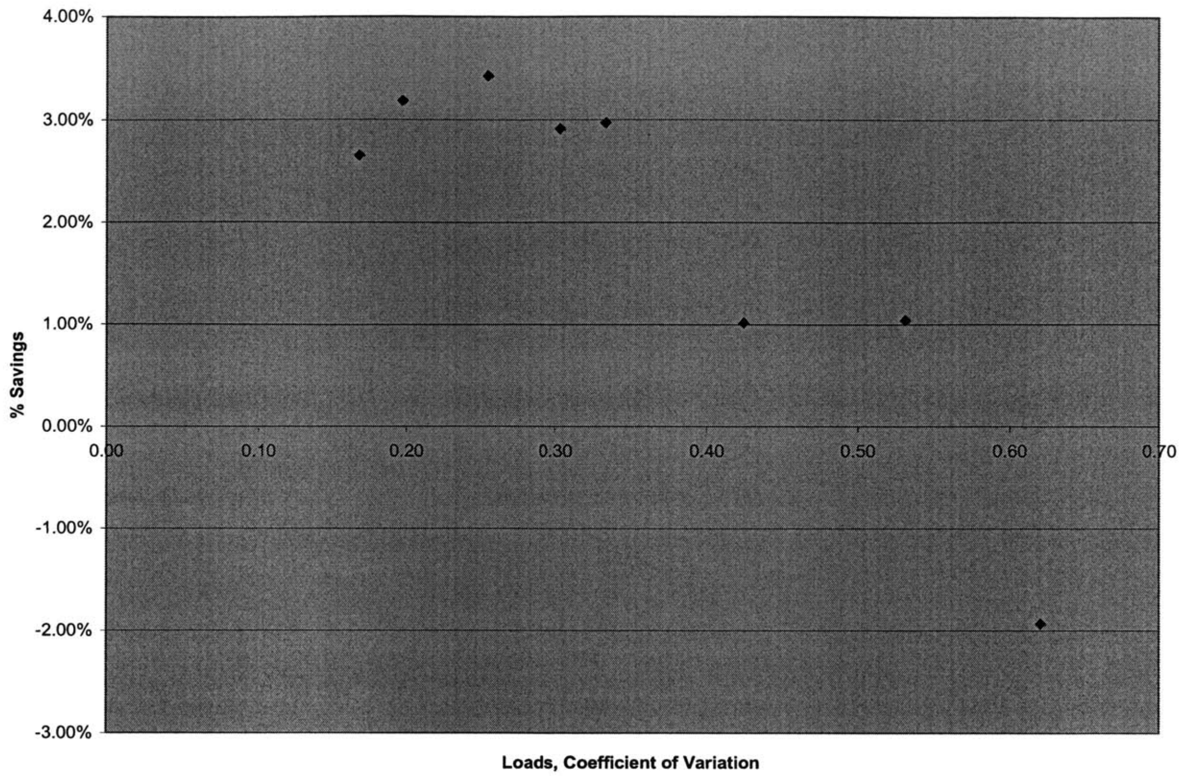

Looking at Table 4-17, we note that the five locations with the highest savings had savings that were consistently positive. This result is encouraging since it shows that if a shipper expects future demand to follow roughly the historical distribution of loads, he can still expect to achieve savings. If we study the figures in the Appendix that show the generated distribution of demand for truckload capacity and compare them to Table 4-18, we can see that in most cases the generated distributions cover a much wider range of outcomes than the historical ranges for the data. This further increases the confidence that the savings are not just the result of fortuitous circumstances.

Chapter Five

5.1 Conclusion

The findings in Chapter Four demonstrate that, given the assumptions of the model discussed, a flexible contract for transportation procurement can provide value to the shipper. The main challenge is to design a contract whose parameters can guarantee higher savings to the shipper and at the same time provide the carrier with sufficient incentives to enter such an agreement.

One way of achieving this would be to offer to pay to the carrier a high penalty rate for unused guaranteed capacity in exchange for low guaranteed rate. A carrier might be interested in such an arrangement provided he is adequately compensated. This will minimize the uncertainty he faces because he will have commitment from the shipper who will pay either the guaranteed rate or a sizeable penalty that will cover the re-routing of unused capacity. In addition, the carrier will be able to develop dedicated fleets that serve the same destinations for a given shipper, and in this way allow its drivers to follow the same routes and better plan their schedules. This will increase driver satisfaction, lower driver turnover and partly lower the capacity problem which is due not only because of lack of investment in new equipment over the past few years but also because the supply of skilled drivers is very limited.

A major advantage of having a large portion of the shipper's transportation needs covered by guaranteed truck commitment will be the increased level of service offered by the carriers. As was mentioned before, the percentage of loads turned down by carriers could be significant. Due to uncertainty in demand, carriers accept that shippers will not always fulfill their commitment and use the trucks they have requested. Similarly, shippers do not expect carriers to respond to all of their requests for trucks but only to a high percentage of them, 70-80% (Caplice and Sheffi, 2003). If, however, the shipper knows that he will pay a penalty in case he does not use a certain base level of capacity, he will expect the carrier to respond with a much higher level of service approaching 100% or have a reciprocating penalty costs that the carrier will have to pay if he fails to provide the needed capacity.

Having a guaranteed higher service level may be attractive to those companies that currently employ private fleets in order to ensure better service and timely delivery of their products, without necessarily having costs as their primary concern. If using a flexible contract can offer them the same service level but at a lower cost, this approach is definitely an area they should consider and explore further.

5.2 Suggestions for further research

The current model suffers from several limitations. It aggregates lanes at facility level in order to smooth out fluctuations along individual lanes. In order to test its results more conclusively, it needs to be extended to a lane level because carriers might be willing to know the shipper's network structure in more detail in order to plan their subsequent movements accordingly.

In addition, it will be useful to conduct a more complex simulation that varies several inputs simultaneously and study, for example, the compound effect of different spot rate, penalty cost and expansion factor. A further calibration of the current model might be to have a better estimate for the historical spot and contract rates that account for the volatilities of number of loads at different locations.

The analysis was based on data from a CPG manufacturer from a single year. A possible extension of the model would be to test its findings with additional data from the same company. Ideally, we should have several consecutive years and test whether we would have been able to achieve any savings by estimating the level of guaranteed commitment from historical data and applying it forward. Alternatively, we should test the approach to companies from other industries in order to collect more data and get a better idea of the conditions under which such an approach is justifiable and when it does not make sense.

It is not clear whether the carrier will be willing to accept the condition to provide a variable capacity up to a certain level when there is no penalty paid by the shipper for this part of the deal. The carrier might request from the shipper to pay upfront an extra fee per truck from the variable capacity as compensation for having this capacity ready in case the shipper needs it. Having such a fee will make the contract less attractive for the shipper. In addition, setting the fee will depend to a large extent on the competition between carriers who are willing to enter such a flexible contract.

Bibliography

Bear Stearns, (2004). The Supply-Chain Indicator, Fourth-Quarter 2003 Shippers Survey Results.

Brach, M. A. (2003). *Real Options in Practice.* Hoboken, NJ, John Wiley & Sons, Inc.

Caplice, C and Sheffi, Y. "Chapter 21: Combinatorial Auctions for Truckload Transportation," in *Combinatorial Auctions*, MIT Press, Eds., Peter Cramton, Yoav Shoham, and Richard Steinberg, forthcoming.

Copeland, T., and V. Antikarov (2001). *Real Options - A Practioner's Guide*, Texetere.

Cox, J., Ross, S. and Rubinstein M. (1979) "Options Pricing: A Simplified Approach", *Journal of Financial Economics*, No. 7, pp. 141-183.

Cox, J. C. (1985). *Options Markets.* Englewood Cliffs, NJ, Prentice-Hall, Inc.

Harding, M. (2004). Innovations in Transportation Conference. MIT, Cambridge, MA.

Howell, S. et al. (2001). *Real options: evaluating corporate investment opportunities in a dynamic world.* London, Financial Times Prentice Hall.

Hull, J. (1991). *Introduction to Futures and Options Markets.* Englewood Cliffs, NJ, Prentice Hall.

McCarthy, K. Manager Logistics Services, C. H. Robinson Worldwide. Personal Interview. April 19, 2004.

Mittal, G. (2004) *Real Options Approach to Capacity Planning Under Uncertainty*, Thesis (S.M.) Massachusetts Institute of Technology, Engineering Systems Division, Technology and Policy Program, Cambridge, MA.

Moore, E. William, Janice M. Warmke, and Lonny R. Gorban. 1991. "The Indispensable Role of Management Science in Centralizing Freight Operations at Reynolds Metals Company". *Interfaces.* v21 n1: 107-129.

Mun, J. (2002). *Real Options Analysis.* Hoboken, NJ, John Wiley & Sons.

Plummer, C. (2003). *Bidder Response to Combinatorial Auctions in Truckload Procurement*, Master of Engineering in Logistics Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Stalter, M. Senior Manager Transportation Procurement, Pepsico Beverages and Foods. Personal Interview. April 23, 2004

Standard & Poor's Industry Surveys. (2003) Volume 171, No. 52, Section 1

Whicker, G. SVP Engineering Services, J. B. Hunt Transport, Inc. Personal Interview. April 23, 2004.

Appendix

Simulations

Location A

Forecast: Savings

Summary:
Display Range is from $38,560 to $220,944
Entire Range is from $3,037 to $220,944
After 1,000 Trials, the Std. Error of the Mean is $1,151

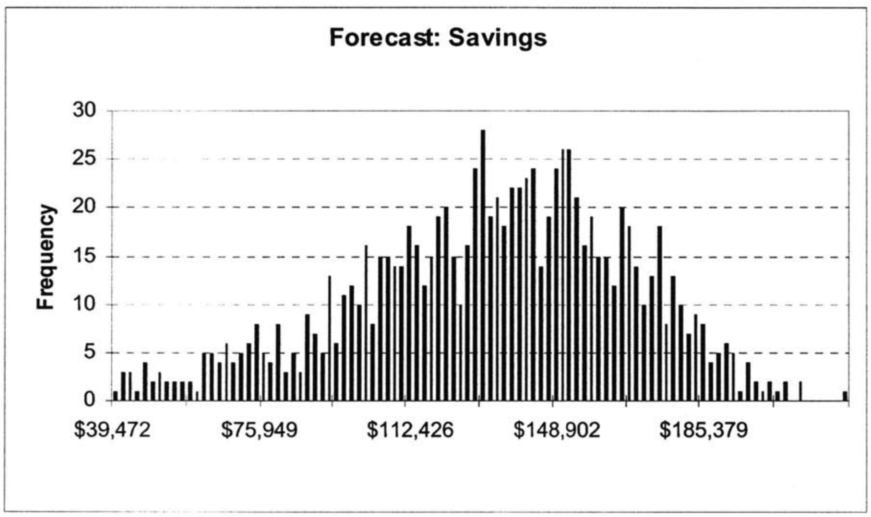

Forecast: %Savings

Summary:
Display Range is from 0.74% to 4.24%
Entire Range is from 0.06% to 4.24%
After 1,000 Trials, the Std. Error of the Mean is 0.02%

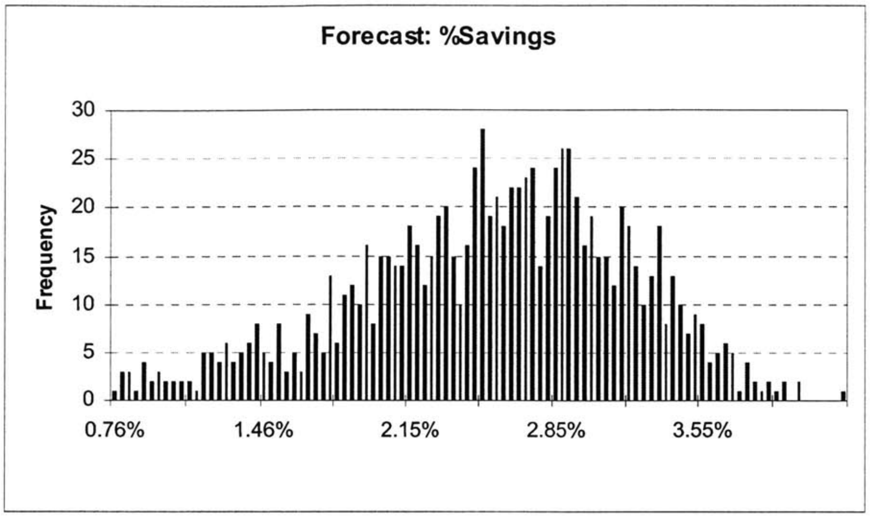

Assumptions: Distribution of Loads

Low Season

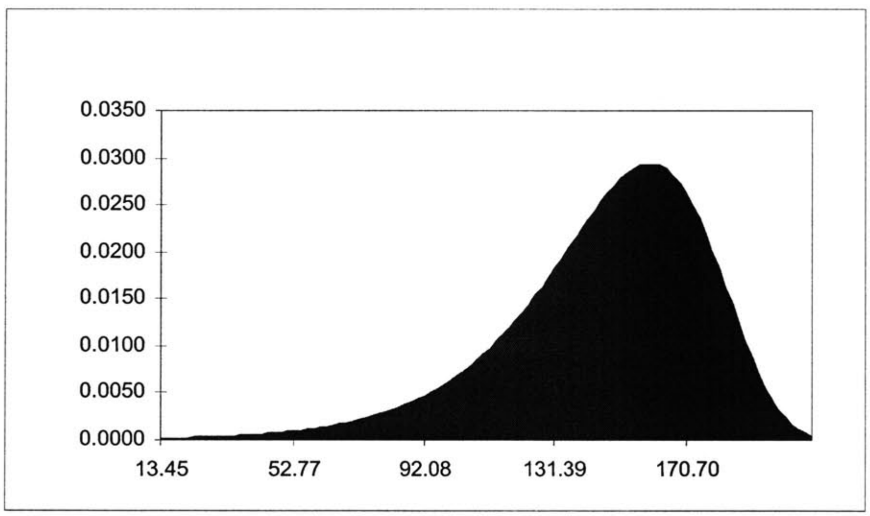

High Season

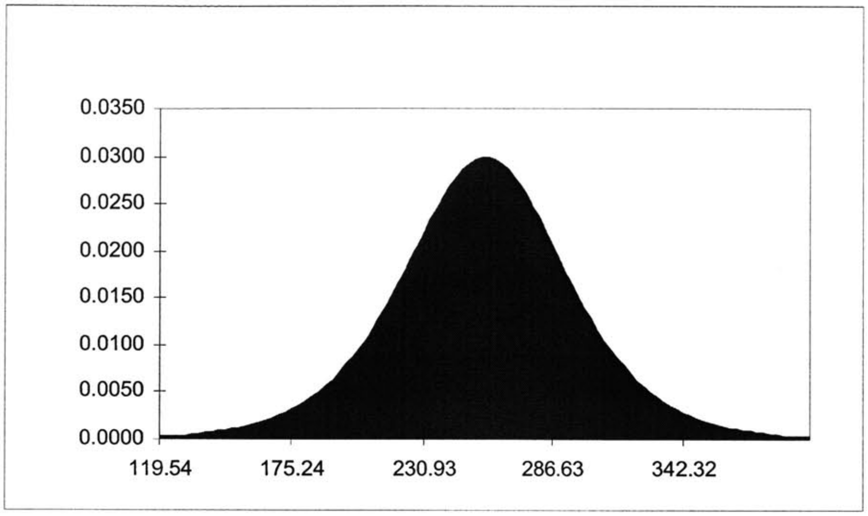

Location B

Forecast: Savings

Summary:

Display Range is from $64,923 to $173,837
Entire Range is from $49,114 to $184,969
After 1,000 Trials, the Std. Error of the Mean is $642

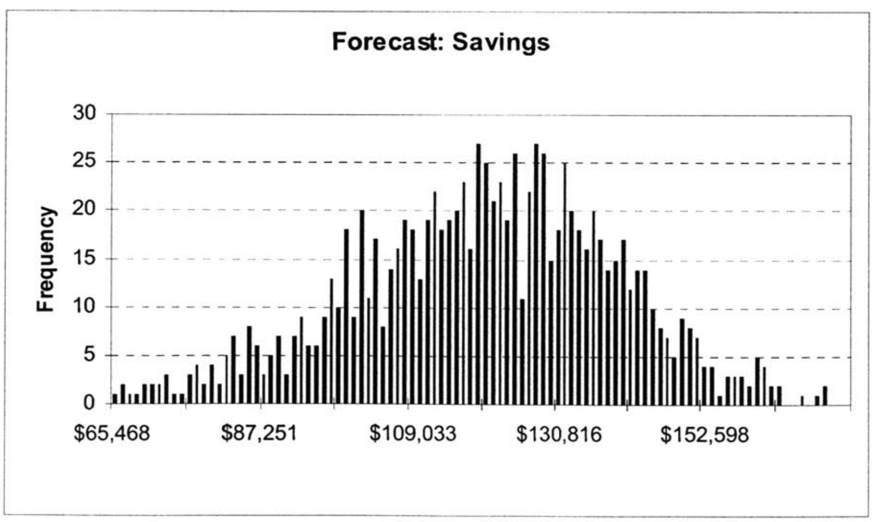

Forecast: % Savings

Summary:
Display Range is from 1.46% to 3.91%
Entire Range is from 1.10% to 4.16%
After 1,000 Trials, the Std. Error of the Mean is 0.01%

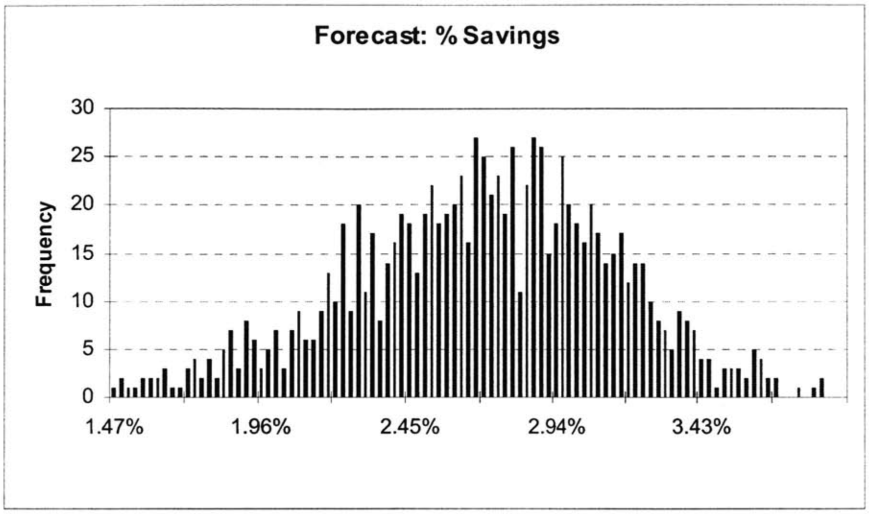

Assumptions: Distribution of Loads

Low Season

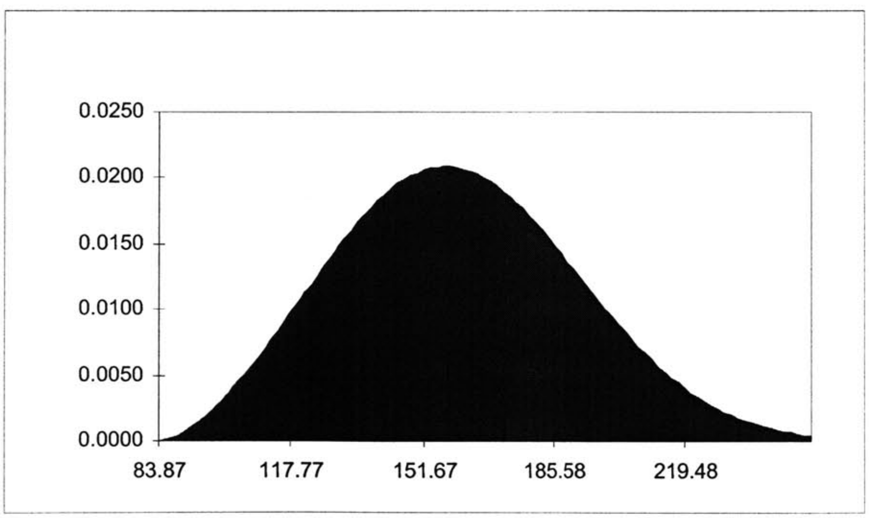

High Season

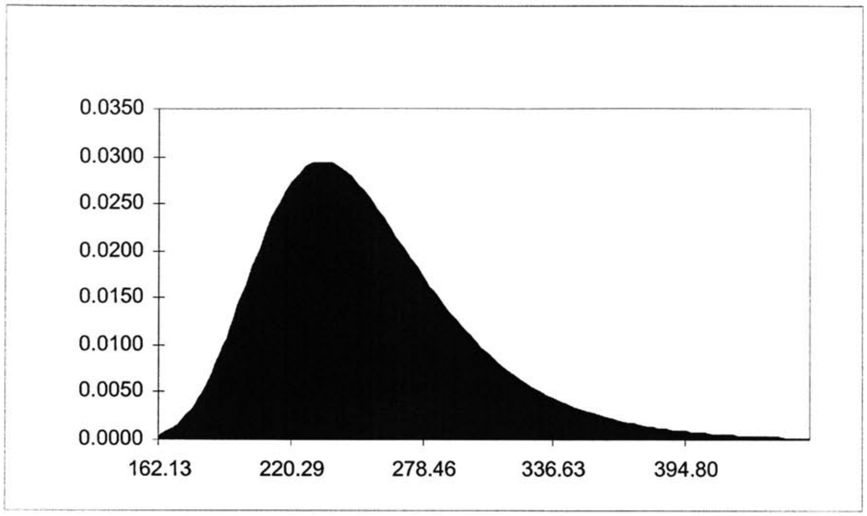

Location C

Forecast: Savings

Summary:

Display Range is from ($56,133) to $88,062
Entire Range is from ($77,459) to $88,062
After 1,000 Trials, the Std. Error of the Mean is $854

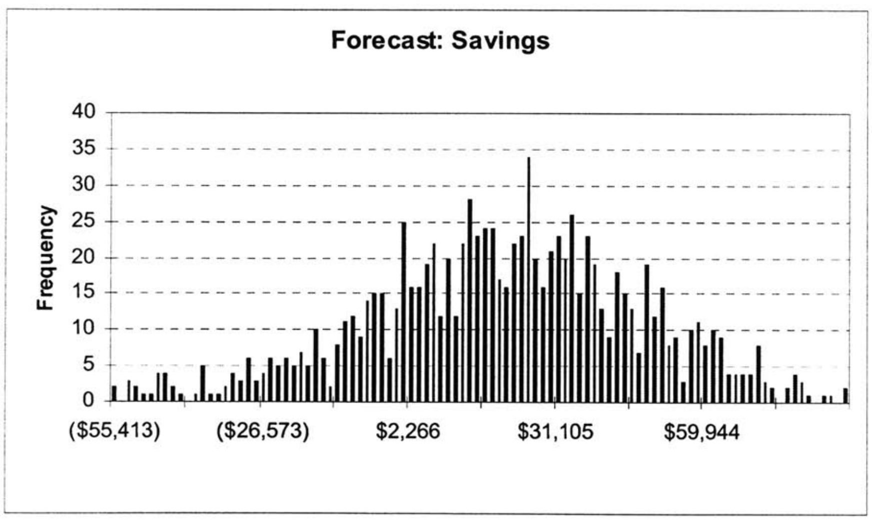

Forecast: %Savings

Summary:

Display Range is from -1.47% to 2.31%
Entire Range is from -2.04% to 2.31%
After 1,000 Trials, the Std. Error of the Mean is 0.02%

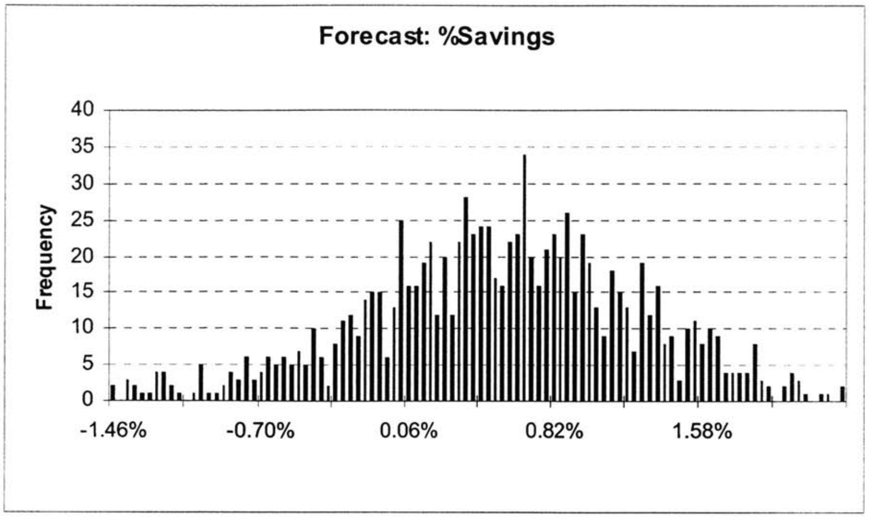

Assumptions: Distribution of Loads

Low Season

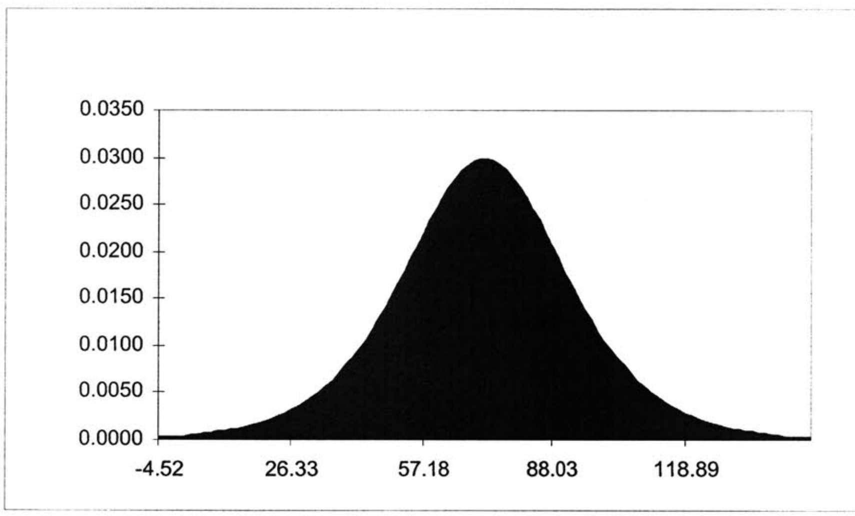

High Season

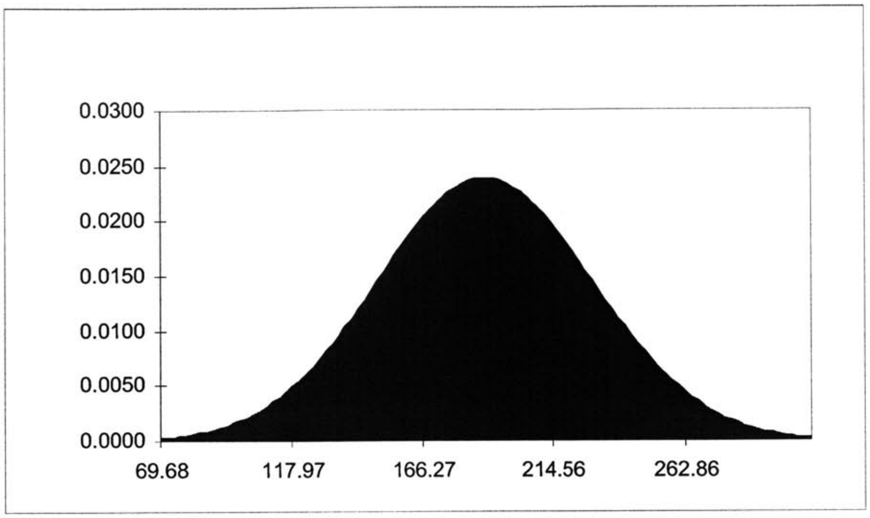

Location D

Forecast: Savings

Summary:
Display Range is from ($40,509) to $47,417
Entire Range is from ($63,766) to $49,147
After 1,000 Trials, the Std. Error of the Mean is $525

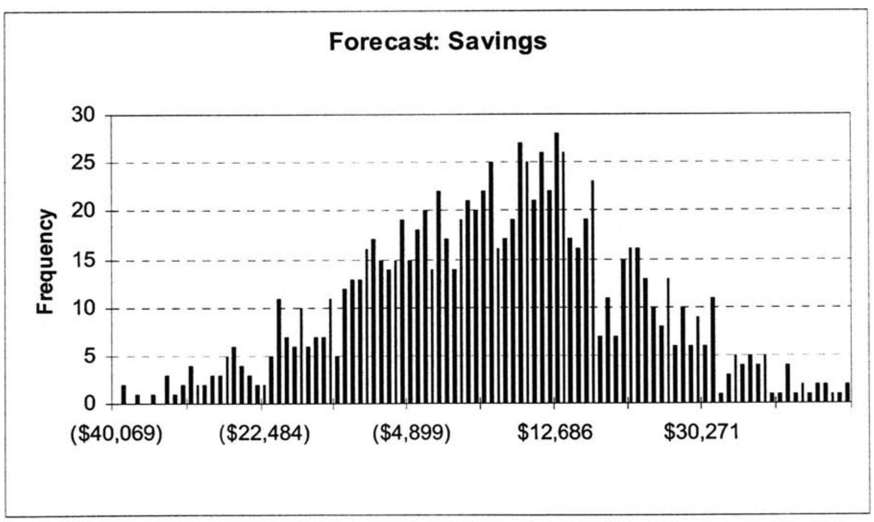

Forecast: %Savings

Summary:
Display Range is from -1.66% to 1.94%
Entire Range is from -2.61% to 2.01%
After 1,000 Trials, the Std. Error of the Mean is 0.02%

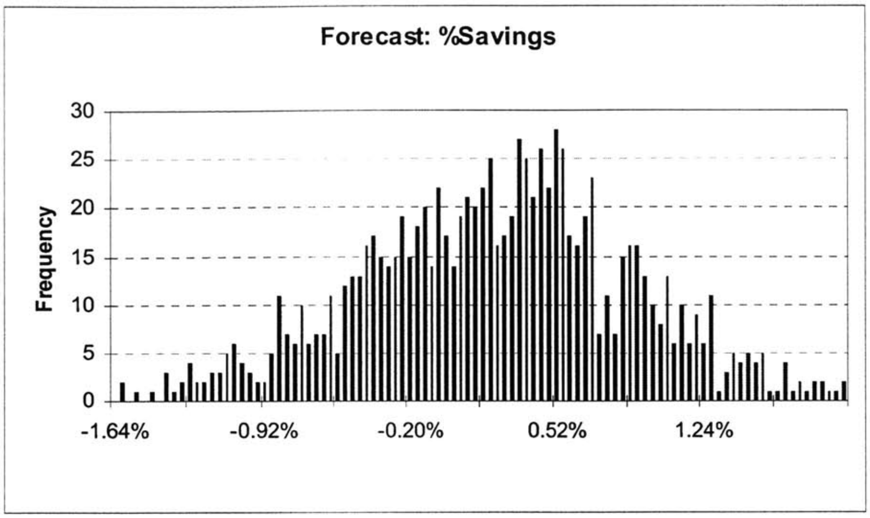

Assumptions: Distribution of Loads

Low Season

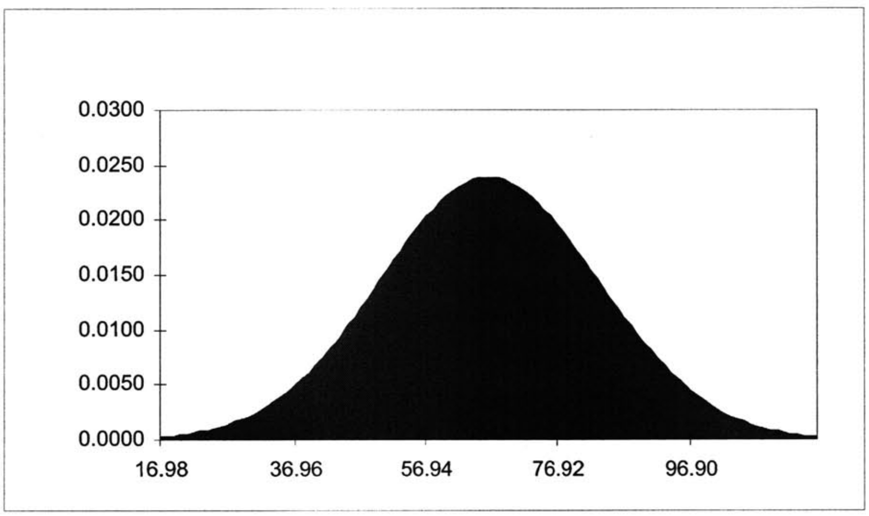

High Season

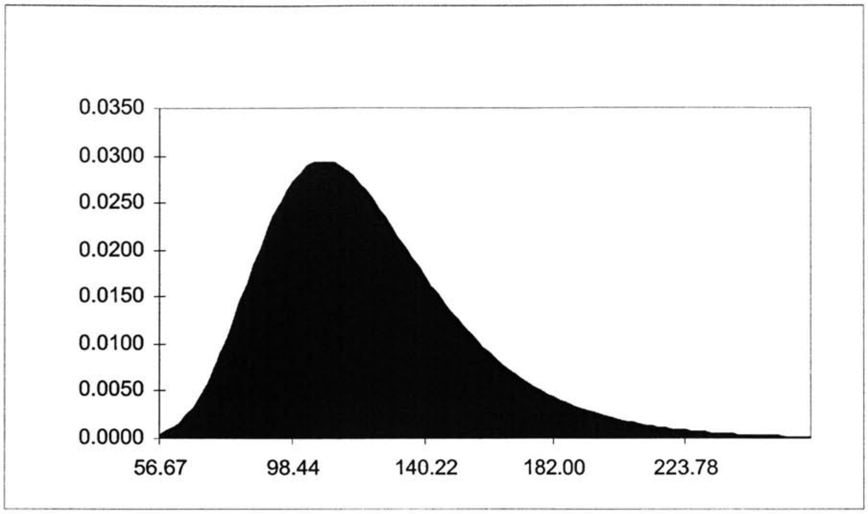

Location E

Summary:
Display Range is from \$31,085 to \$82,544
Entire Range is from \$22,300 to \$83,325
After 1,000 Trials, the Std. Error of the Mean is \$317

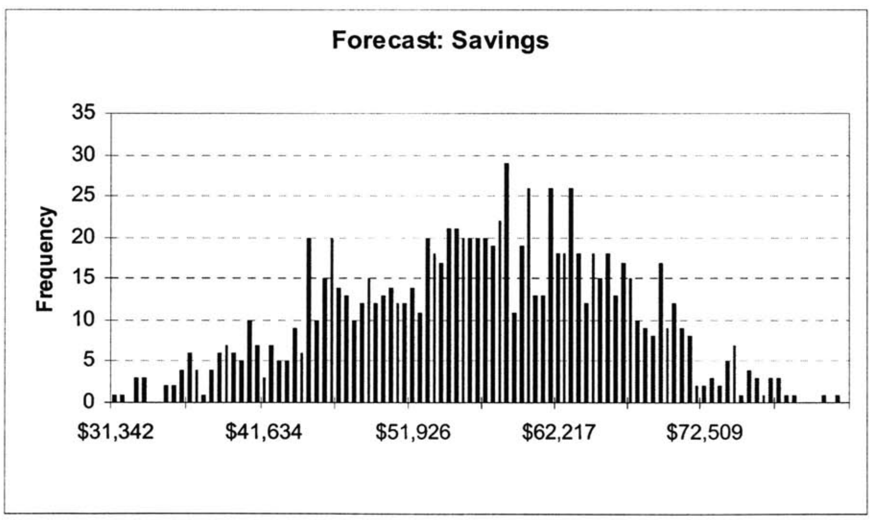

Forecast: %Savings

Summary:
Display Range is from 1.47% to 3.89%
Entire Range is from 1.05% to 3.93%
After 1,000 Trials, the Std. Error of the Mean is 0.01%

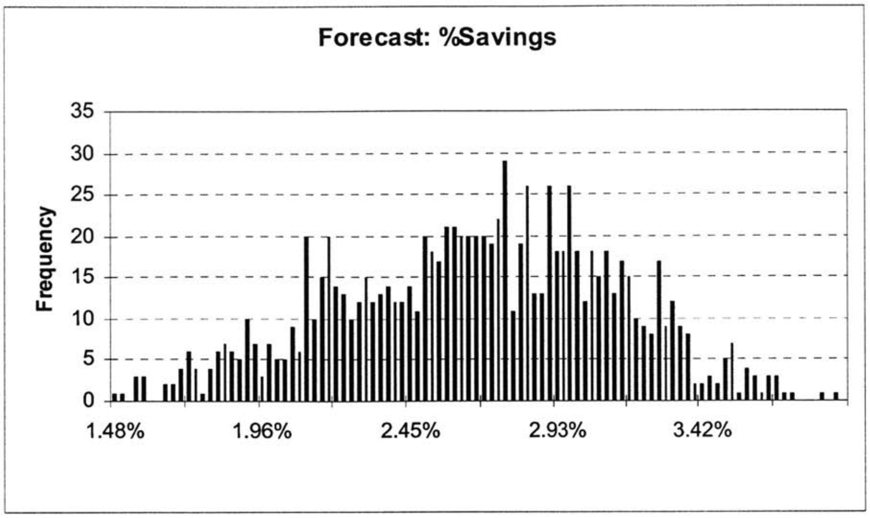

Assumptions: Distribution of Loads

All Year

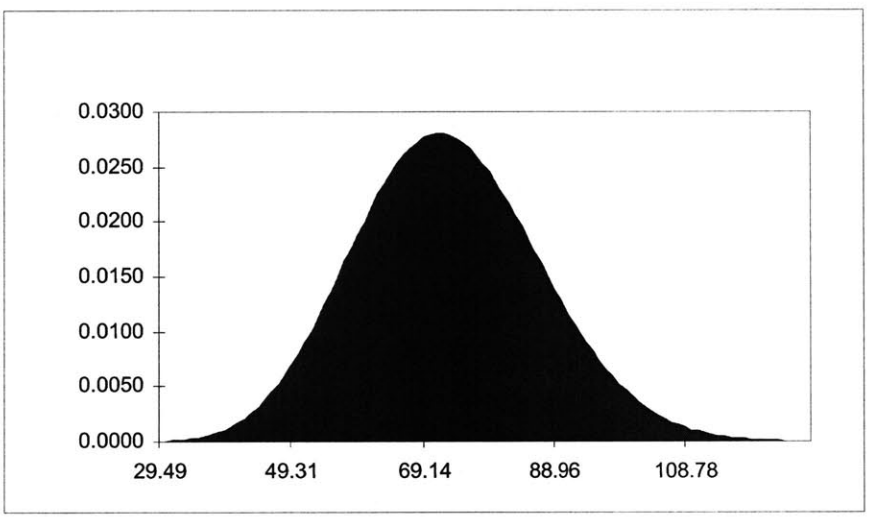

Location F

Forecast: Savings

Summary:

Display Range is from \$39,544 to \$93,120 \$
Entire Range is from \$30,960 to \$94,283 \$
After 1,000 Trials, the Std. Error of the Mean is \$330

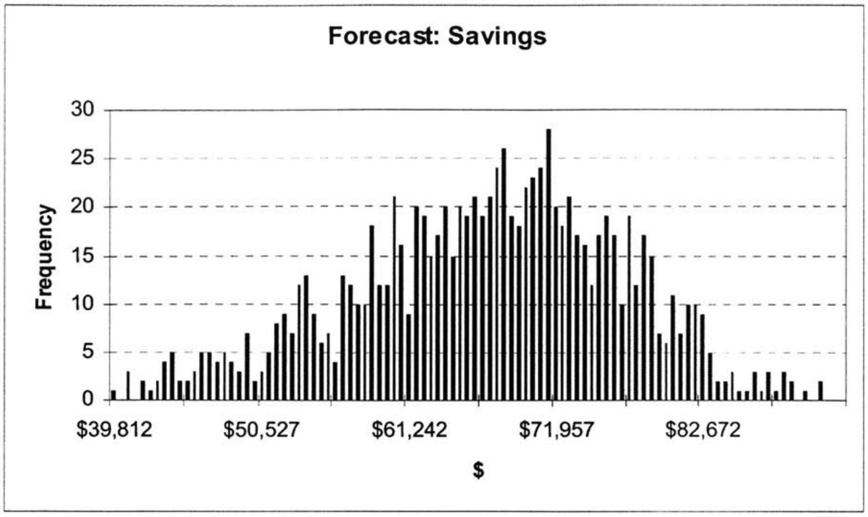

Forecast: %Savings

Summary:

Display Range is from 1.89% to 4.46%
Entire Range is from 1.48% to 4.51%
After 1,000 Trials, the Std. Error of the Mean is 0.02%

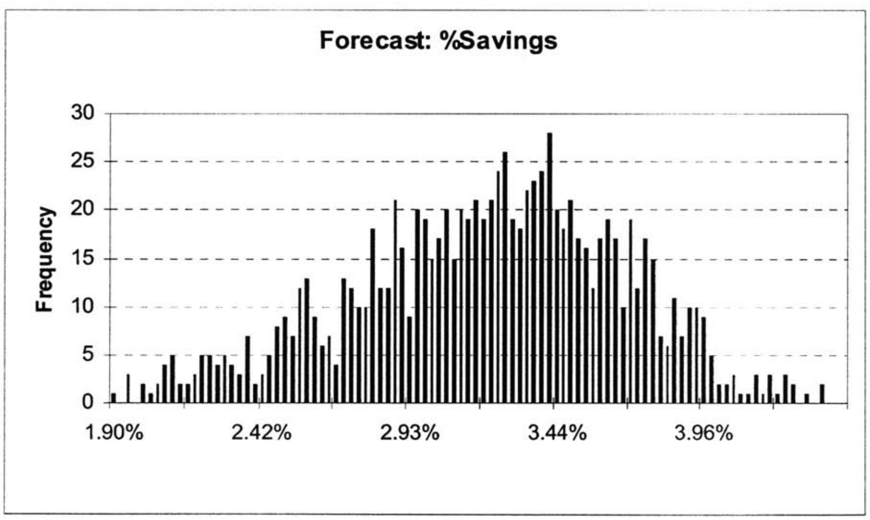

Assumptions: Distribution of Loads

All Year

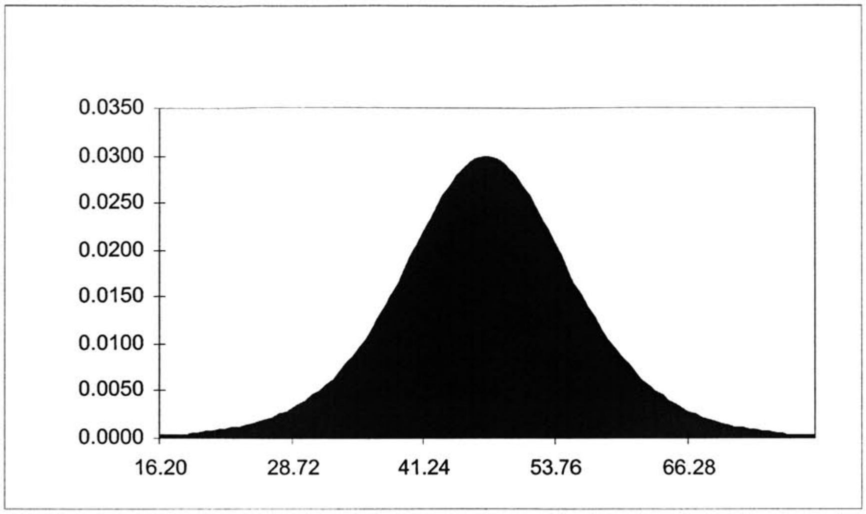

Location G

Summary:
Display Range is from ($71,340) to $11,917
Entire Range is from ($81,634) to $17,635
After 1,000 Trials, the Std. Error of the Mean is $500

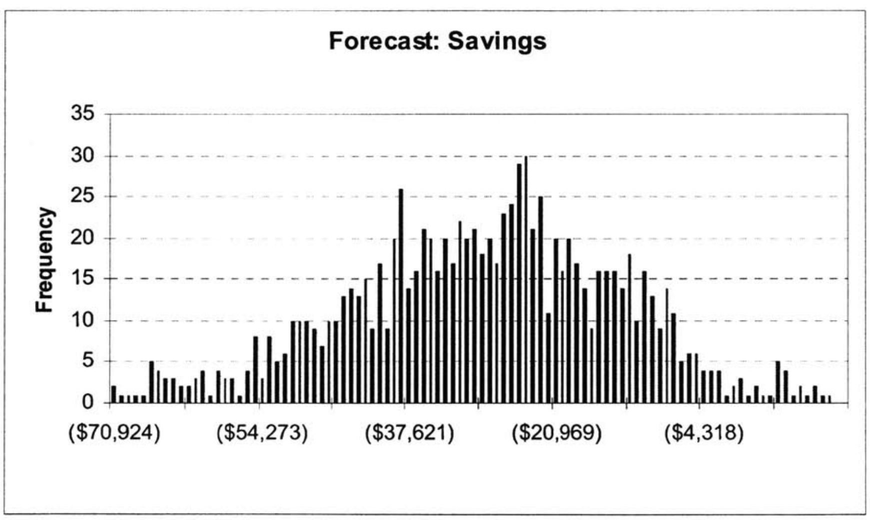

Summary:
Display Range is from -4.51% to 0.75%
Entire Range is from -5.16% to 1.12%
After 1,000 Trials, the Std. Error of the Mean is 0.03%

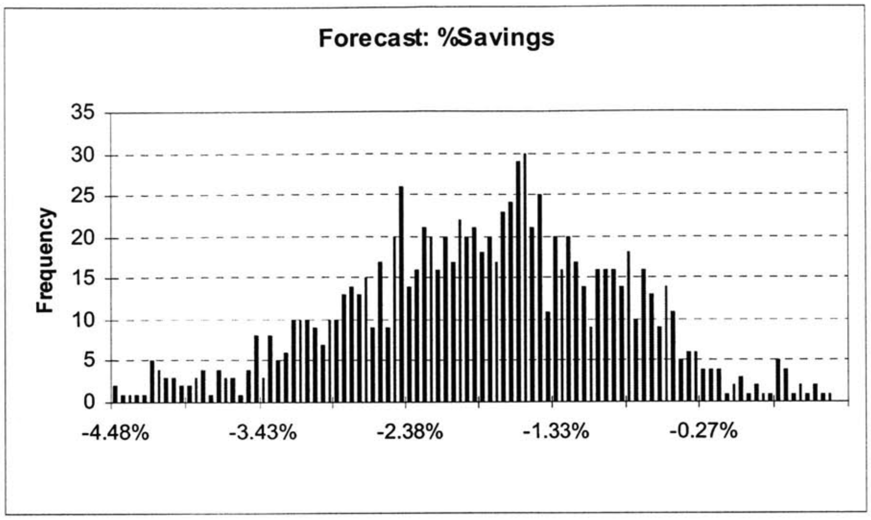

Assumptions: Distribution of Loads

Low Season

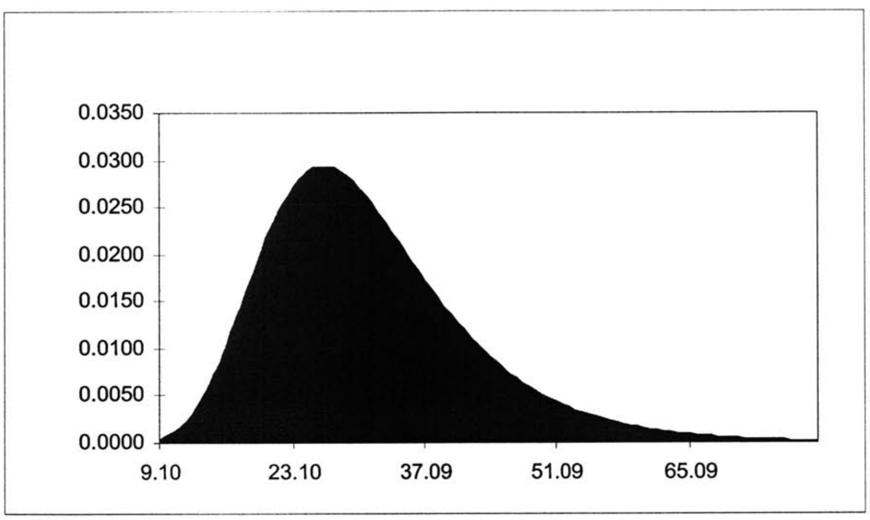

High Season

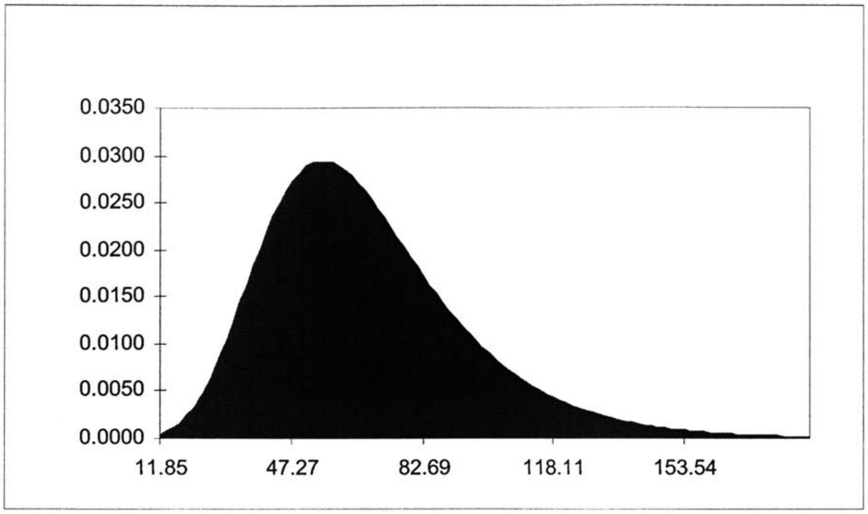

Location H

Summary:
Display Range is from $33,647 to $65,433
Entire Range is from $30,236 to $65,716
After 1,000 Trials, the Std. Error of the Mean is $196

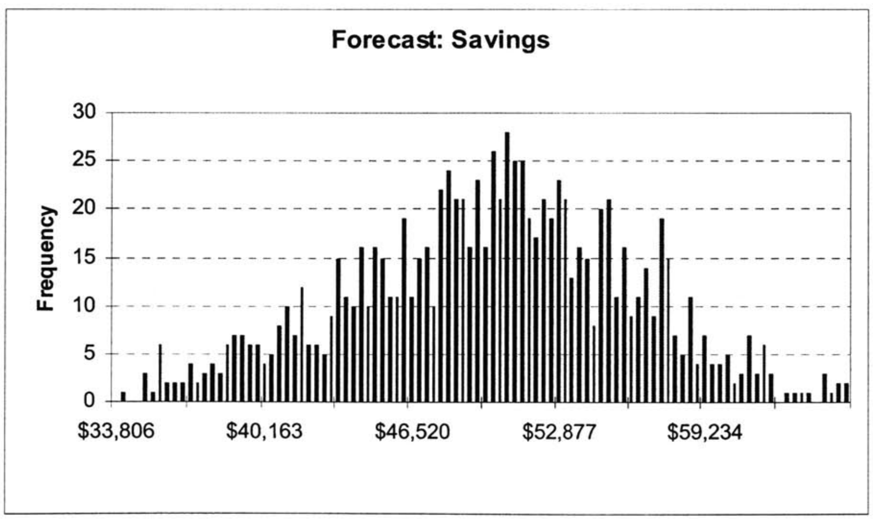

Summary:
Display Range is from 2.20% to 4.28%
Entire Range is from 1.98% to 4.30%
After 1,000 Trials, the Std. Error of the Mean is 0.01%

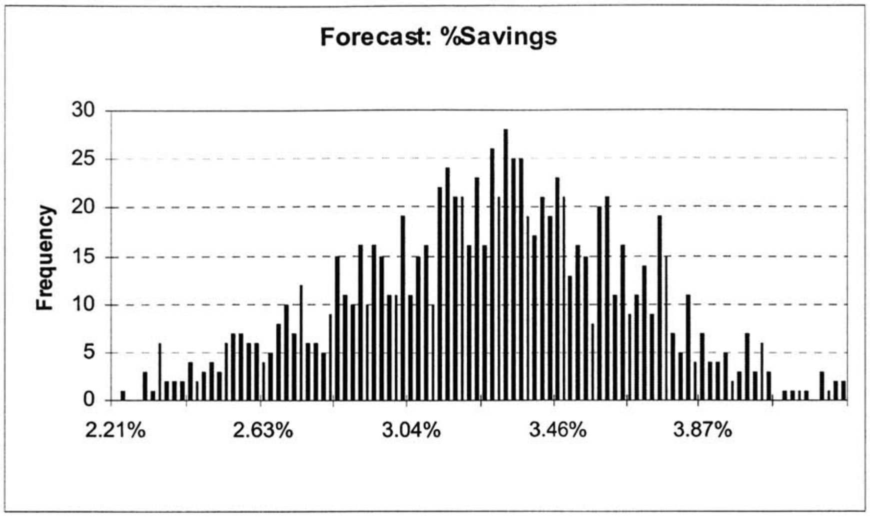

Assumptions: Distribution of Loads

Low Season

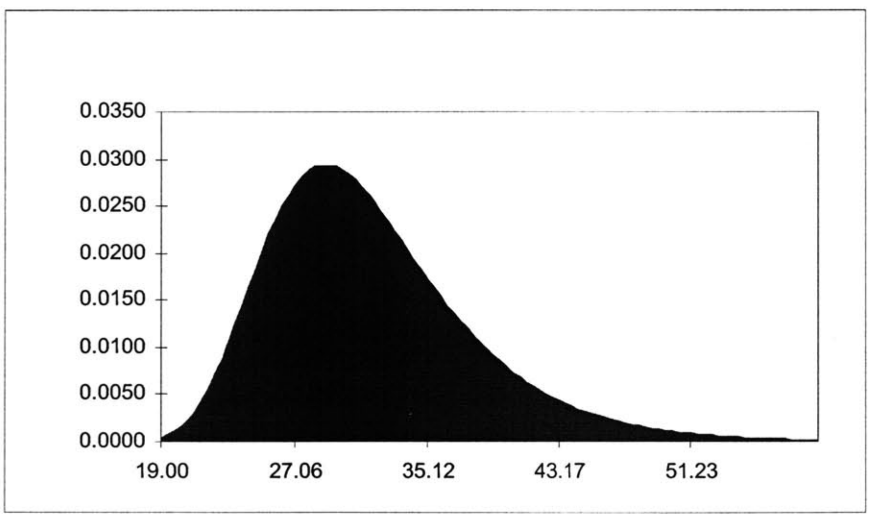

High Season

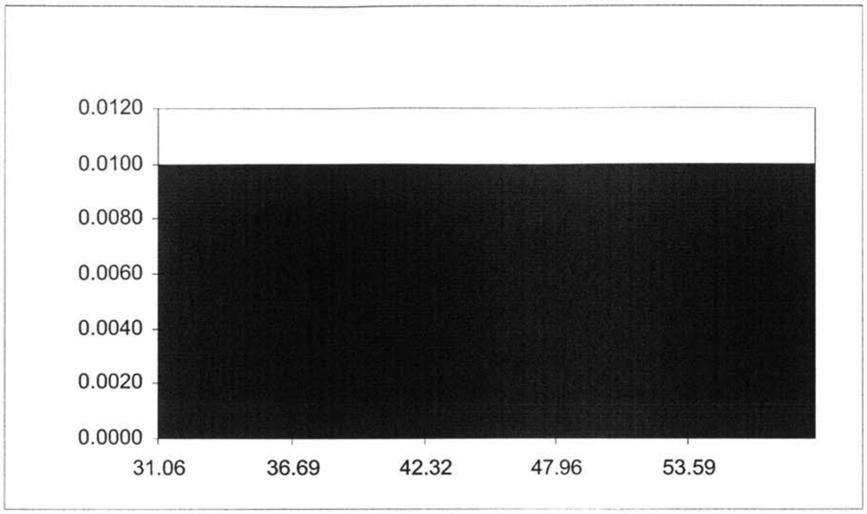
0.0120
0.0100
0.0080
0.0060
0.0040
0.0020
0.0000
31.06
36.69
42.32
47.96
53.59

Distribution of Actual Weekly Loads by Location

Location A

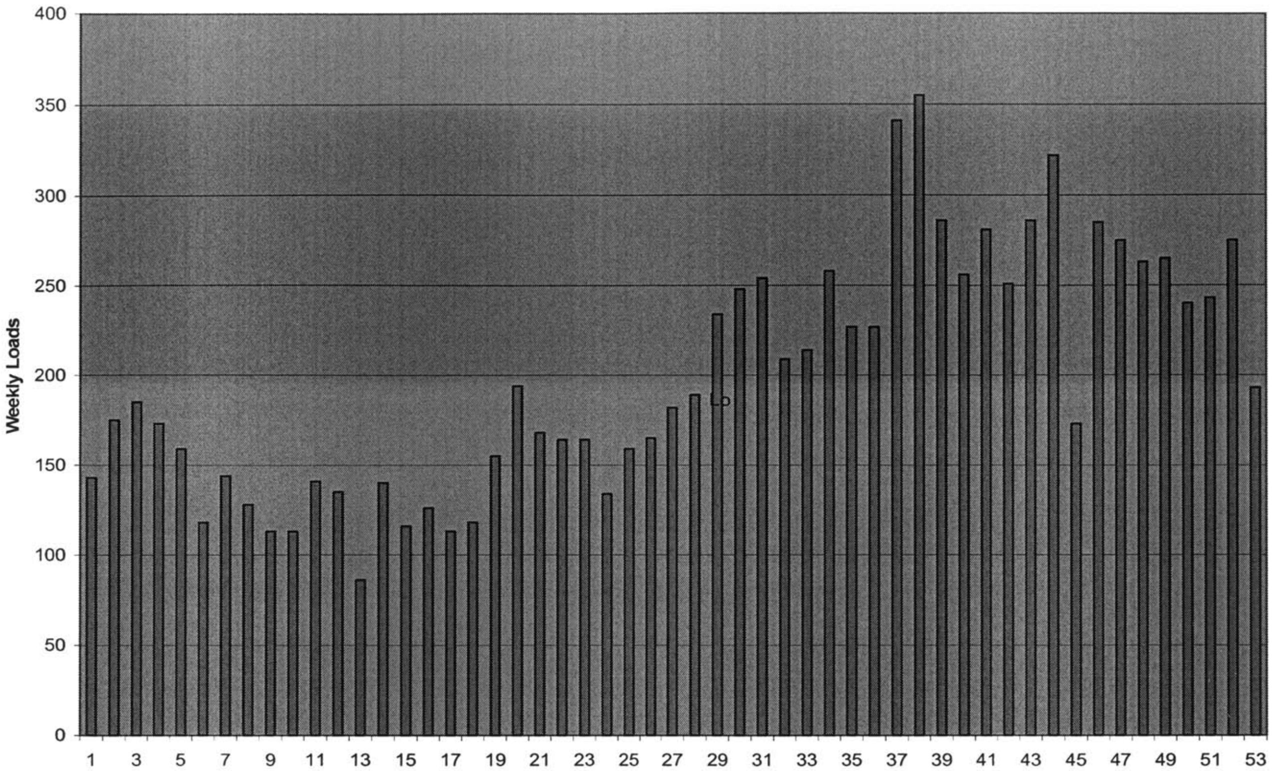

Location B

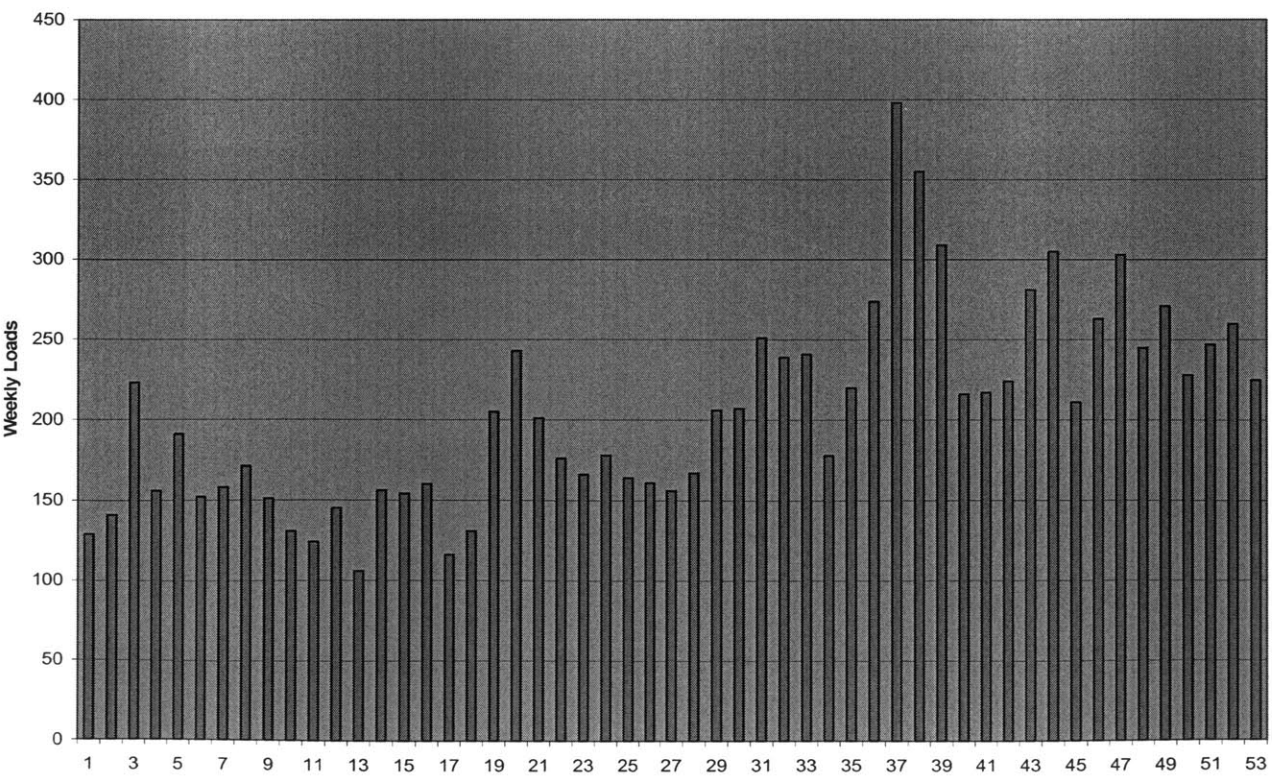

Location C

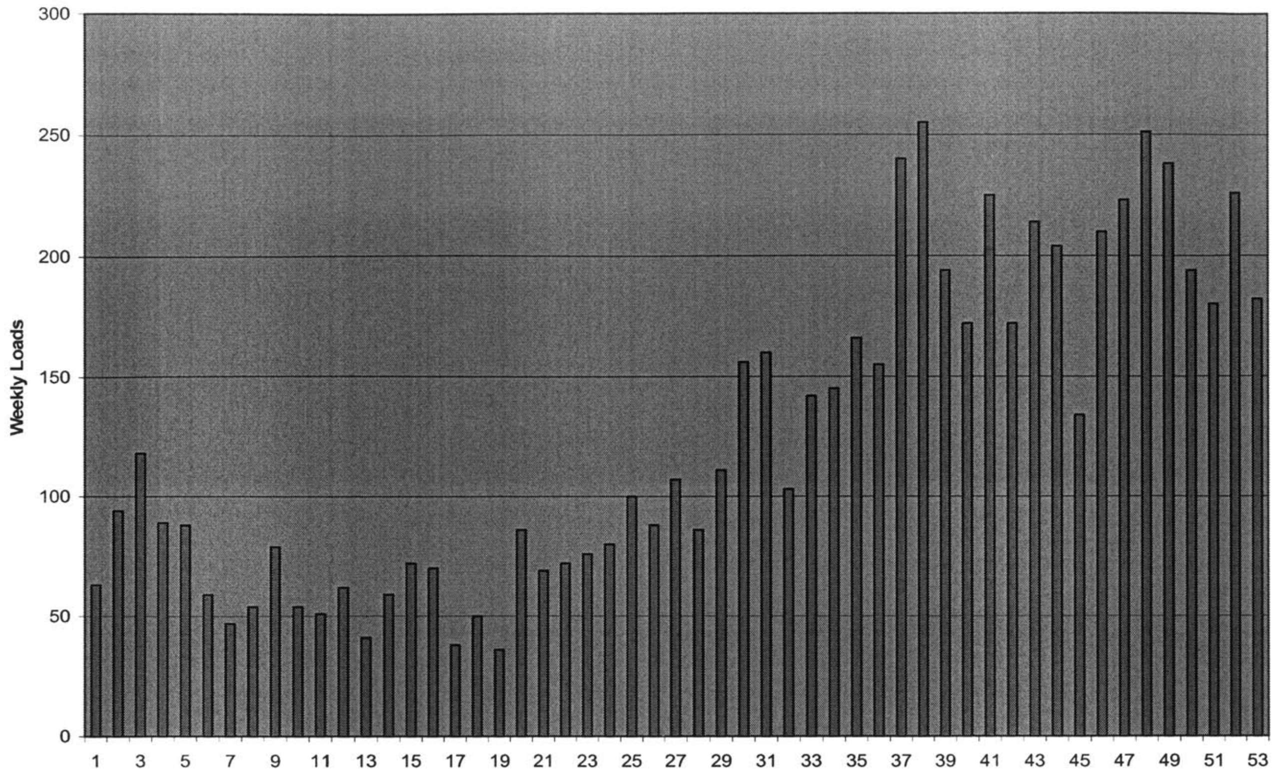

Location D

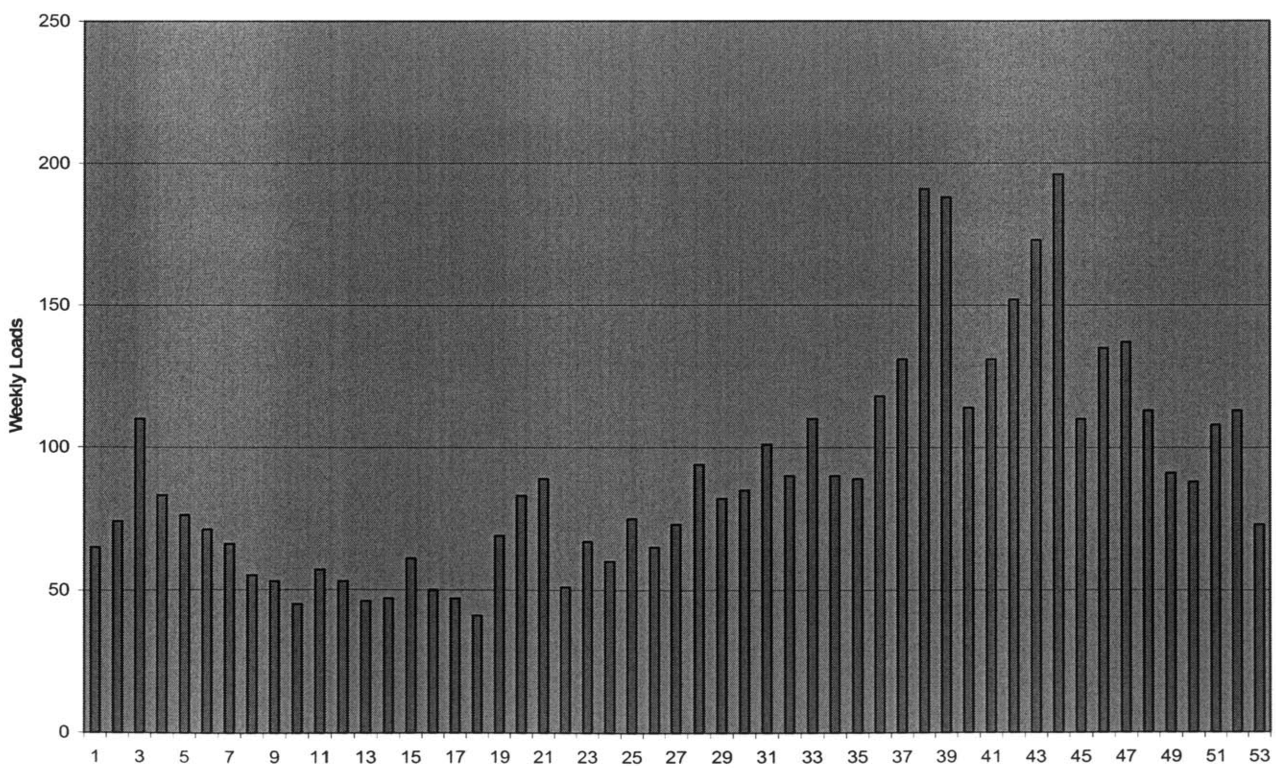

Location E

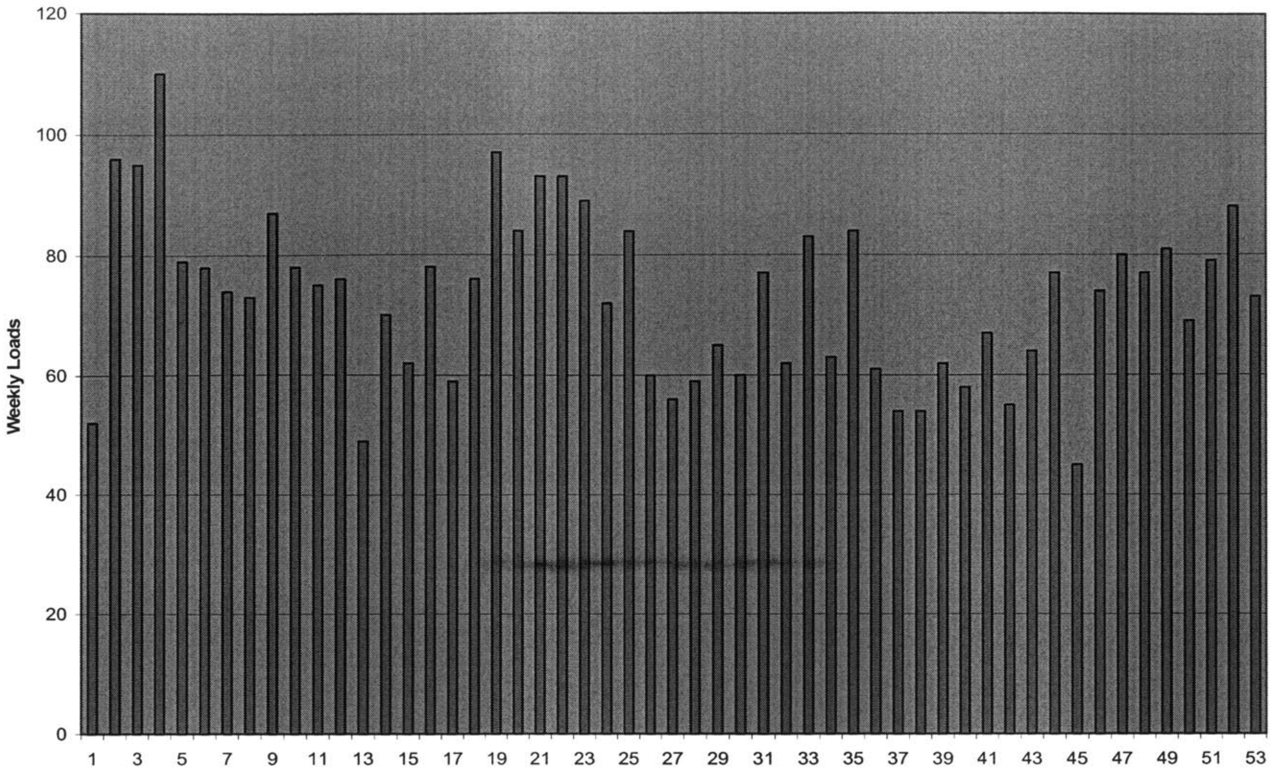

Location F

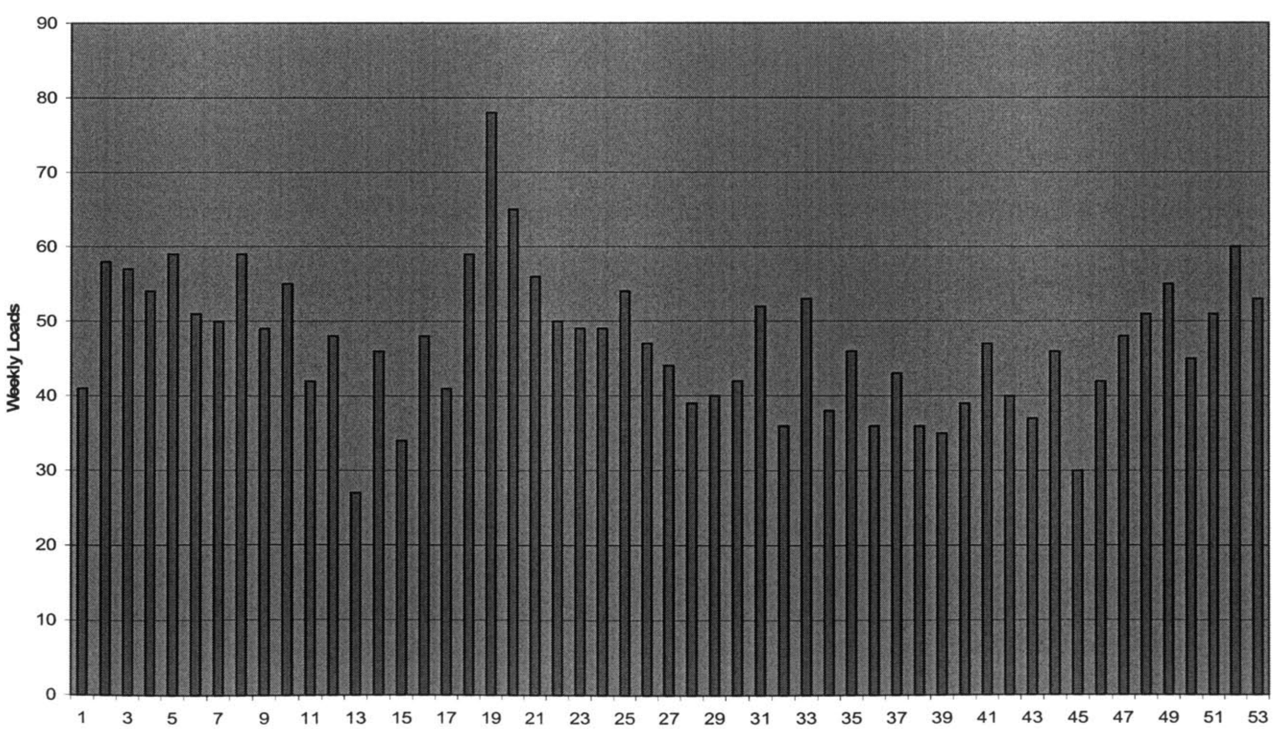

Location G

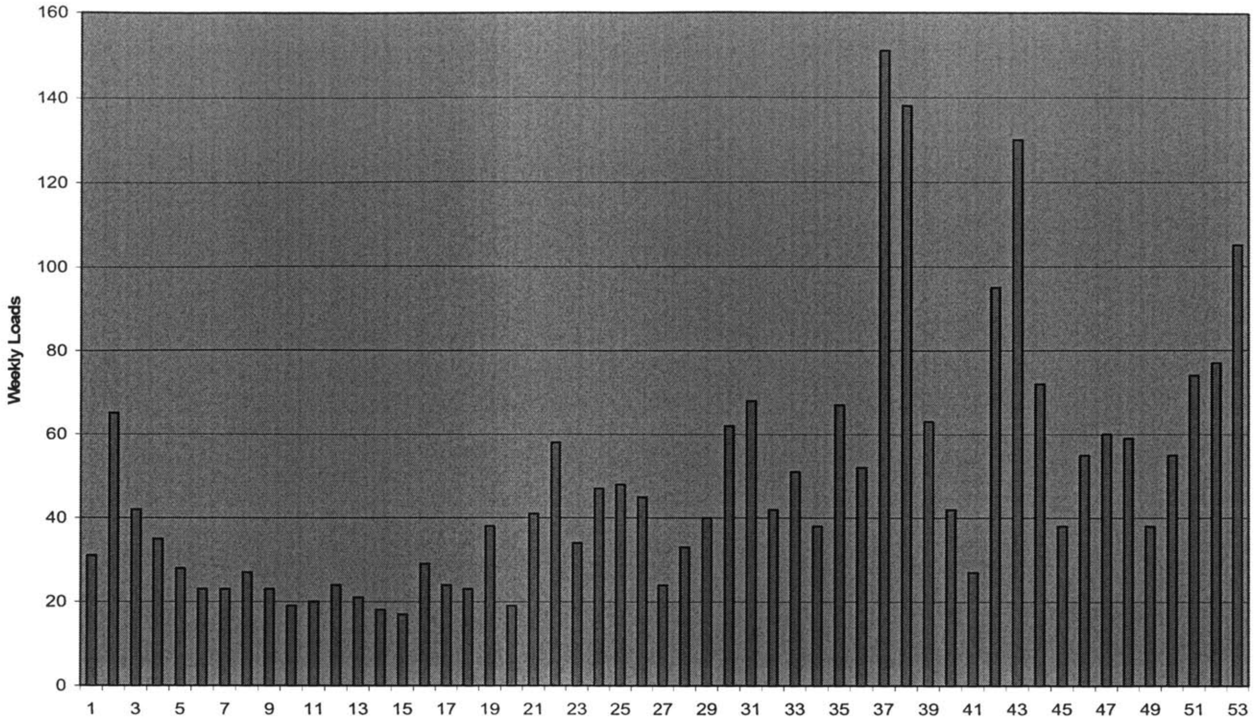

Location H