Time Series Modeling of Spot Energy Prices for Strategic Fuel Management and Gas/Electricity Arbitrage

by

S.R.V.Babu Bangaru

Submitted to the Technology and Policy Program
in partial fulfillment of the requirements for the degree of
Master of Science in Technology and Policy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 1997

Author . . . .
Technology and Policy Program
February 27, 1997

Certified by . . . .
A. Denny Ellerman
Exec. Director, Center for Energy and Environmental Policy Research
Senior Lecturer, Sloan School of Management
Thesis Supervisor

Accepted by . . . .
Richard de Neufville
Professor, Civil and Environmental Engineering
Chairman, Technology and Policy Program

Accepted by . . . .
Arthur C. Smith
Professor, Department of Electrical Engineering
Chairman, Committee on Graduate Thesis

Time Series Modeling of Spot Energy Prices for Strategic Fuel Management and Gas/Electricity Arbitrage

by

S.R.V.Babu Bangaru

Submitted to the Technology and Policy Program
on February 27, 1997, in partial fulfillment of the
requirements for the degree of
Master of Science in Technology and Policy

Abstract

As the electric industry is preparing to embrace competition in a deregulated environment, a savvy owner of a BTU converter (or a generator) is more interested in risk management and economic efficiency. The thesis identifies significant economic value in the day to day management of risks like fuel exposure. The financial firms do not have expertise or resources to help in this area and is solely the arena of plant owner or operator. The strategies, tools developed in this thesis can assist in managing the fuel exposure for a higher profitability by considering the expected payoffs across the commodity markets the BTU converter is dealing with.

The "Spot Energy Price" or "Lambda" model is built as an Auto Regressive Moving Average (ARMA) transfer function noise model with hourly pool load as external input. The model is easy to maintain and uses inputs that can be easily tracked. It bypasses the need for expensive production cost models with enormous and hard to maintain data bases for daily operations such as generator commitment and fuel nomination. An unit commitment algorithm, with inputs such as the day ahead forecasts generated by the λ model, unit dispatch price, operational constraints etc., is developed to economically dispatch a generator. The combination of the λ model and unit commitment algorithm formed a powerful tool to forecast day ahead fuel usage. It also provided a means to evaluate the performance of current methods of fuel exposure management.

A general framework of bid/ask convenience spreads and penalty/discount structure of intra day gas markets is developed for gas/electricity arbitrage and assessing fuel exposure impact.

Thesis Supervisor: A. Denny Ellerman
Title: Exec. Director, Center for Energy and Environmental Policy Research
Senior Lecturer, Sloan School of Management

Acknowledgments

I sincerely thank the following people for their encouragement, inspiration, friendship, and support during my memorable stay at MIT.

Denny Ellerman, Greta Ljung, Richard Tabors, Maria Illic, Richard de Neufville, Gail Hickey at MIT. Jonathan Lowell, Ann Cattarin, Rob Denehey, Steve Rourke, and others with the Portifolio/Risk Management, Fuel, and Generation Marketing groups at New England Electric Systems.

Mary Jane, Ed, Bob, Sreeni, Michael, Horst, Andrew, Daniel, Shiva, Richard, Doug, J, Erica, Karen, Leong, Stacey, Brian, Kathy, John, and several other friends at school, work, dorm and outside.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The role of an electric generator in the evolving spot market for electricity will be that of a BTU converter – with fuel as input and electricity generated as output. However, the fuel and electricity are two separate markets, and an energy producer must manage the risks associated with both markets. Moreover, for a profit maximizing power producer the marginal benefit from electric power and fuel must be roughly equal.

This is a relatively new way of conducting business and certainly a far cry from the way an electric utility, considered as a natural monopoly, operated with a given return on investment as set by the regulating agency. The cost based recovery system in which operating units are compensated at the replacement cost of fuel is changing into a bid based system where the units are compensated at either the unit's bid price or market clearing price at that instant. The bid price itself need not be justified on a cost basis. This decoupling of value of power (unit earnings) from operational costs creates new opportunities to increase efficiency and cut costs.

The ideas developed here, though applicable to other kinds of power generation, are primarily tested for a natural gas powered unit. The reasons are two fold. One reason is that natural gas is increasingly used for power generation world wide. The second reason, probably more important, is that the newer gas powered generators are more flexible in their operation with only a few hours of minimum run/down time requirements.

However, unlike other fuels such as coal and oil, natural gas is generally supplied to the generator site via pipelines and it is not common to store gas on-site. Some pipeline operators need as much as 24 hours advance notice of the amount of gas transported through their pipeline. Hence, the concern for fuel exposure is greater for a gas powered unit. Fuel exposure results when there is a discrepancy between the gas nominated for burns ahead of time and the actual usage.

Wall Street firms, with their expertise in financial risk management and commodity trading, are equipping the industry with powerful tools. Over a longer time horizon, fuel exposure and other risks[1] can be managed using various instruments such as long term OTC[2] contracts, NYMEX futures contracts, commodity swaps etc. In the day to day operation, there is a lot of potential for gas/electricity arbitrage that increases the economic value of a BTU converter. The financial firms do not have expertise or resources to help in this area and is solely the arena of plant owner or operator. The three common ways to manage the *daily* fuel exposure are to

- "must run" the generator according to a specific schedule set ahead of time
- buy or sell gas on the intra day market[3], usually at a premium over the inter day spot index, as discrepancy arises during the day
- change dispatch price for the generator during the day in order to control the number of hours the unit runs economically

Section 1.4 discusses the performance of these above methods as currently used for fuel management. Over shorter time periods, managing the physical risk can take precedence over other risks such as price. However, the bigger goal here is to maximize overall profitability of the BTU converter. Managing the fuel exposure alone with out consideration to the economic value of using the gas available at hand may result in lower profits. A comparison is to be made between the expected payoffs

[1]other risks include basis, transportation or pipeline capacity, gas quantity, embedded options etc. These are discussed in the following sections.

[2]Over The Counter or customized

[3]unlike the inter day spot market where gas is priced today for tomorrow's physical delivery, intra day market is for same day delivery

in the electricity and gas markets. The tools developed in the thesis can assist in this regard.

An introduction to gas unit operation is presented in Section 1.1. Then, the changing landscape of electric industry, as it is deregulated, is examined and relevant issues are briefly discussed in Section 1.2. The economic implications of the fuel exposure management methods, beyond the sole function of matching gas flows, are brought out in Section 1.3. New strategies that methodically consider gas/electricity arbitrage and other economic costs are then identified. Section 1.5 is dedicated to detailed discussion of gas markets as they operate today. While Section 1.5.1 discusses the formation of gas prices, Section 1.5.2 goes into the details of inter day and intra day gas markets and sets up a general framework of convenience spreads and penalty/discount structure to deal with fuel exposure impact assessment and gas/electricity arbitrage.

1.1 Gas Fired Generator Operation

This section is a lead to the generator operation and covers issues related to formation of dispatch prices, gas nomination and timing etc. The hypothetical generator is located in the New England area of Eastern United States.

Each generator prepares a form called NX-12A that contains unit specific characteristics. The form reports the generator's rated Summer and Winter capacity in MW, the minimum and maximum loading levels the unit is allowed to operate, ramp times, minimum run time, down time and start-up time from cold, hot conditions etc. Also given are the block heat rate data, start-up heat input and other fixed costs for hot/cold start-ups, maximum number of starts allowed per week etc.

The NX-12A form can be updated by the owner of the plant when necessary. Some of the operational constraints data such as the minimum down and run times, minimum loading are not necessarily due to the physical characteristics of the generator but are a result of optimization involving:

- unit start-up costs from cold and hot conditions

- plant cycling, wear and tear and longevity
- profit/loss function of expected gain during the on-peak hours and any loss during off-peak hours
- emission allowance for the unit. The NOx emissions are higher during start-up and when operating in lower loading blocks

The NX-12A data are used by the computer models that perform economic dispatch for the power pool. The economic dispatch models take into account the fuel cost ($/MBtu) along with the unit characteristics in determining the least cost way to meet the load or demand. The fuel purchases themselves can be either firm or spot.

1.1.1 Formation of The Dispatch Price

A spot price is synonymous with a floating market price. For example, an inter day spot contract is priced today for tomorrow's physical delivery. The alternative form of purchase called firm contract provides the buyer a fixed amount of gas at a fixed price for each day of the contract, the price and quantity being set at the time of contract purchase. The most common firm contract is a monthly contract priced with respect to the corresponding NYMEX gas futures contract price. In other situations, the utility may contract a fixed pipeline capacity to transport the traded spot or firm gas. Further discussion on gas markets and prices is deferred until Section 1.5.1.

In any case, once the gas price is determined, dispatch price ($/MBtu) for the following day for each generator is calculated as = gas price + any dispatch adders. A common adder is due to emission limits. The dispatch price for the following day has to be quoted to NEPEX[4] by 2 P.M. today. The quoted dispatch price remains in effect for the whole day.

Along with the dispatch price calculated as above, usually up to two additional dispatch prices can also be submitted. Historically, these additional dispatch prices

[4]New England Power Exchange, the central dispatch group for New England utilities power pool

are allowed to accommodate dual fuel units that are capable of switching between alternate fuels during the day. One price can be for firm gas contract while another for spot. Or alternatively, one price can be the inter day spot price while the other two prices are expected intra day gas prices. This option to quote multiple dispatch prices for a single generator is a powerful tool to actively engage in intra day gas/electricity arbitrage and minimizing the fuel exposure impact. The Section 1.3 further explains the use of this option.

1.1.2 Gas Nomination and Timing Issues

Natural gas, as mentioned earlier, is supplied via gas pipelines to the generator site. The pipeline companies set up proper volumes and pressures so that the right amount of gas flows to the correct site. The gas user or trader submits a gas nomination to the pipeline operator with the following information: Date, the (upstream) contract delivering gas to the pipeline, the quantity of gas, and the delivery point where the gas is intended to flow to.

Gas nominations for the following Gas day[5] are required by 10 AM today for the Algonquin (AGT), Iroquois (IRQ), ANR SW leg (ANR), and Appalachia Columbia (TCO) pipeline companies, by 11 AM for the Tennessee (TN) and 12 noon for Trans Canada (TCPL) pipeline companies. Additional nomination updates during the day are accepted by Algonquin which enables intra day transactions between parties connected through them.

The fuel department surveys the gas market by calling various marketers till around 9 AM. After market survey, the fuel department gathers gas usage needs from plant operators who *sense* the plant operation for the following day. After making their own assessment of the day ahead plant operation, fuel department sells excess firm gas[6] in the spot market or buys additional spot gas and submits nomination information to the pipeline company.

[5]Gas day is also a 24 hour day but runs from 8 A.M. of one calendar day to 8 A.M. of the following calendar day.

[6]if there is a firm contract in force as described in Section 1.5.1

The end-user is strongly encouraged to use the nominated quantity with out fail. Overdrawing by more than a small percentage of the nominated quantity is discouraged through penalties besides outright impossible at times. Thus, there is an high incentive to predict the gas usage accurately.

1.2 Restructuring of The Electric Industry – Implications

The fundamental change is the creation of a electricity spot market that is bid based rather than cost based. In a bid based system, the economic dispatch is based on individual generator's bid price. The compensation for a successful unit, i.e. when the unit is committed or called for generation, could be either the unit's bid price or the market clearing price at that instant. The market clearing price or the spot energy price is the bid price of the next unit of generation at this moment in time. Either way, the price of power is no longer based directly on the recovery cost of investment. The significant change to notice is that the earnings potential for a unit is dependent on market characteristics as captured by the market clearing price.

In the current system of cost based unit dispatch, system lambda (denoted by λ) is the incremental cost of the next unit of generation at this moment in time and is measured in dollars per Mega Watt Hour ($/MWH). The hourly system lambda is expected to be highly representative of the market clearing price or the spot energy price in the future system of bid based unit dispatch[7]. The dispatch cost calculated from the dispatch price and block heat rates is similar to generator bid price. The difference being that the generator bid price can be any arbitrary number that does not have to be based on the dispatch price (or fuel cost). This difference introduces a whole new dimension of gaming, market power and manipulation issues in regards to quoting the best generator bid price. These gaming issues are currently investigated by the industry to develop appropriate policies and business practices that curb such

[7]In the state and federal filings by electric utilities, certain justifications of expected future electricity prices are based on the historical system lambda data.

opportunities.

Another relevant issue at unit level is regarding multiple dispatch (or bid) prices and option to switch dispatch prices intra day. This option, now practically free, gives tremendous flexibility to take advantage of any arbitrage opportunities that arise during the day besides allowing efficient management of fuel exposure. The bid based system may impose limitations on the exercise of this free option or impose a fee or both. However, if only one generator bid price is allowed, then the fuel exposure may be covered by transacting in the intra day gas market.

On the financial side of the new power markets, electricity future contracts and market indices allow the unbundling of physical and financial aspects of power transactions. Power swaps and options, power basis swaps and other tools are engineered to manage various identified risks.

1.3 Thesis Objectives

As the electric industry is preparing to embrace competition in a deregulated environment, various strategies and tools are developed for risk management of power generators. While majority of the tools are focused on performance over a time scale of days to months or years, a plant owner would also be concerned with managing the day to day risks like fuel exposure.

Table 1.1 summarizes various opportunities available and certain important factors to be considered for gas/electricity arbitrage over different time lengths. Various combinations of these transactions can be used to build a specific strategy. For example, a company with a long term firm pipeline capacity contract may be interested to fully use the contracted capacity. The company might be interested in buying firm gas long-term and utilize the pipeline capacity to sell on the daily spot market. Alternatively, the company might be interested in selling firm gas + capacity long-term while the company itself buys gas on the daily spot market to cover its positions and needs. The end result in either case is similar to a natural gas swap.

The focus of this thesis is on the last two columns of Table 1.1. The three common

	Horizon of Interest			
	Year or longer	Month	Inter day (day ahead)	Intra day (hour to hour)
Buy	OTC contracts	NYMEX contracts	spot from marketers	spot from marketers or end-users
Sell	OTC contracts	NYMEX contracts	spot to marketers	spot to marketers or end-users
Misc. factors for consideration	firm pipeline capacity	futures prices	daily market index for gas, power	hourly power index
	seasonal storage	long term fuel needs	quoted dispatch prices	bid/ask convenience spread
	unit maintenance	Market movement	unit availability and weather forecasts	forced outages

Table 1.1: Opportunities for Gas/Electricity Arbitrage

ways to manage fuel exposure, mentioned at the beginning of this chapter, also fall under these two columns. The thesis develops strategies to improve on the current ad hoc management[8] of fuel exposure. The thesis demonstrates that forecasting plant dispatch is a tractable problem; a strategy based on fundamental economic principles and simple computer models can aid the gas nomination process while reducing error. Improvement is achieved by methodically considering the opportunity costs and convenience spreads associated with the commodity markets the BTU converter is dealing with.

1.3.1 Strategies and Risk Management Tools for a Gas Fired Generator

The strategy is best represented in a block diagram as in Figure 1-1. The owner of a gas fired generator would normally procure the amount of gas expected to be burnt day ahead in the inter day spot market. The dashed arrow represents the possibility that the firm might have a long-term firm gas contract and might engage in selling in the inter day market.

[8]The Section 1.4 discusses performance of the current methods.

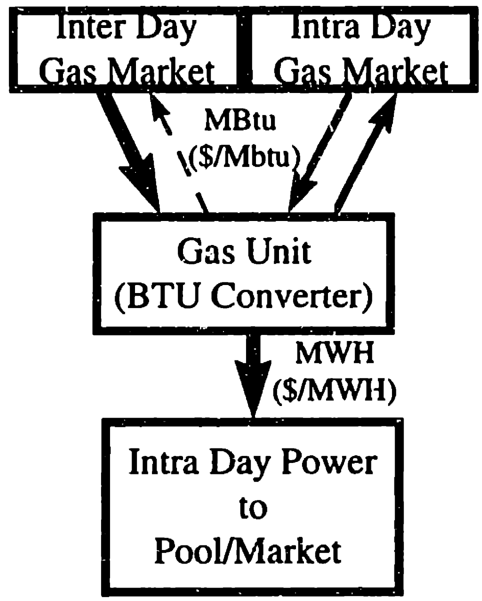

Figure 1-1: Representation of Gas/Electricity Arbitrage

The quoted dispatch cost for each generator relative to the spot energy price determines how much the specific unit is going to run and hence the amount of gas burns. The BTU converter is making a profit when ever its dispatch cost is below the spot price or the spark spread is positive. The process is illustrated through Figure 1-2. The plot shown is of the hourly spot price for electricity. If DC_m is the dispatch cost for the generator, ideally, the generator would be dispatched when ever the market clearing price is above DC_m. While the dispatch cost indicates the value of gas at hand, the spot energy price determines the payoff in the power market. The payoff in the power market is given by $\Sigma(\lambda - \text{dispatch cost}) * \text{MWH}$.

The "Spot Energy Price" or "Lambda" model is built in Chapter 2 to forecast hourly spot prices for up to 48 hours ahead. These forecasts are used to project unit commitment based on the heuristic "commit when pool λ is greater than the

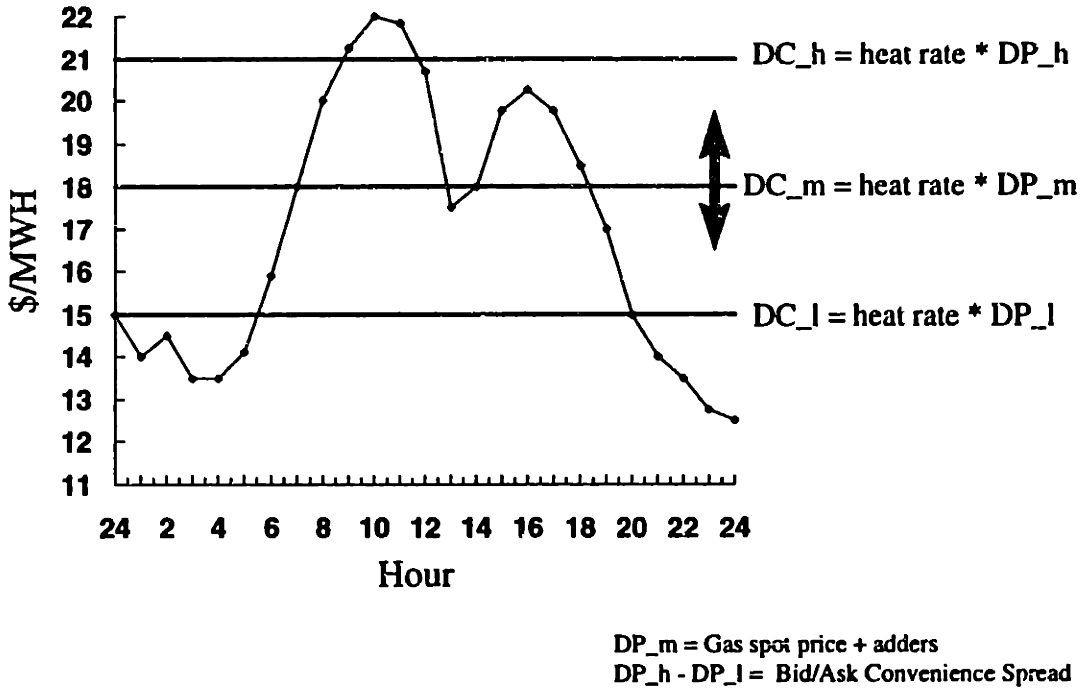

Figure 1-2: Representation of Gas/Electricity Arbitrage

dispatch cost for the generator" or when ever the spark spread is positive and the various operational constraints are met. In the simple case with a single dispatch cost and no option to switch prices intra day, the projected unit dispatch would aid in determining the amount of gas to be procured on the inter day spot market. Any consequent fuel exposure would have to be offset by transacting in the intra day gas market. This scenario with one dispatch cost is the subject of discussion for Chapter 3.

Further improvement in the payoff is possible when multiple dispatch prices are allowed and there is an option to switch dispatch prices intra day. Changing the dispatch price controls the unit operation and hence the volume of gas burned intra day. Effectively, moving the dispatch price up or down translates to volume of gas into or out of the intra day gas market. The effect of switching to a higher or lower dispatch price is shown figuratively in Figure 1-3. Now it is possible to actively take advantage of any intra day price swings by switching dispatch price to *create* a under/over nomination condition and trade the resulting gas quantity for profit. Thus it is possible to actively perform intra day gas/electricity arbitrage or control the potential losses.

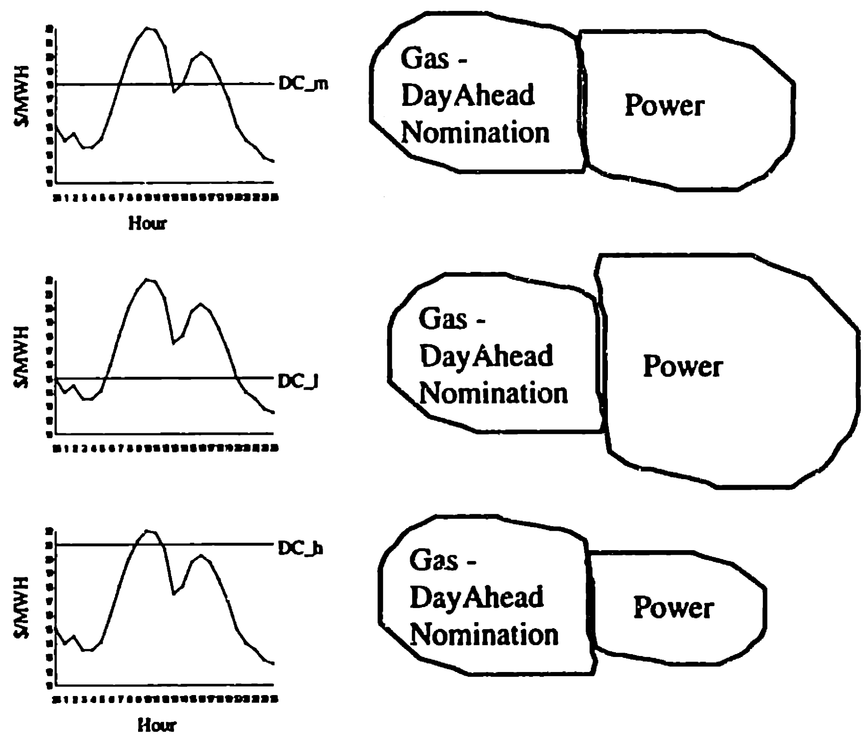

Figure 1-3: Representation of Gas/Electricity Arbitrage

When the intra day switching is unrestricted, it is possible to change to the best dispatch price for any given market conditions. However, the current NEPOOL system allows only up to three dispatch prices for each unit. Moreover, these three prices are to be quoted by 2 A.M. the noon before. In Figure 1-2, DP_h, DP_m and DP_l are the three dispatch prices quoted for one day. DC_h, DC_m and DC_l are the corresponding dispatch costs for the unit. The dispatch price spread can be adapted constantly to track the bid/ask convenience spread. As discussed in Section 1.2, this multiple dispatch price option may be restricted in use or made available for a fee only.

One implicit assumption made here has to do with the electricity revenues. Each generator owner might have entered into certain committed load contracts with the electricity consumers. The load contracts obligate the supplier to meet the demand from the firm's generators or otherwise. Now, when the dispatch cost is increased to take advantage of the high intra day gas prices, the load obligation has to be met by buying power in the market or from a power pool. The assumption then is that the revenue generated by the load contract is approximately same as that paid to full fill the contract. Strictly speaking, this is true only when the load contract is indexed to the hourly electricity spot prices.

1.4 Managing the Fuel Exposure

Fuel exposure results because the firm is obligated to use all the gas nominated day ahead while the actual usage can be different. In the current *ad hoc* process of gas nomination and management, human intuition and mental models play a major role. The current strategies employ the three methods mentioned in the beginning of this chapter to either eliminate or reduce fuel exposure.

The discrepancy between the gas nominated and actual burns is totally eliminated when a generator is "must-run." With must-run, it is specified far in advance exactly how the unit is to be committed during the day irrespective of what the dispatch economics recommend. Then it is straight forward to procure the required amount of fuel gas in advance. Besides, the current method to determine must-run schedule is mostly by trial and error, and learning from the recent past mistakes. Other reasons to must-run a generator include the following:

- limit the unit cycling, start-up costs, mechanical wear and tear etc.
- the unit could be needed for pool security and area protection
- the unit is undergoing compliance testing for emissions etc.

However, this passive strategy does not necessarily maximize profit because the flexibility to respond to changing economics during the day of operation is lost. The must-run strategy is inefficient because it controls a redundant number of dispatch variables while ignoring arbitrage opportunities to make profit or reduce losses. The same task of offsetting the fuel exposure can be accomplished satisfactorily by controlling only one variable like the dispatch price or meeting the gas needs intra day.

When the unit is marked for "economic dispatch" instead of must-run, the actual usage of gas in general differs from the amount nominated ahead of time. The difference in quantities is covered by transacting in the intra day market. Alternatively, the dispatch price for the generator is switched intra day in order to control the number of hours the unit runs. Section 1.1.1 introduced how multiple dispatch prices are allowed for generators. Switching to a higher dispatch price would price the unit out of economics resulting in fewer hours of operation. Switching to a lower dispatch price would have an opposite effect and increase the hours of operation. In other words, the dispatch is adjusted to the nomination. At times, a combination of intra day gas transactions and dispatch price switch is used.

After all these methods to reduce fuel exposure, imbalances in the gas flows do occur in practice due to various reasons such as a delay in communication between the plant operators and fuel department that might bring the unit on-line sooner than expected. Such imbalances between the actual gas burned and gas delivered are cashed out at the end of each month as follows: Let the cumulative gas imbalance for the month = (quantity delivered - quantity burned); a cash *inflow* equal to (imbalance * cash-out gas price) is paid by the gas company when overnominated, where as a cash *outflow* equal to (imbalance * cash-out gas price) + **maximum** interruptible transportation cost is paid to the gas company when undernominated. The average transportation cost charged is ≈43 cents/MBtu of gas! The gas price used in the cash-out is a weighted average of various published daily indices for the pipelines, gas is transported through. This cash-out gas price calculated at the end of month can differ from the gas price used initially in calculating the dispatch price. The settlement policy of paying for transportation cost via. the longest route when undernominated could be an incentive to overnominate gas.

Sometimes, associated with the firm contract of pipeline capacity and physical delivery, the pipeline companies may be extending embedded options that allow considerable daily imbalances cashed out as above. The pipeline operators can accommodate such imbalances, for example, by packing the pipelines higher than usual. These options would be highly valuable for a utility that tends to think of over nomination as a form of contingency storage and hedge against unforeseen risks. The embedded options are often not formal agreements and extended as business privilege to a large customer. Alternatively, the gas pipeline company may have factored in a price through a high demand charge in the fixed price contract for providing such flexibility. Either way, the option is not totally free because of the differences between the price used for settlement and actual price paid, and the transportation charges. An explicit, guaranteed option is likely to be expensive given the tightening demands for pipeline capacity and its seasonal nature.

Figure 1-4 is a plot showing the distribution of daily gas imbalances over the entire year 1994. The gas delivered is the amount of gas nominated day ahead (inter day) **plus** any other intra day transactions made. The plot represents the performance and end-result of current gas management methods.

The distribution has a mean $\approx$150 MBtu, and a standard deviation of $\approx$5200 MBtu. The pipeline operators were somewhat lax and allowed considerable daily imbalances[9] probably due to the firm capacity contract with the electric company. In case of inter day spot deals, large deviations from the nominated quantity are not possible. Interestingly, the imbalances in the plot are weighted towards overusage as indicated by the long tail in the bottom half, positive skewness (1.233) and high kurtosis (5.685 + 3) statistics. The monthly imbalances are calculated from the daily data and presented in Table 1.2. The table shows a tendency of underuse in Summer months and overuse other times. It also represents the extent to which embedded options were used to cover gas flow imbalances each month. Further analysis presented in Section 3.2 suggests that there may have been profit making opportunities that

[9]The cash-out gas price data are not available to estimate the total price of imbalance option.

Month	Imbalance (MBtu)
February	0
March	0
April	-3,361
May	-10,441
June	-20,104
July	-22,116
August	33,155
September	38,342
October	16,988
November	32,065

Table 1.2: Gas Flow Imbalances by Month

were not exploited by the current system.

Of significant interest would be a comparison of the day ahead gas nomination with daily burns. Unfortunately, the current procedures of book keeping does not separate the day ahead nomination data from the intra day transaction data. If such data were available, then one can look for any recurring biases in the nomination process and create new policies to minimize such errors. For instance, if there is a pattern or strong autocorrelation among the daily under/over nomination error sequence and/or there is a strong correlation with the day-to-day price changes, the nomination process would warrant scrutiny. A quick check of the final imbalances in Figure 1-4 did not reveal any such covariances or autocorrelation. However, this does not rule out the possibility of strong correlations in the day ahead prediction error. Never the less, it is comforting that the intra day transactions or dispatch price switching or combination of the both are offsetting any prediction biases.

1.5 The Gas Markets

The details of inter day and intra day gas markets, formation of gas index and storage options for natural gas are discussed here in greater detail.

1.5.1 The Gas Price

Ideally, Spot price in New England would be = Henry Hub[10] index price + transportation cost. But the New England price is rarely so because of the peculiarities of local supply and demand for natural gas and the gas marketer's profit margins. The resulting difference in gas prices between the various index locations is referred to as the basis differential. The basis differential is due to commodity risk as well as transportation risk. While both risks are highly seasonal, the transportation risk is significantly higher than commodity risk during Winter months. For example, interruptible contracts are almost unavailable during Winter when the capacity is always used. It may happen that gas is available but transportation is not possible and vice versa. For particular locations, these basis differentials can result in large bid/ask convenience spread as further discussed in Section 1.5.2.

In order to mitigate some of these risks, the utility may sometimes enter into long term firm pipeline capacity contracts. In this case, the utility would buy gas at the well head and use the contracted pipe line capacity to transport the gas. Then the gas price would be:

$$\begin{aligned}\text{Gas price (\$/ MBtu)} = {} & \text{commodity price} + \text{demand charge} + \\ & \text{fuel charge} + \text{other O\&M associated costs}\end{aligned}$$

Commodity price (\$/MCFD[11]) is a charge determined at the well head. Most commonly, this is a fixed price for each month. Demand charge,also a fixed amount, is the cost of using pipeline capacity. The pipelines use compressors to maintain pressures that support gas flows. The fu charge is accrued due to compressor's use of the gas transported as its fuel. Alternatively, gas and pipeline capacity n be purchased or sold on the spot market as a bundle.

[10]designated delivery site for NYMEX natural gas futures

[11]Million Cubic Feet per Day

1.5.2 Inter Day, Intra Day Gas Markets, Penalty/Discount Structure, Arbitrage and Price Impact

The inter day spot market is active with a large number of participants. Reliable market indexes exist; one can buy or sell large quantities of spot gas inter day with minimal or no impact on the price. On the other hand, intra day markets operate over regions where pipeline companies can accommodate intra day gas movements. A regional intra day gas market has fewer market players, has no published market index, but never the less provides ample opportunities to manage fuel exposure and arbitrage.

The high volatility of intra day gas market creates arbitrage opportunities of two kinds. Firstly, there is a magnitude difference between the inter day and intra day prices. This could reflect the utilities of various intra day market participants and change in demand characteristics from day ahead to intra day at a particular location. It is not uncommon when the intra day prices went up to as high as \$12/MBtu when the inter day spot index for the same day was only \$3 or \$4/MBtu. Of course, the spot prices for the following day soared high indicative of the high forward correlation between today's intra day prices with tomorrow's inter day prices.

Secondly, there is a gas convenience penalty/discount structure that reflects the illiquid nature of the intra day market. Each day the intra day price goes up or down or remains the same relative to the inter day price. Similarly, a fuel exposure may occur or a perfect forecast is made and no exposure occurs. With no fuel exposure and no price change, obviously, there is no action to be taken to realize higher profits or reduce losses. If there is a fuel exposure, it is due to either under or over nomination. With a fuel exposure but no change in prices, one simply off-sets the exposure by transacting accordingly in the intra day market. Again profit is maximized. The penalty/discount structure becomes important when there is a fuel exposure and a price change.

The daily fuel exposure series and the inter to intra day price changes can each be random[12], but there can be a correlation (or dependency) between the price changes and fuel exposure. This correlation is expected because a large proportion of natural gas users are electricity producers, especially in Winter. Another reason for the correlation can be the perceived price impact by a large generator operator.

When tne fuel exposure is due to undernomination (i.e. gas has to be bought to off-set the exposure), and the price change is upwards (i.e. intra day price is higher than the inter day price), then the gas is purchased at a premium price and the generator owner is penalized. The day in which this happens is called a "penalty day" or "premium day." Perhaps the price change is downwards when the fuel exposure is due to undernomination, then gas can be bought at a discount price and the generator owner is making a profit. The day in which this happens is called a "discount day" or "arbitrage day." On a discount day, more gas is bought at a cheaper price to produce more electricity, most likely for a profit in the electric market, thus making a perfect gas/electricity arbitrage. Similarly, a discount day occurs when there is overnomination and the price change is upwards; a penalty day occurs when there is undernomination and the price change is downwards.

The final profit depends on the magnitude of price change, and the kind of days generator sees often. If there is a price impact, it reduces the profit or even changes an otherwise discount day to premium day provided the magnitude of price impact is large. Similarly, multiple dispatch price option would minimize losses or change an otherwise premium day to a discount day through dispatch price switching.

This penalty/discount structure can be understood as a bid/ask convenience spread. The convenience spread is the difference between inter and intra day prices for gas. Each day the generator owner would be interested in either exclusively buying or selling intra day gas. When gas is purchased, the spread is called bid spread; when gas is sold, the spread is called ask(ed) spread. Thus on a arbitrage day, the generator gains or enjoys the spread as a discount as if for providing the market liquidity. On a penalty day, the generator incurs the spread as a cost for transacting in the intra day

[12]The forward correlation between today's intra day prices with tomorrow's inter day prices can still subsist with purely random inter to intra day price changes.

Fuel Exposure Due to	Inter to Intra Day Price Change	
	Upward	Downward
Undernomination	Premium or Penalty Day Bid spread paid	Arbitrage or Discount Day Bid spread gained
Overnomination	Arbitrage or Discount Day Ask spread gained	Premium or Penalty Day Ask spread paid

Table 1.3: Penalty/Discount Structure of Intra Day Markets

market. Hence the terminology bid/ask convenience spread. All these combinations are represented in a matrix in Table 1.3.

Maintaining the price changes or convenience spreads as separate bid or ask spread would allow us to analyze and answer questions regarding market or price impact, symmetry of penalty/discount structure, etc. For example, it is useful to know if one spread is wider than the other, or on average if the individual spreads are positive or negative. Defining bid spread = (intra day bid price) - (inter day index), it is profitable if the bid spread is negative on average. Similarly, defining ask spread = (intra day ask price) - (inter day index), it is profitable if the ask spread is positive on average. In Section 3.4.1, actual spreads are estimated from real transaction data are used in assessing the impact of fuel exposure. These spreads are also useful in training or adapting the strategies developed in Section 1.3 to maximize the expected pay-offs.

The arbitrage opportunities discussed so far are not pure arbitrage in the sense that the same commodity is not bought cheap and sold high. Although such transactions do occur, it is hard to imagine frequent opportunities like that exist in a market of few players who are well connected in a small region. If intra day gas can be sold at \$2.5/MBtu and bought at \$2/MBtu, it is a pure arbitrage opportunity irrespective of what the inter day index is. Any such real spreads disappear quickly and are also limited by the quantity available. One situation, such a spread can be seen is as a price impact: When a major quantity of gas is offered for sale intra day, say due to forced outage of the generator, the price procured for that quantity can be below the current market price. But if the generator comes back on-line in a few hours and the gas has to be bought back, it may only be available at a higher price than it is sold at. Some one else made pure arbitrage profit at the generator owner's expense. This is called moving the markets against oneself! Hopefully, such situations do not arise often.

1.5.3 Storage for Natural Gas

On-site storage of natural gas for power generation is rare and considered expensive. Off-site Storages can be of daily or seasonal type. For New England utilities, both seasonal and daily storages are available in NJ, PA, and TN.

Salt dome type seasonal storages typically inject/store during the cheap Summer months and extracted during Winter. Storages are contracted similar to pipeline capacity – firm or interruptible and are found to cost as much as $1.50/MBtu. Daily storages are mostly used by marketers although some interruptible contracts can be available for end users.

Another kind of gas storage is pipeline storage. Though pipeline storage sounds attractive, it is least available when it is most needed during Winter. There is usually no fixed demand charge for the storage capacity and is available on a first come first served basis. Typical storage costs are 6 cents/MBtu for injection and extraction and 1 cent/MBtu per day of storage.

And finally, some industry experts consider dual-fuel and co-firing capabilities of a generator as a form of storage. With the additional capability, one can switch fuel or change the ratios of burns mixture to manage the quantity of fuel used.

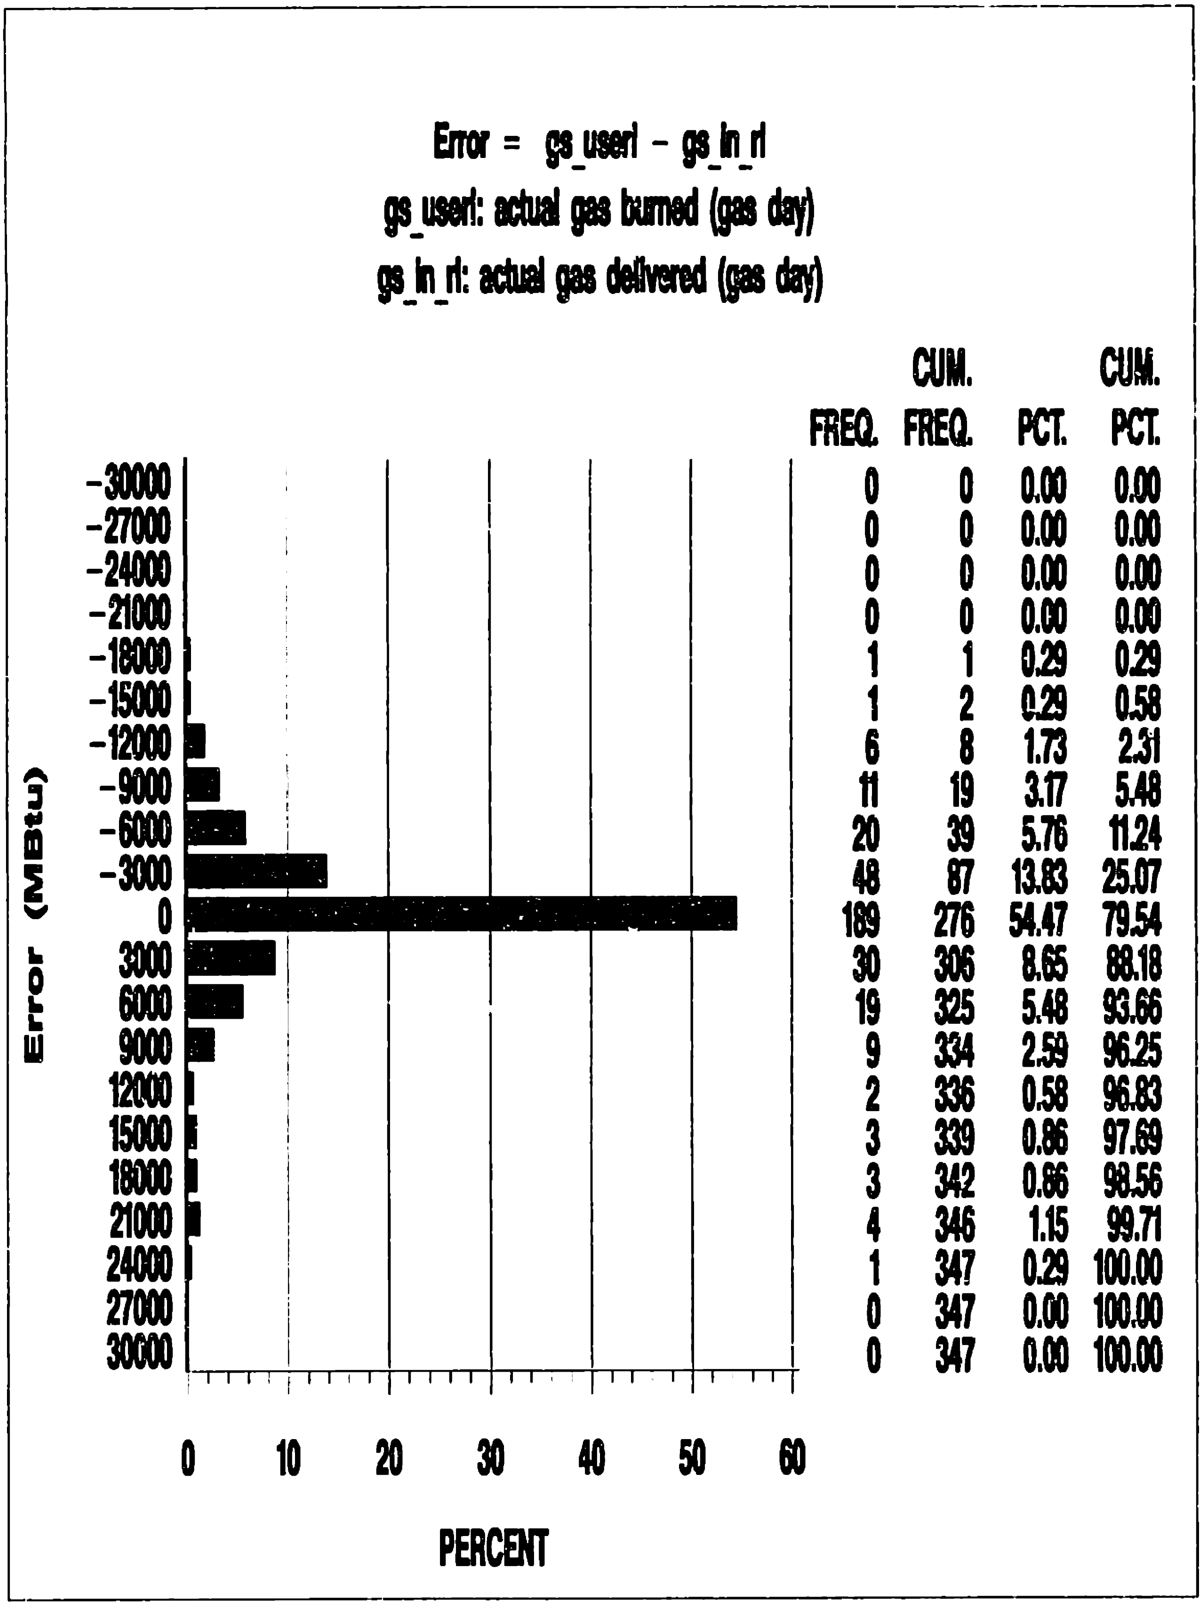

Error (MBtu)	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-30000	0	0	0.00	0.00
-27000	0	0	0.00	0.00
-24000	0	0	0.00	0.00
-21000	0	0	0.00	0.00
-18000	1	1	0.29	0.29
-15000	1	2	0.29	0.58
-12000	6	8	1.73	2.31
-9000	11	19	3.17	5.48
-6000	20	39	5.76	11.24
-3000	48	87	13.83	25.07
0	189	276	54.47	79.54
3000	30	306	8.65	88.18
6000	19	325	5.48	93.66
9000	9	334	2.59	96.25
12000	2	336	0.58	96.83
15000	3	339	0.86	97.69
18000	3	342	0.86	98.56
21000	4	346	1.15	99.71
24000	1	347	0.29	100.00
27000	0	347	0.00	100.00
30000	0	347	0.00	100.00

Figure 1-4: Gas Flow Imbalances

Chapter 2

Spot Energy Price Model

The "Spot Energy Price" or "Lambda" model is a model of system lambda, where system lambda (denoted by λ) is the incremental cost of the next unit of generation at this moment in time and is measured in dollars per Mega Watt Hour ($/MWH). This number is updated every few seconds and originates from the NEPEX's[1] automatic economic dispatch system. NEPEX then calculates a single number for each hour, from all the real-time lambdas generated during that hour, which is considered to represent the cost of electricity for that hour. From here on, this hourly number is referred as the "system lambda[2]" or simply λ. NEPEX also files this λ with FERC[3] each year.

Figure 2-1 is a plot of actual hourly λ (in $/MWH) and pool load (in MW) over a week or 168 hours. From the definition of λ, it is a function of the available generation and load at that moment. The hourly load, of course, is a complex function of weather, consumer characteristics, building properties etc. The cost of generation is dependent on the participating generator's fuel cost, unit characteristics such as heat rate curves, ramp times, dispatchable or must take operation, scheduled, unscheduled outages etc.

Electric utilities and power pools maintain an enormous data base of generator characteristics, power contracts, load forecasts and other factors mentioned above.

[1]New England Power Exchange

[2]Refer to Section 2.4.1 for a discussion on limitations of this hourly λ and other kinds of λs used occasionally in the industry.

[3]Federal Energy Regulatory Commission

These data bases feed complex production cost models[4] that can simulate the economic dispatch and estimate the hourly $\hat{\lambda}$. This $\hat{\lambda}$ can be pretty good forecast provided the inputs to the production cost model are accurate enough. However, the level of effort necessary to update and maintain such a database is just hideous. Moreover, detailed data as required may not be readily available for competitor's generators.

Thus the challenge is to create a model that requires minimal maintenance, uses factors that can be easily tracked yet produce forecasts of acceptable accuracy. Though the literature search revealed no previous work carried out to model hourly electricity spot prices, one can see the parallels with other fields[5] including stock prices. The most recent I heard is about a company marketing neural network models for predicting hourly spot prices. The following sections describe the model building process in Section 2.2, discussion of various drivers of λ in Section 2.1 followed by model evaluation and testing in Section 2.3. The final section in this chapter, Section 2.4 discusses the limitations due to various data used to build the model, and the model itself.

2.1 Factors that Affect Lambda

Though it is tempting to include all the known λ drivers in the model because it produces a good fit on the historical data, a large number of input variables reduce the forecasting power and widen the confidence band of predicted $\hat{\lambda}$. The minimum factors that have a significant and direct impact on the marginal cost need to be identified by reason, intuition, and by way of model statistics like t-stats, AIC, BIC etc. One should be able to understand how a particular factor affects the variable modeled, check the understanding against expert knowledge besides that it passes the statistical tests. Such a factor will add value to the forecasts.

At a broad level, the drivers of λ either fall under the demand or supply side. The demand side factors are load and several variables that drive the load. These

[4]such as Polaris, ProdCost, ProSim

[5]There is significant amount of ongoing research to forecast hourly load; see bibliography for some references

are discussed in Section 2.1.1. The supply side factors include unit availability, unit characteristics, outages etc. These are discussed in Section 2.1.2.

2.1.1 Demand Side Drivers

Looking at Figure 2-1 one can expect a high degree of correlation between the pool load and λ. The intuition is confirmed by Figure 2-2, which is a scatter plot of hourly load and λ, and the statistical estimates of covariance and cross-correlation coefficient are given below:

$$\mathrm{Cov}(\lambda_t, \mathrm{load}_t) = 9.248$$

$$\sigma_\lambda = 2.686, \sigma_{load} = 4.639$$

$$\mathrm{Corr}(\lambda_t, \mathrm{load}_t) = 0.742$$

Weather is considered to be a prime driver of the hourly load and hence the generation cost. NEPOOL[6], for example, stores over 30 years of data from 10 weather stations. The weather is represented by raw measurements of temperature, humidity, wind speed, cloud cover, dew point and derived variables like wind chill factor, average temperature, cold/hot Degree Days, temperature-humidity index etc. Simple linear regression of lambda over these weather variables did not possess a significant explanatory power.

Other factors like thermal properties of structures affect the load on a slightly longer time scale. The thermal build-up of buildings, for example, can be captured by a moving average of daily temperatures. However, the influence of these factors is reflected in the actual hourly load data, and load forecasts generated by NEPEX.

One example is cloud cover — it had no significant t-stat in the simple regression described earlier. However, some utility load models are said to use cloud cover with a 5% weight in their models. Weather centers use cloud cover to calculate the amount

[6]New England Power Pool

of Sun light(in lumens) reaching the surface of the earth. Overall, hourly pool load[7] is the single variable used in actual model building.

2.1.2 Supply Side Drivers

A large number of generators of all kind, and size participate in a central economic dispatch so that the current load can be met with the lowest cost of generation available from the entire set of generators. There could be instances often when a generator is running below its normal rating but the next cheapest available unit of energy is not coming from this unit but from a different generator with competing heat rate blocks. Although, the fine detail of unit characteristics and a myriad other variables seem to play a role in determining the marginal cost, just as in the case of weather driving the load, the current and past λ reflect the influence of these factors.

In a large pool of generators, significant autocorrelation of the λ data and a high degree of cyclicity can be expected. Thus the complex production cost models with hundreds of input variables may not be necessary to predict the day ahead marginal costs with reasonable accuracy. A model that adaptively updates few model parameters based on present and past λ data are all that may be required. In fact short term power marketers strike deals partly due to their intuitive understanding of simple measures such as moving averages, trends, and expected load, unit outages. More details on the forecasting power of such a model, and the statistics of λ are presented in the section on model building . The instances when such a model can be inaccurate is when some thing happens suddenly with no prior indication — like a major unit outage.

Unit outages can occur either scheduled or unscheduled. Scheduled outages are mostly for reasons like unit maintenance or testing. An unscheduled or forced outages occur usually because a unit trips due to a mechanical malfunction. A scheduled outage or bringing online of a base loaded nuclear unit, say of 1000 MW, can move the average daily λ up or down by as much as \$2\$/MWH and remain there. A

[7]Refer to Section 2.4.2 for a discussion on limitations of this hourly load data

forced outage of a major unit, on the other hand, can double or triple the hourly λ temporarily until a cheaper unit picks up the slack. The forced outages are inherently random which makes it impossible to model it with an hour to hour accuracy. Refer to Section 2.4.3 for a discussion on how unit outages can be included in the model and any limitations with the available data

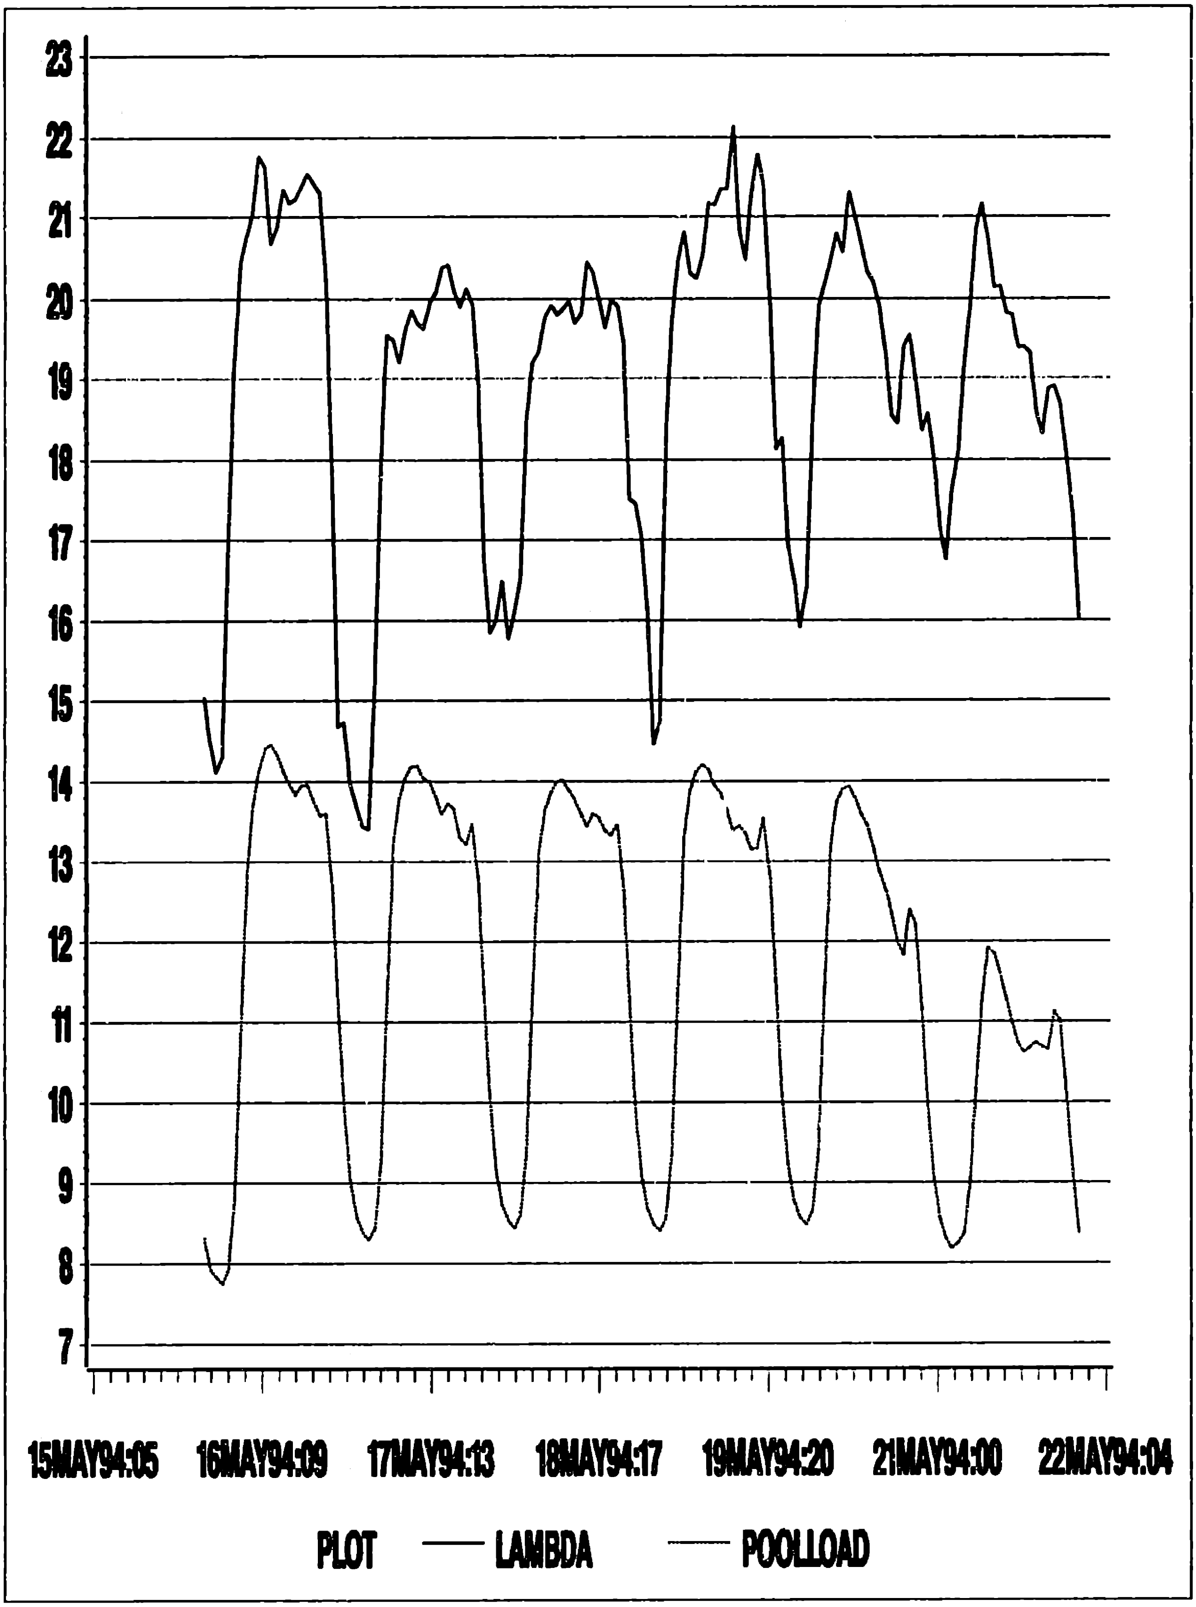

Figure 2-1: Typical hourly Lambda and pool load

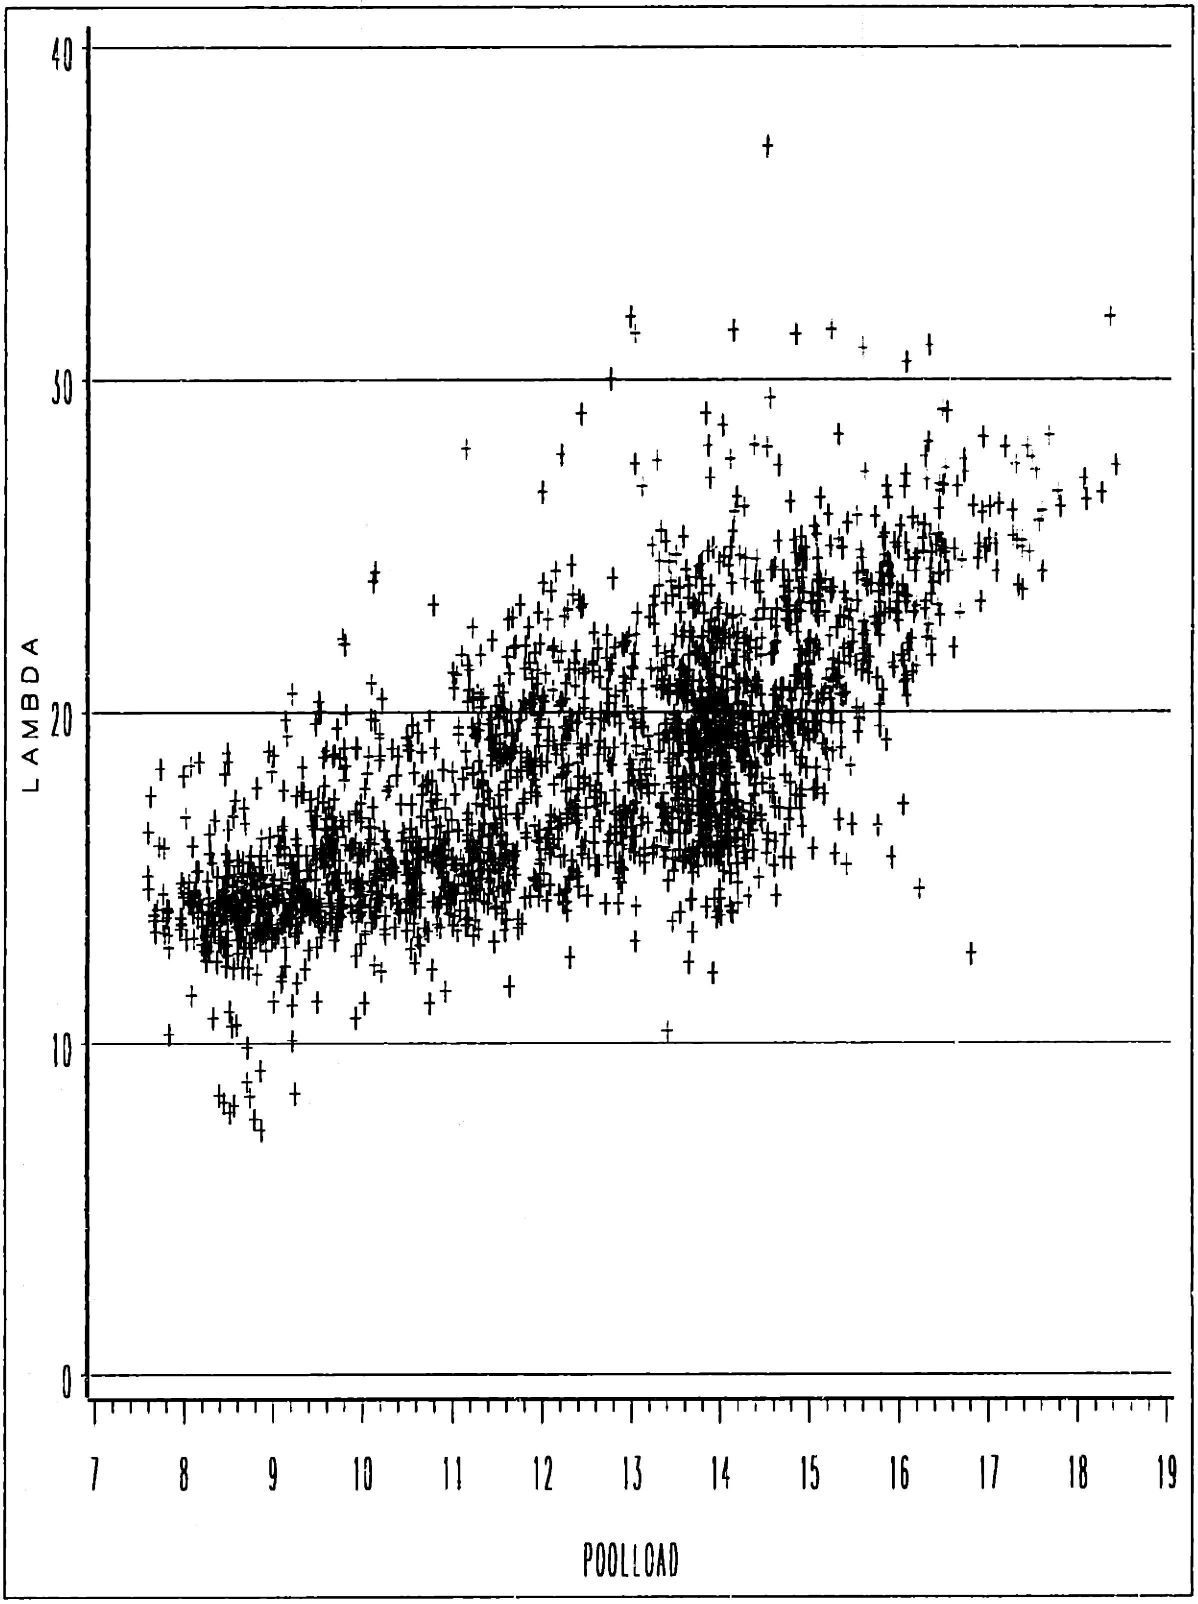

Figure 2-2: Scatter plot of Lambda Vs. pool load

Moments	
Number of observations	8,760
Mean	19.569
Standard Deviation	4.639
Skewness	1.288
Kurtosis	4.70
Quantiles (Def=5)	
Maximum	59.26
Median	18.975
Minimum	6.69
Mode	18.52

Table 2.1: Relevant statistics of hourly λ in year: 1994

2.2 Model Description

The hourly system λ is highly cyclical and has identifiable patterns that repeat with day to week periodicity. Various drivers that influence λ also exhibit same kind of cyclicity. A part of the estimated spectrum of λ is shown in Figure 2-3. Clearly, the spectrum is dominated by the Period = 12 and 24 hour components. The complete spectrum shows peaks of decreasing amplitude at hormonics of 12 hour period. Now, Figure 2-4 is same as Figure 2-3 with the 12, 24 hour period components removed. There are hardly any other strong periodicities left now.

Statistics on λ population are presented in Table 2.1 and Figure 2-5. The statistics are estimated over the 8,760 data points in year 1994 and affirm the general intuition of the short-term generation marketers that hourly $\lambda >$ \$35/MWH is highly likely to be an outlier; 3*the estimated standard deviation + mean is close to \$35 from below. This information is used in the one-stage model described in the model identification Section 2.2.1. One statistic of interest can be the distribution of the hour-to-hour changes in λ. The standard deviation estimate of the hour-to-hour steps comes out to be ≈ 1.887. And almost always $\lambda >$ \$35/MWH when λ changes more than 5.66[8] from hour-to-hour. Figure 2-6 is the distribution of these $\Delta\lambda = \lambda_t - \lambda_{t-1}$,

[8] 5.66 = 3 * standard deviation of hour-to-hour jumps in a whole year.

2.2.1 Model Identification and Estimation

The model is required to produce day ahead (48 hours ahead) forecasts $\hat{\lambda}$. Various models for the the response series, λ, were identified and tested to select the two final models that differ in the way outliers are handled. One is a *two-stage model* where the input series are included in two stages while the other one is a *one-stage model.* The comparison of models is deferred until Section 2.3.

Both the λ models are Auto Regressive Moving Average (ARMA) transfer function noise models with other external series as inputs. The ARMA (p,q) model has a Auto Regressive (AR) part of order p, and a Moving Average (MA) part of order q. For the identified model, p = (1)(24) implies the λ is modeled as an average value plus some fraction of deviation from this average value 1 hour ago (*recent effects*), and 24 hours ago (*periodic effects*), plus a random error; q = (1)(12) implies that observations 1 hour apart and 12 hours apart are correlated, so that the *memory* of the process is just 12 periods. The mathematical form of the ARMA terms in the model is:

$$\Phi(B)\lambda_t = \Theta(B)\epsilon_t \tag{2.1}$$

where λ_t is the response series

ϵ_t is the noise or innovation series

$\Phi(B) = (1 - \phi_1 B)(1 - \phi_{24}B^{24})$ are the AR terms

$\Theta(B) = (1 - \theta_1 B)(1 - \theta_{12}B^{12})$ are the MA terms

B is the lag operator i.e. $\theta * B^2 X_t = \theta * X_{t-2}$ and

$\phi_1, \phi_{24}, \theta_1, \theta_{12}$ are the coefficient parameters

In addition to the past values of the response series (AR terms), and past errors (MA terms), the models also use the current and past values of other input series.

Two-stage Model

Stage 1 of the *two-stage model* has three input series:

- current and lagged hourly pool load

- a binary variable that is unity for weekend or holiday and zero otherwise
- a binary variable that is unity for a shoulder day (Fri, Mon) and zero otherwise

The effect of pool load on λ is modeled as a linear function of the current and past one hour lagged values of pool load along with the weekend and shoulder day binary variables. The mathematical form of the model with the input series is:

$$(1-\phi_1 B)(1-\phi_{24}B^{24})\lambda_t = (1-\theta_1 B)(1-\theta_{12}B^{12})\epsilon_t + (\omega_{1,0}-\omega_{1,1}B)X_{1,t} + \omega_{2,0}X_{2,t} + \omega_{3,0}X_{3,t} \tag{2.2}$$

where $X_{1,t}$ is the hourly pool load input series

$X_{2,t}$ is the weekend binary variable

$X_{3,t}$ is the shoulder day binary variable and

$\omega_{1,0}, \omega_{1,1}, \omega_{2,0}, \omega_{3,0}$ are the coefficient parameters

For estimation of model parameters in Stage 1, the model uses actual hourly load and lambda for the past seven days (168 hours). Conditional Least Squares (CLS) method is used to estimate the model parameters that minimize the error variance. CLS method assumes that the past unobserved errors are equal to zero. Further details regarding the estimation algorithm can be readily found in SAS/ETS User's Guide or any advanced text book on time series analysis. See bibliography for suggested references.

After Stage 1 estimation, a new input series is created using the model residual series (residual = actual - fitted). The new binary variable takes on a value of one whenever a calculated residual is outside $3*\sigma_{resid}$ where σ_{resid} is the standard deviation of the residual distribution. When the new binary variable thus created is a non-zero vector, i.e. some of the residuals are outside the $3*\sigma_{resid}$ range, we reestimate the model parameters in Stage 2, this time including the new input series. Otherwise, we are ready to generate forecasts, $\hat{\lambda}$ with the model from Stage 1. The idea is that any residual outside the $3*\sigma_{resid}$ band is probably due to an outlier; and the new dummy variable marks out these outliers.

The use of residual series for model diagnosis and the results from various tests are presented in Section 2.3. The Stage 2 model is similar to the one discussed in Stage 1 except for the addition of the new input series to isolate the effect of probable outliers. The mathematical form of Stage 2 model is given by equation 2.3.

$$(1-\phi_1 B)(1-\phi_{24}B^{24})\lambda_t = (1-\theta_1 B)(1-\theta_{12}B^{12})\epsilon_t + (\omega_{1,0}-\omega_{1,1}B)X_{1,t}$$
$$+\omega_{2,0}X_{2,t} + \omega_{3,0}X_{3,t} + \omega_{4,0}X_{4,t} \qquad (2.3)$$

where $X_{4,t}$ is the dummy variable indicating possible outliers and
$\omega_{4,0}$ is the coefficient parameter

One-stage Model

This model is essentially same as the Stage 1 of the *two-stage* model with the addition of another binary variable. This new binary variable takes on a value of one whenever $\lambda >$ \$35/MWH[9]. Thus the mathematical form of *one-stage* model is same as equation 2.3.

Another one-stage model was also tested where the outlier binary variable is unity whenever the hour-to-hour jump in λ is more than 5.66. This model, however, was outperformed by the above two models.

2.2.2 Generating Forecasts

To generate ex-ante forecasts of the response series, $\hat{\lambda}$, say for up to 48 hours ahead, we need to provide day ahead forecasts of the various input series to the model. The NEPEX hourly load forecasts can be used for the pool load series. The weekend and shoulder day series are deterministic and can be generated easily. The final input series which flags possible outliers is set to zero for the prediction period. The 48 hour forecasts, and forecast confidence intervals are generated using CLS method.

[9]Refer to the beginning of Section 2.2 on why \$35/MWH is chosen

Forecasts for periods greater than 48 hours can be generated but the reliability will deteriorate rapidly. On the other hand, if the parameter estimation is carried out on fewer than seven days (168 hours) of actual data, forecast accuracy can decrease as well. Now we turn to the model evaluation and diagnostic checking using various criteria.

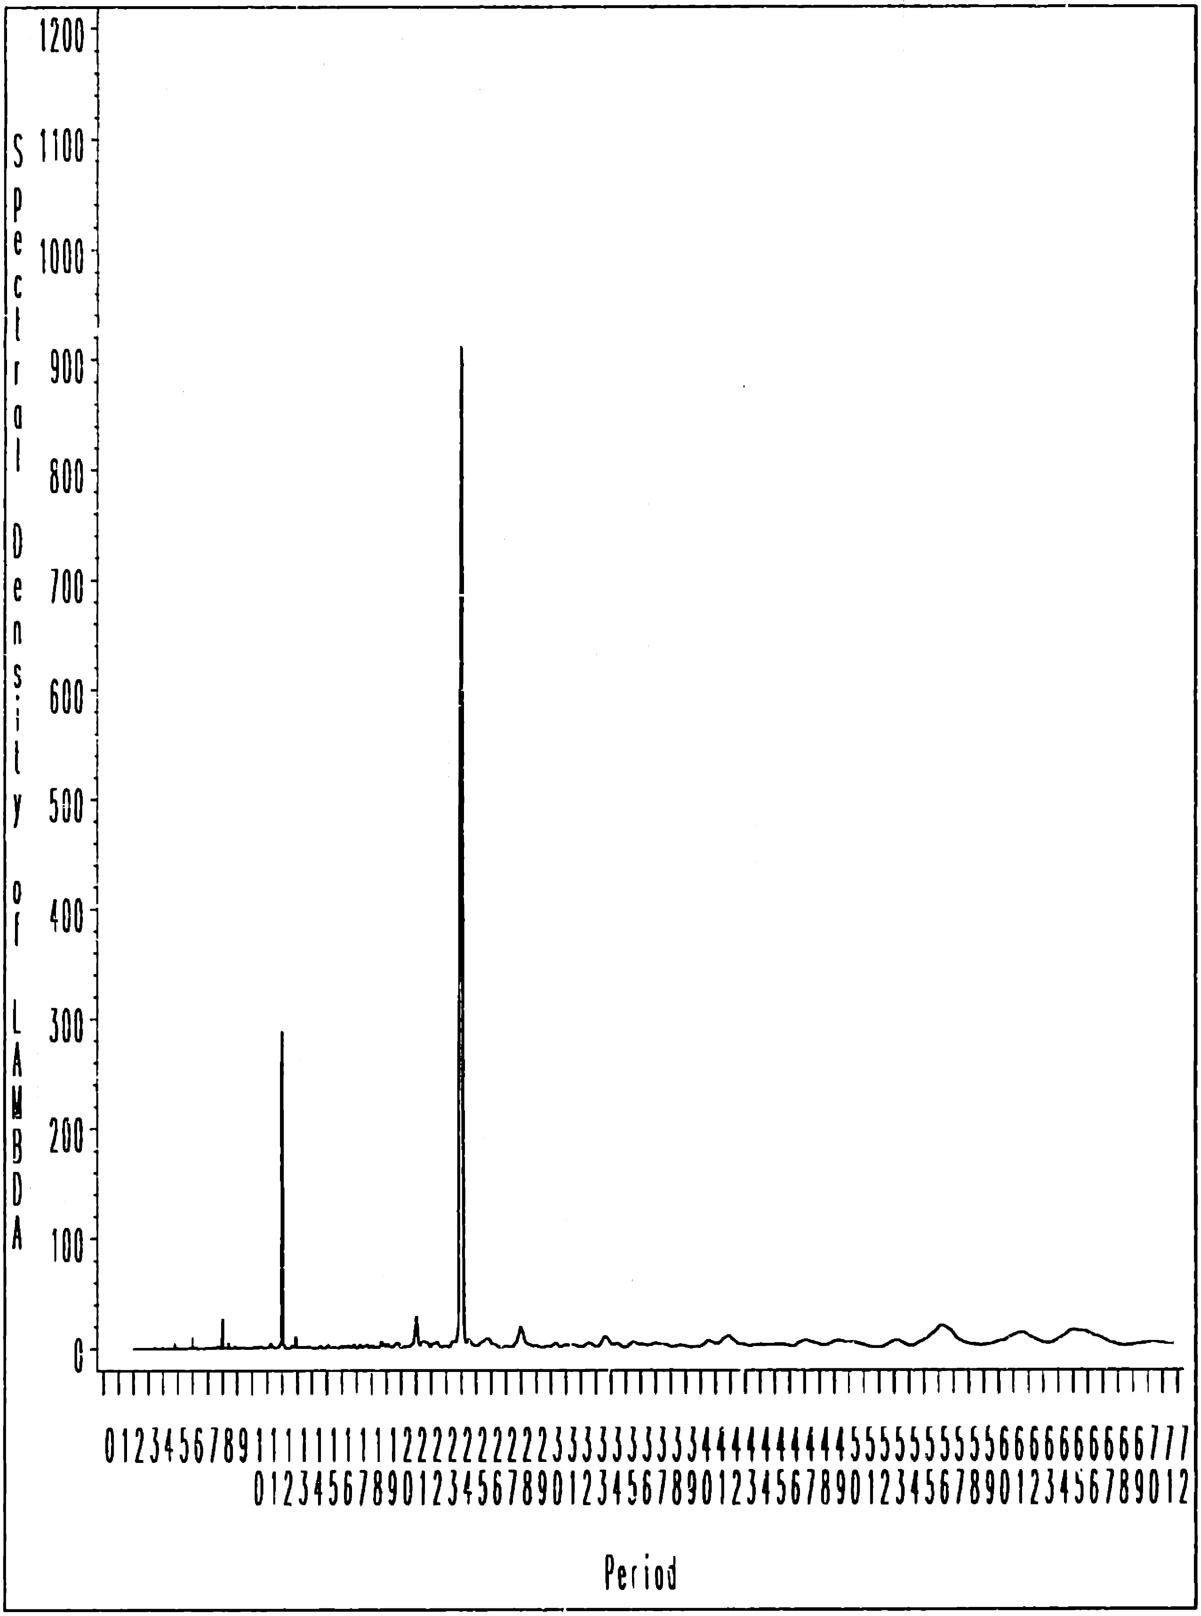

Figure 2-3: Spectrum of λ

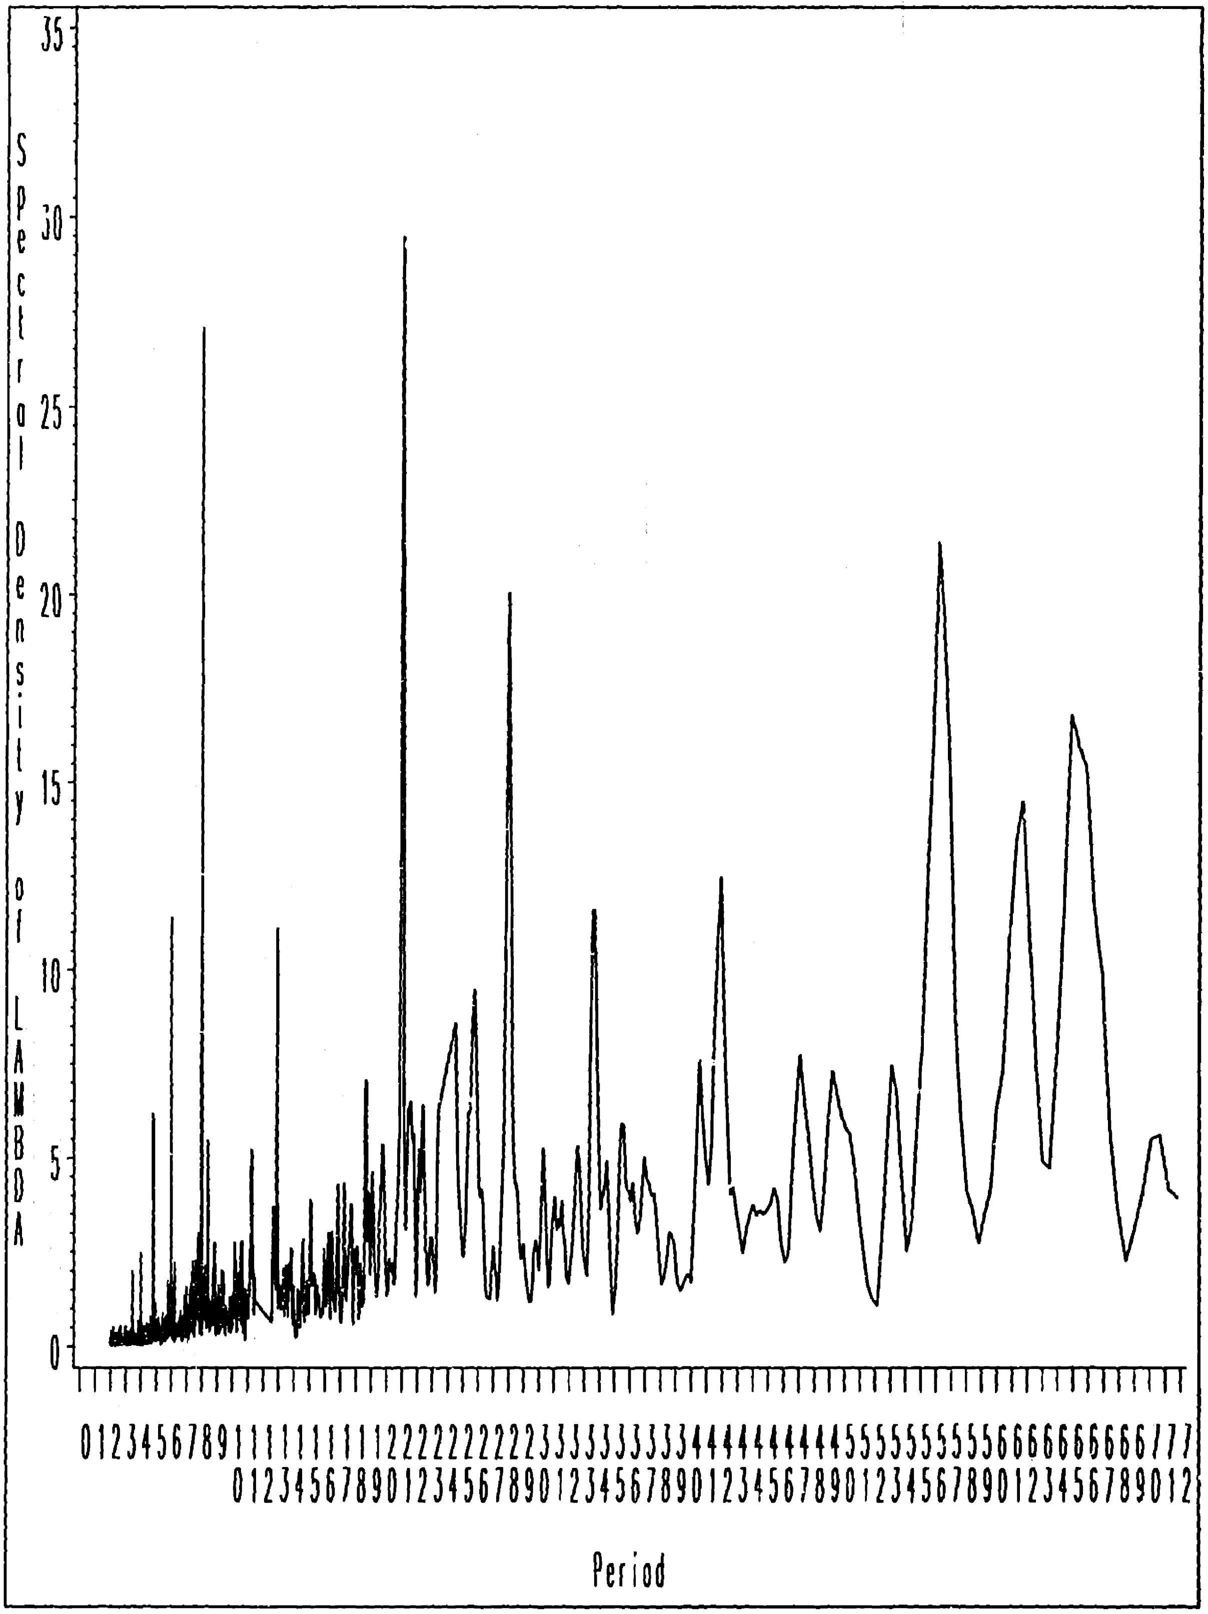

Figure 2-4: Spectrum of λ when T = 12, 24 hour periods are removed

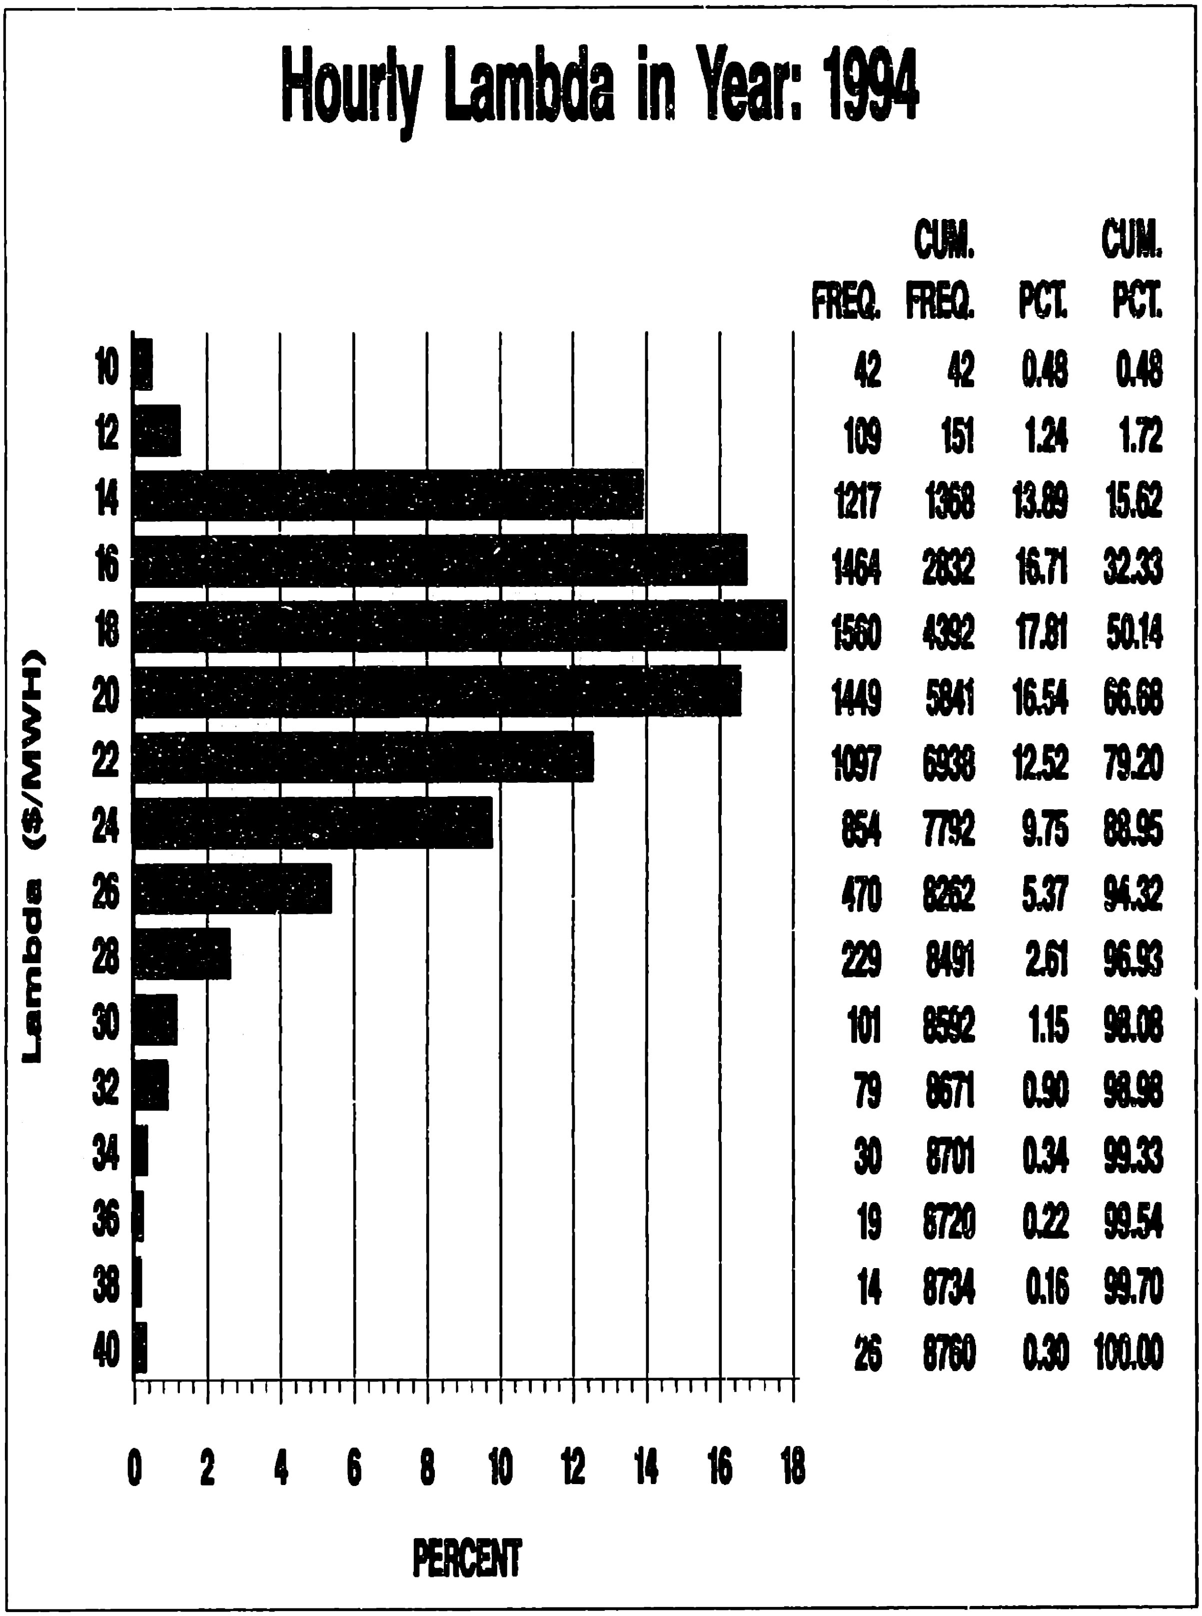

FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
42	42	0.48	0.48
109	151	1.24	1.72
1217	1368	13.89	15.62
1464	2832	16.71	32.33
1560	4392	17.81	50.14
1449	5841	16.54	66.68
1097	6938	12.52	79.20
854	7792	9.75	88.95
470	8262	5.37	94.32
229	8491	2.61	96.93
101	8592	1.15	98.08
79	8671	0.90	98.98
30	8701	0.34	99.33
19	8720	0.22	99.54
14	8734	0.16	99.70
26	8760	0.30	100.00

Figure 2-5: Distribution of of hourly λ in year: 1994

Hour-to-Hour Changes in Lambda

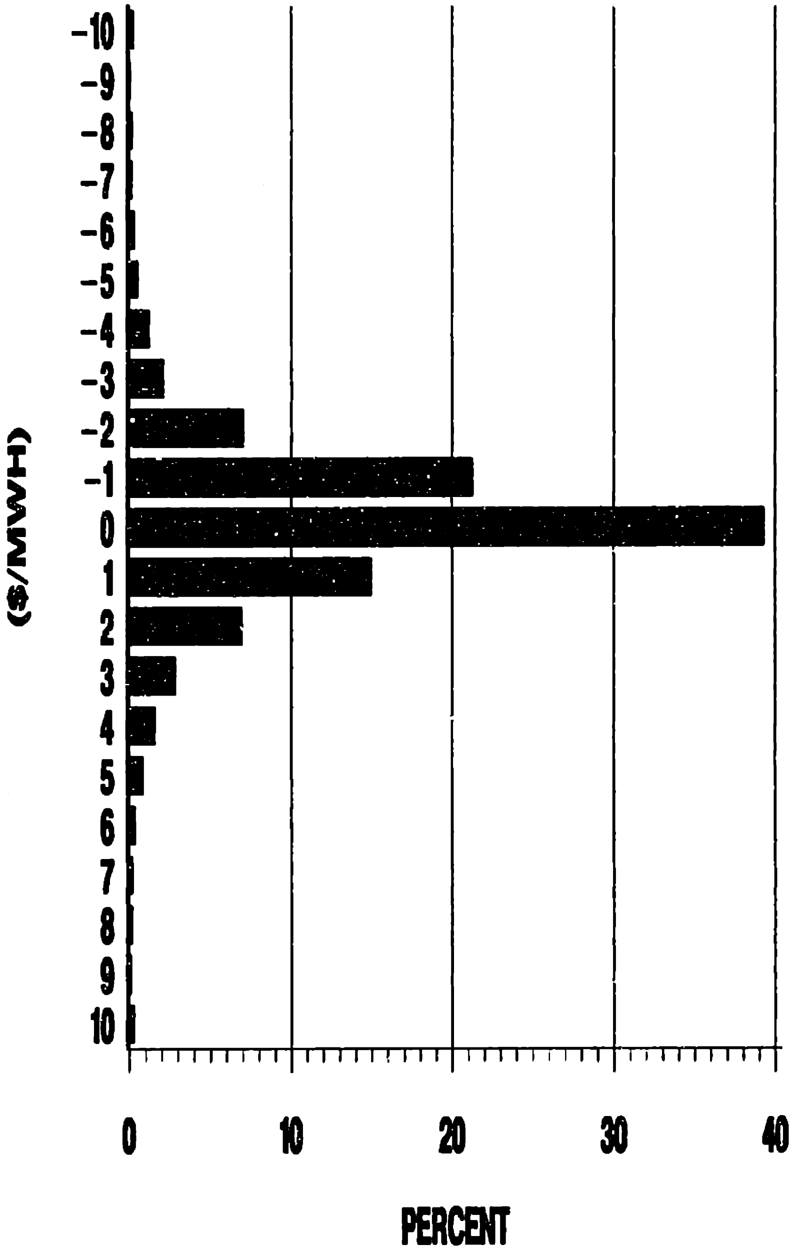

($/MWH)	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-10	20	20	0.23	0.23
-9	3	23	0.03	0.26
-8	12	35	0.14	0.40
-7	13	48	0.15	0.55
-6	25	73	0.29	0.83
-5	42	115	0.48	1.31
-4	104	219	1.19	2.50
-3	181	400	2.07	4.57
-2	612	1012	6.99	11.55
-1	1859	2871	21.22	32.78
0	3440	6311	39.27	72.05
1	1307	7618	14.92	86.97
2	605	8223	6.91	93.88
3	245	8468	2.80	96.68
4	135	8603	1.54	98.22
5	70	8673	0.80	99.02
6	30	8703	0.34	99.36
7	14	8717	0.16	99.52
8	12	8729	0.14	99.66
9	7	8736	0.08	99.74
10	23	8759	0.26	100.00

Figure 2-6: Distribution of $\Delta\lambda = \lambda_t - \lambda_{t-1}$

2.3 Diagnostic Checking and Model Testing

While model selection criteria like AIC and BIC help in determining the most appropriate values of p and q, there is still a good deal of discretion and experience of the model builder are necessary for model identification. AIC and BIC contain a penalty factor for the number of estimated parameters which discourages over-fitting. However, they do not provide guidelines on how to identify the actual AR or MA factors. Identification of these factors is done using the ACC, PAC, IAC plots, spectral analysis, t-statistics, guessing and basic understanding of the underlying process. These steps were outlined in Section 2.2 earlier.

Two other important diagnostic checks are discussed here in Sections 2.3.1 and 2.3.2. One test makes use of the the model residuals calculated as residual error = actual - fitted. The other utilizes the forecasts generated ex-post[10]. The forecast or prediction error is calculated as error = actual - (ex-post forecast) and should not to be confused with residual error. More than a diagnostic check, the ex-post forecast method is used to choose between alternative models.

2.3.1 Residual Tests

If the fitted model is adequate, the calculated residual series must be like a structureless white noise process. This implies that

1. the mean of the residuals should be close to zero,
2. the variance of the residuals should be approximately constant, and
3. the autocorrelations of the residuals should be negligible

These above conditions are tested and the results are shown in the following several diagrams. For each day of forecast in year 1994, mean and standard deviation of the model residuals are estimated. Figure 2-7 depicts the distribution of the estimated means for the *one-stage* model, summary statistics of the means series for both the models are presented in Table 2.2.

[10]forecasts of historical data that can be verified against the actuals

Moments		
	One-Stage Model	Two-Stage Model
Number of observations	357	357
Mean	-0.0027	0.0025
Standard Deviation	0.039	0.038
Skewness	0.213	0.202
Kurtosis	3.246	3.940
Quantiles (Def=5)		
Maximum	0.157	0.157
Median	-0.002	-0.003
Minimum	-0.147	-0.189
Mode	-0.037	-0.039

Table 2.2: Relevant statistics for Distribution of Means of Model Residuals

Figure 2-8 depicts the distribution of the estimated standard deviation for the *one-stage* model, Table 2.3 gives the summary statistics of the standard deviation series for both the models. *One-stage* model shows slightly better stability. As will be shown in the following Section 2.3.2, the performance is only slightly better.

The final Figure 2-9 is the autocorrelation plot of one sample residual series. It is nice to see that the residual series is almost white noise indicating a good model fit.

2.3.2 Ex-post Forecasts

The ex-post forecasting power of a model is one of the major criteria in choosing between the alternative models. Day ahead (up to 48 hours ahead) forecasts of λ are generated for 357 days in the year 1994[11]. Then the 48 hours ahead forecasts are separated into 0 to 24 hours ahead and 24 to 48 hours ahead sets. The 0 to 24 hours ahead forecasts are referred to as Day-1 forecasts, and 24 to 48 hours ahead forecasts are referred to as Day-2 forecasts.

The following Figures 2-10, 2-11, and 2-12 show the hourly plots of actual λ, predicted $\hat{\lambda}$, and pool load for three selected days in 1994 from the final model. Figure 2-10 is predicted $\hat{\lambda}$ for Day-1 or the following 24 hours while Figures 2-11,

[11] a total of 8,568 hours

Moments		
	One-Stage Model	Two-Stage Model
Number of observations	357	357
Mean	1.286	1.381
Standard Deviation	0.376	0.494
Skewness	0.854	1.004
Kurtosis	0.875	0.519
Probability the distribution is Normal	0.943	
Quantiles (Def=5)		
Maximum	2.573	2.91
Median	1.221	1.275
Minimum	0.565	0.565
Mode	2.482	0.791

Table 2.3: Relevant statistics for Distribution of Standard Deviation of Model Residuals

Prediction Power	Forecast Period			
	One-stage Model		Two-stage Model	
	Day-1	Day-2	Day-1	Day-2
% of time $\lvert\lambda - \hat{\lambda}\rvert \leq \$1/MWH$	56.20	52.27	56.10	51.84
% of time $\lvert\lambda - \hat{\lambda}\rvert \leq \$2/MWH$	76.15	72.13	76.02	71.83

Table 2.4: Prediction Power of λ Models Compared

and 2-12 show the Day-2 or day ahead predictions. Plots like these give a visual clue on the performance of various models. Further reference is made to these figures in Section 2.4.

On a macro level, one can look at the distribution of prediction errors for all the 8,568 hours. Figures 2-13 and 2-14 show such distributions for the *one-stage* model. The predictive power of both these models is compared in Table 2.3.2.

2.4 Conclusions

The λ model turned out elegant, simple and as demonstrated through application to generator dispatch in the following chapter, the prediction power is quite satisfactory. To appreciate the model performance it is important to understand its limitations as well. While some frailty can be attributed to the various data input or lack of key data, certain errors point to other advanced modeling techniques which might further improve the forecasting performance.

For example, in Figure 2-11 there is an upward drift in the pool load indicative of unstationarity, while hourly λ is showing a high variance as like in Figure 2-12. Though the trend is reflected in the day ahead forecasts, the volatility of lambda is poorly predicted. Broadly speaking, the reasons can be:

- enough information is not input to explain the variance and/or
- ARMA model's assumptions of constant variance and mean are violated

In order to better capture the volatile behavior of λ, a higher sampling rate can be tried. Since λ is available more frequently than once every hour, taking more samples per hour might provide the necessary information. However, at times λ itself may not a reliable indicator of the true market value of electricity. These issues are addressed in Section 2.4.1.

Other major problem is due to lack of hourly available capacity or unit outage data that can bear significant impact, at least transiently, on the hourly operation and hence the λ. This is the topic of discussion for Section 2.4.3. Finally, the pool load data has its own peculiarities as discussed in Section 2.4.2.

Turning to statistical properties, the bursts of high volatility can be indicative of swithching regimes or bubbles. Such variance shifts do cause problems for parameter estimation and subsequent forecasting. A class of models in time series called GARCH (Generalized Autoregressive Conditionally Heteroscedastic) models are most appropriate to handle when variance is changing with time or is a function of another time series. These are advanced models and are not tried to model lambda so far.

2.4.1 Limitations of System Lambda and Other Lambdas

In general, the hourly system lambda is accepted to represent the marginal cost of generation for that hour. However, it is necessary to know the limitations of this system λ. The limitations arise because not all the generators take part in the automatic economic dispatch run by NEPEX. In general thermal units with Automatic Generation Control (AGC) mechanism and those units marked for automatic dispatch by the pool participate in the economic dispatch. Units not marked for automatic dispatch include those undergoing various tests, and the generators that are "must run." Must run units run according to a specific schedule set by the generator's owner, usually because the owner believes it is more profitable to do so. Then there are peaking units that are manually dispatched during the peak hours of demand. These peakers are typically Internal Combustion Units (ICU) with fast start up and shut down times and run on diesel oil. When these expensive ICUs (with incremental cost as high as $100 per MWH) are manually dispatched, the system λ is not updated. This means the system λ errs to the low side.

Another situation when the λ data are unreliable is during periods of capacity deficiency when the pool is unable to carry the full level of operating reserves dictated by the reliability criteria. In that situation NEPEX follows actions laid out in what is referred to as Operation Procedure #4[12]: Actions during capacity deficiency.

The actions include running steam turbines and ICUs to their Maximum Claimed Capability (MCC), turn on emergency generation or load management resources and/or reducing the level of reserve. Hydro and pumped storage units are suspended from economic dispatch and temporarily committed to the fullest extent possible. Since the fuel cost for hydro and pumped storage is reported as zero, the current λ during OP#4 drops to *zero!*. Note that the ICUs are running at MCC at the same time. MCC is above the Reserve Claimed Capability (RCC) which is usually higher than the Normal Claimed Capability (NCC) for a generator. NCC is the capacity rating in MW where the generator usually operates most of the time, where as MCC

[12]OP#4 for short

is usually above the name plate rating and it is not recommended to operate at this level for extended periods of time. Although OP#4 arises several times a year but rarely, if ever, sheds load.

However, these may not be serious drawbacks for the applications considered in this thesis. For economical dispatch of a typical gas unit, it is important to know the cross over point when the system marginal cost exceeds the unit's marginal cost by a certain extent; beyond that it is profitable to run the unit and it is not essential to accurately know if the system λ is \$80 or \$85 per MWH. This becomes more evident in the later chapters where the applications are discussed in detail.

Several other λs do exist and they represent a snap shot in time as opposed to the one average number for the whole hour. For example, *a hourly lambda* could be the first lambda calculated every hour for each hour of the day. A very different kind of λ called "Ownload Lambda" is also calculated by each participating member of NEPOOL. Although this number is important for the operation of power pools today, it is irrelevant for the thesis purpose. The "Ownload Lambda" is a simulated cost of last incremental unit for each individual member of the pool when the member operates their own generators (using a production cost model on a computer) as if they did not belong to the NEPOOL. In the future world of hourly bidding with a market clearing price, the settlement of costs will not require the simulation of "Ownload Lambda." The hourly system lambda is expected to be highly representative of the market clearing price in a bid based system and hence the spot energy price. In fact, in the state and federal filings by electric utilities, certain justifications of expected future electricity prices are based on the historical system lambda data.

2.4.2 Limitations of Hourly Pool Load

The actual load data[13] is gathered real time by NEPEX. In the event of OP#4[14], the load data are reconstructed off-line adjusting for any load reductions. A potential limitation is that there is no understanding of if there is a time lag between the λ_t

[13]also referred to as Log 9 data in the industry
[14]described in Section 2.4.1

and load_t.

The model uses the 24 and 48 hour ahead pool load forecasts as generated by NEPEX. The forecasts are made available around 8 A.M. each day followed by infrequent updates of the day ahead (or 48-hours ahead) forecasts during the current day. The load foreacsts are generated on a "*same weather day*" principle. NEPEX collected load curves from the past three years in a book and classified them according to weather peculiarities. Up on receiving the latest weather forecast, human eyes look for similar patterns in weather in the book and recognizes a same weather day to generate a load forecast. This is how forecasts were done since several years and used for calculating the net generation requirements, economic dispatch and system operation.

The limitation is that the historical daily load forecast data are only saved as a hard copy. This data was never entered electronically and hence the ex-post model testing used the actual load data in place of forecasts. A reduction in the ex-ante forecasting power of the model can hence be expected. This increase in error reflects the inaccuracy of the hourly load forecast. Further discussion on this limitation is presented while the model applications are considered.

2.4.3 Unit Outages and Limitations of Generation Data

The outage information can be captured in many ways. One approach is to track the aggregated hourly generation by fuel type. The actual hourly generation (MWH) data are available from NEPEX directly but the problem could be with forecasting generation. Also, Phase II Hydro Quebec energy is not reflected in the generation data by fuel type because NEPEX does not consider this as generation. NEPEX, however, calculates the hourly net generation requirement as = (Forecasted load) - (tie line inflows) - (pumped storage generation).

Another approach is to closely track the availability of certain major units or aggregate capacity (MW) data by fuel type and input to the model. Although capacity and generation data are quite similar, tracking MW data could be more intuitive to those using the model because in day to day operation people think in terms of capacity (MW) lost as opposed to generation (MWH) lost. The unit availability data[15] is collected via four satellite links for billing purposes. The Annual Maintenance Schedule lists the planned outages of all major units while actual and forced outages are reported in the daily "Morning Report". Thus a outage factor called *Unavailable* MW can be calculated as = Planned outage + Forced.

The outage factor, although considered as essential, is not included in the final model because the historical data by fuel type, required for model building and testing, could not be obtained in time.

[15]also referred to as Limitation and Constraint data

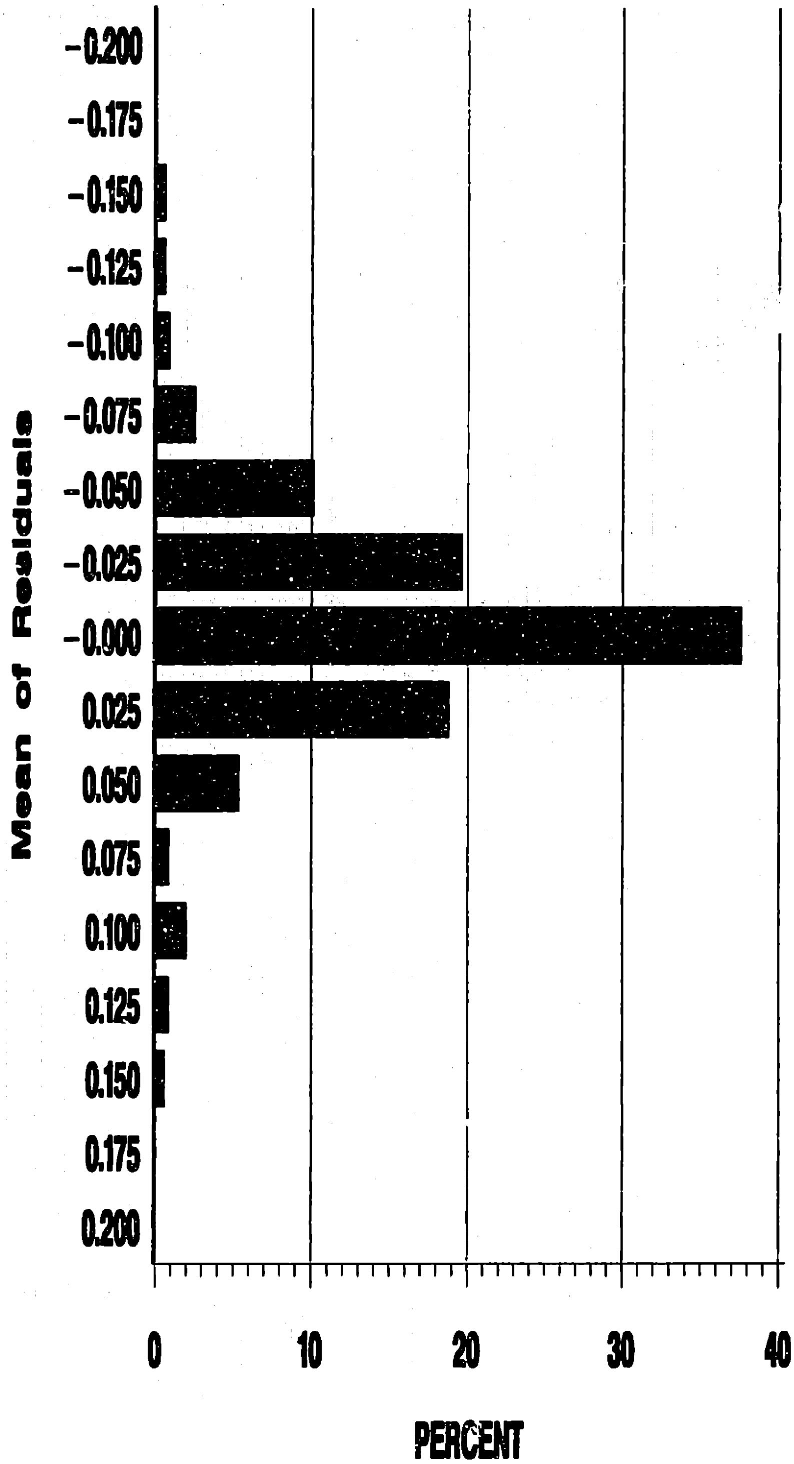

FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
0	0	0.00	0.00
0	0	0.00	0.00
2	2	0.56	0.56
2	4	0.56	1.12
3	7	0.84	1.96
9	16	2.52	4.48
36	52	10.08	14.57
70	122	19.61	34.17
134	256	37.54	71.71
67	323	18.77	90.48
19	342	5.32	95.80
3	345	0.84	96.64
7	352	1.96	98.60
3	355	0.84	99.44
2	357	0.56	100.00
0	357	0.00	100.00
0	357	0.00	100.00

Figure 2-7: Distribution of Means of Model Residuals

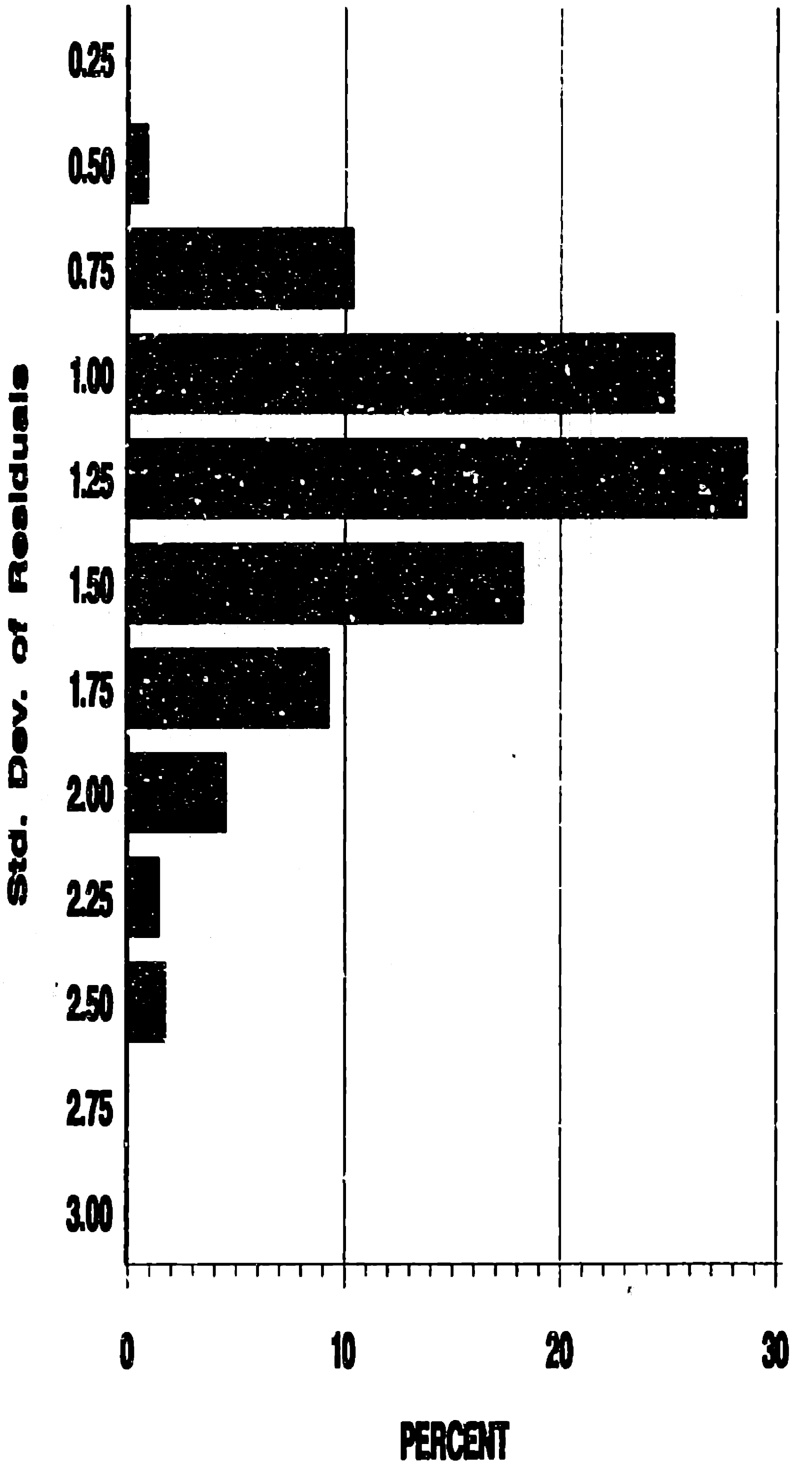

Std. Dev.	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
0.25	0	0	0.00	0.00
0.50	3	3	0.84	0.84
0.75	37	40	10.36	11.20
1.00	90	130	25.21	36.41
1.25	102	232	28.57	64.99
1.50	65	297	18.21	83.19
1.75	33	330	9.24	92.44
2.00	16	346	4.48	96.92
2.25	5	351	1.40	98.32
2.50	6	357	1.68	100.00
2.75	0	357	0.00	100.00
3.00	0	357	0.00	100.00

Figure 2-8: Distribution of Standard Deviation of Model Residuals

```
                         Autocorrelation Plot of Residuals

Lag Covariance Correlation -1 9 8 7 6 5 4 3 2 1 0 1 2 3 4 5 6 7 8 9 1          Std
  0   2.024621    1.00000  |                    |********************|           0
  1 -0.0046516   -0.00230  |                 .  |   .                |    0.077382
  2   0.112637    0.05563  |                 .  |*  .                |    0.077383
  3  -0.050946   -0.02516  |                 . *|   .                |    0.077622
  4  -0.051506   -0.02544  |                 . *|   .                |    0.077671
  5  -0.134160   -0.06626  |                 . *|   .                |    0.077721
  6   0.124186    0.06134  |                 .  |*  .                |    0.078058
  7  -0.054073   -0.02671  |                 . *|   .                |    0.078346
  8  -0.243662   -0.12035  |                 .**|   .                |    0.078401
  9  -0.089822   -0.04436  |                 . *|   .                |    0.079499
 10  -0.106056   -0.05238  |                 . *|   .                |    0.079647
 11  -0.073823   -0.03646  |                 . *|   .                |    0.079853
 12 0.00062887    0.00031  |                 .  |   .                |    0.079953
 13  -0.056083   -0.02770  |                 . *|   .                |    0.079953
 14   0.084434    0.04170  |                 .  |*  .                |    0.080011
 15  -0.014475   -0.00715  |                 .  |   .                |    0.080141
 16  -0.143524   -0.07089  |                 . *|   .                |    0.080144
 17   0.029666    0.01465  |                 .  |   .                |    0.080519
 18   0.030005    0.01482  |                 .  |   .                |    0.080535
 19 0.00039558    0.00020  |                 .  |   .                |    0.080551
 20  -0.028918   -0.01428  |                 .  |   .                |    0.080551
 21   0.021629    0.01068  |                 .  |   .                |    0.080566
 22  -0.141318   -0.06980  |                 . *|   .                |    0.080575
 23   0.036595    0.01808  |                 .  |   .                |    0.080936
 24 -0.0096132   -0.00475  |                 .  |   .                |    0.080960
                                   "." marks two standard errors
```

Figure 2-9: Autocorrlation plot of a Sample Residual Series

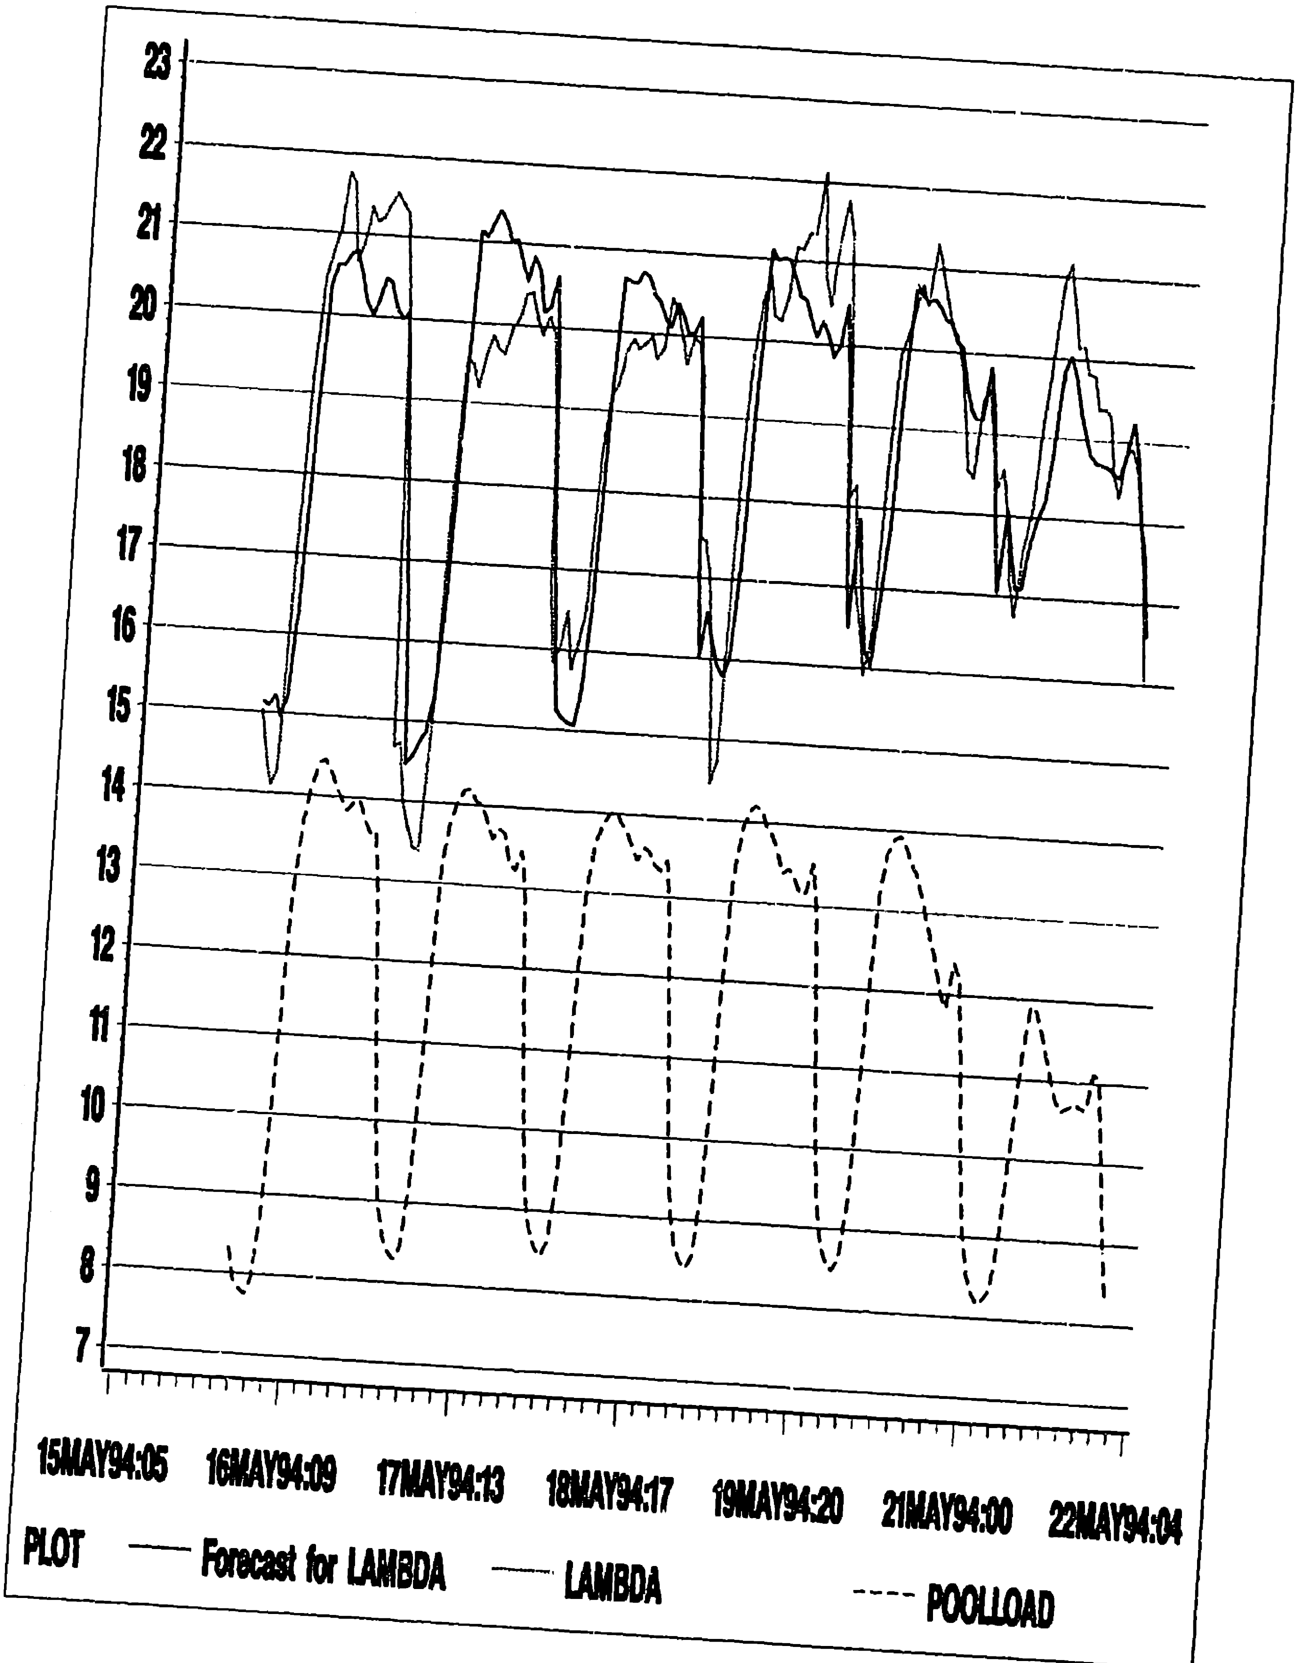

Figure 2-10: Predicted $\hat{\lambda}$, Actual λ, and Pool Load

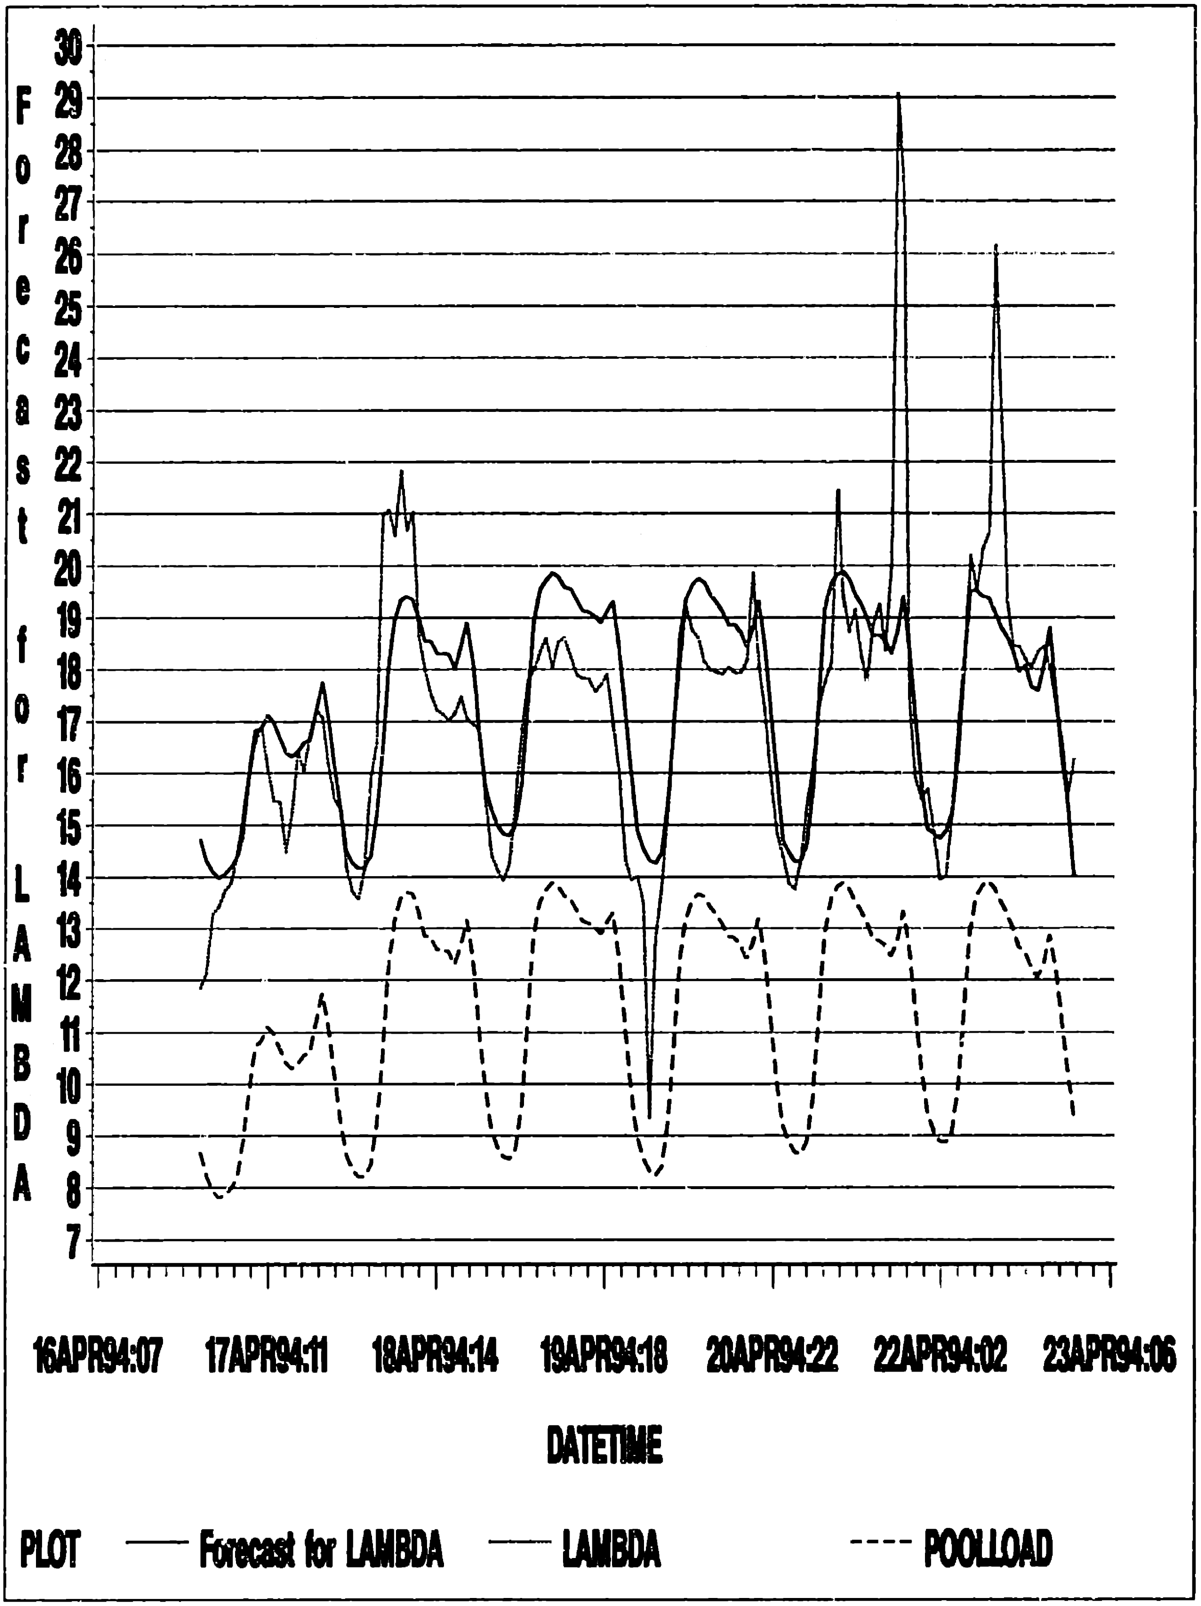

Figure 2-11: Predicted $\hat{\lambda}$, Actual λ, and Pool Load

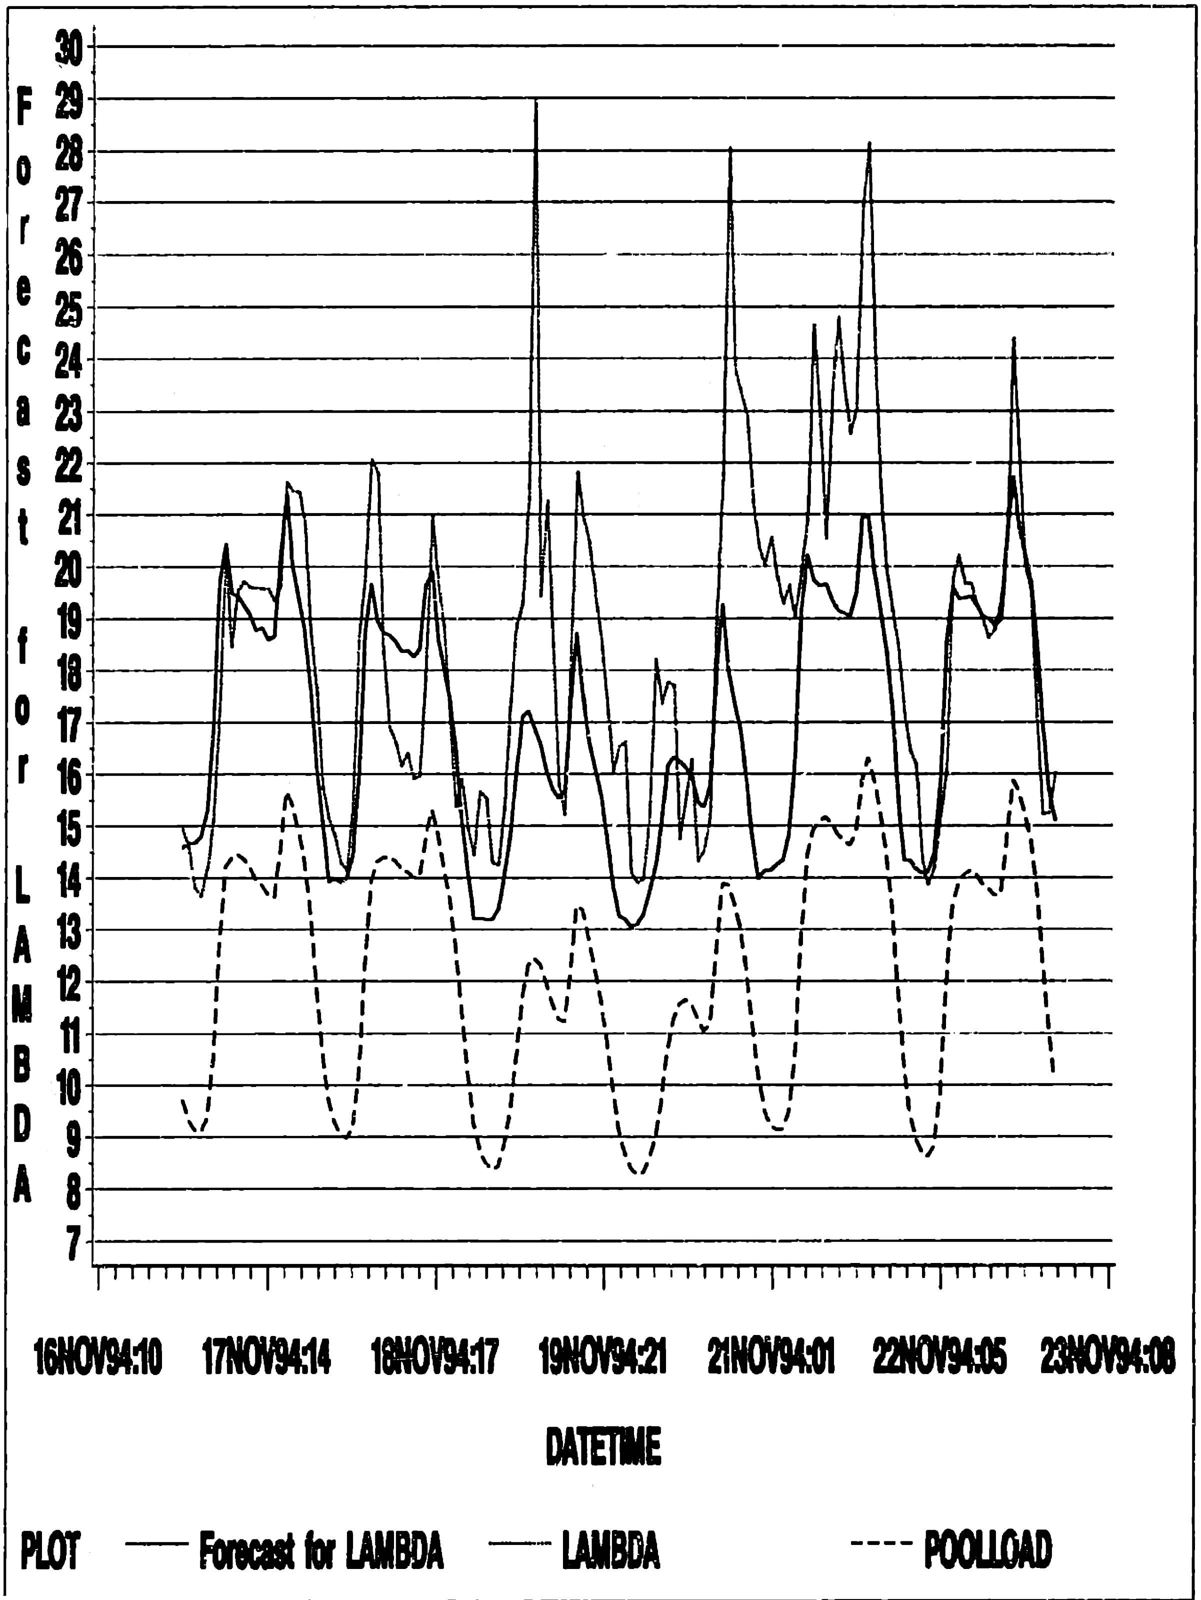

Figure 2-12: Predicted $\hat{\lambda}$, Actual λ, and Pool Load

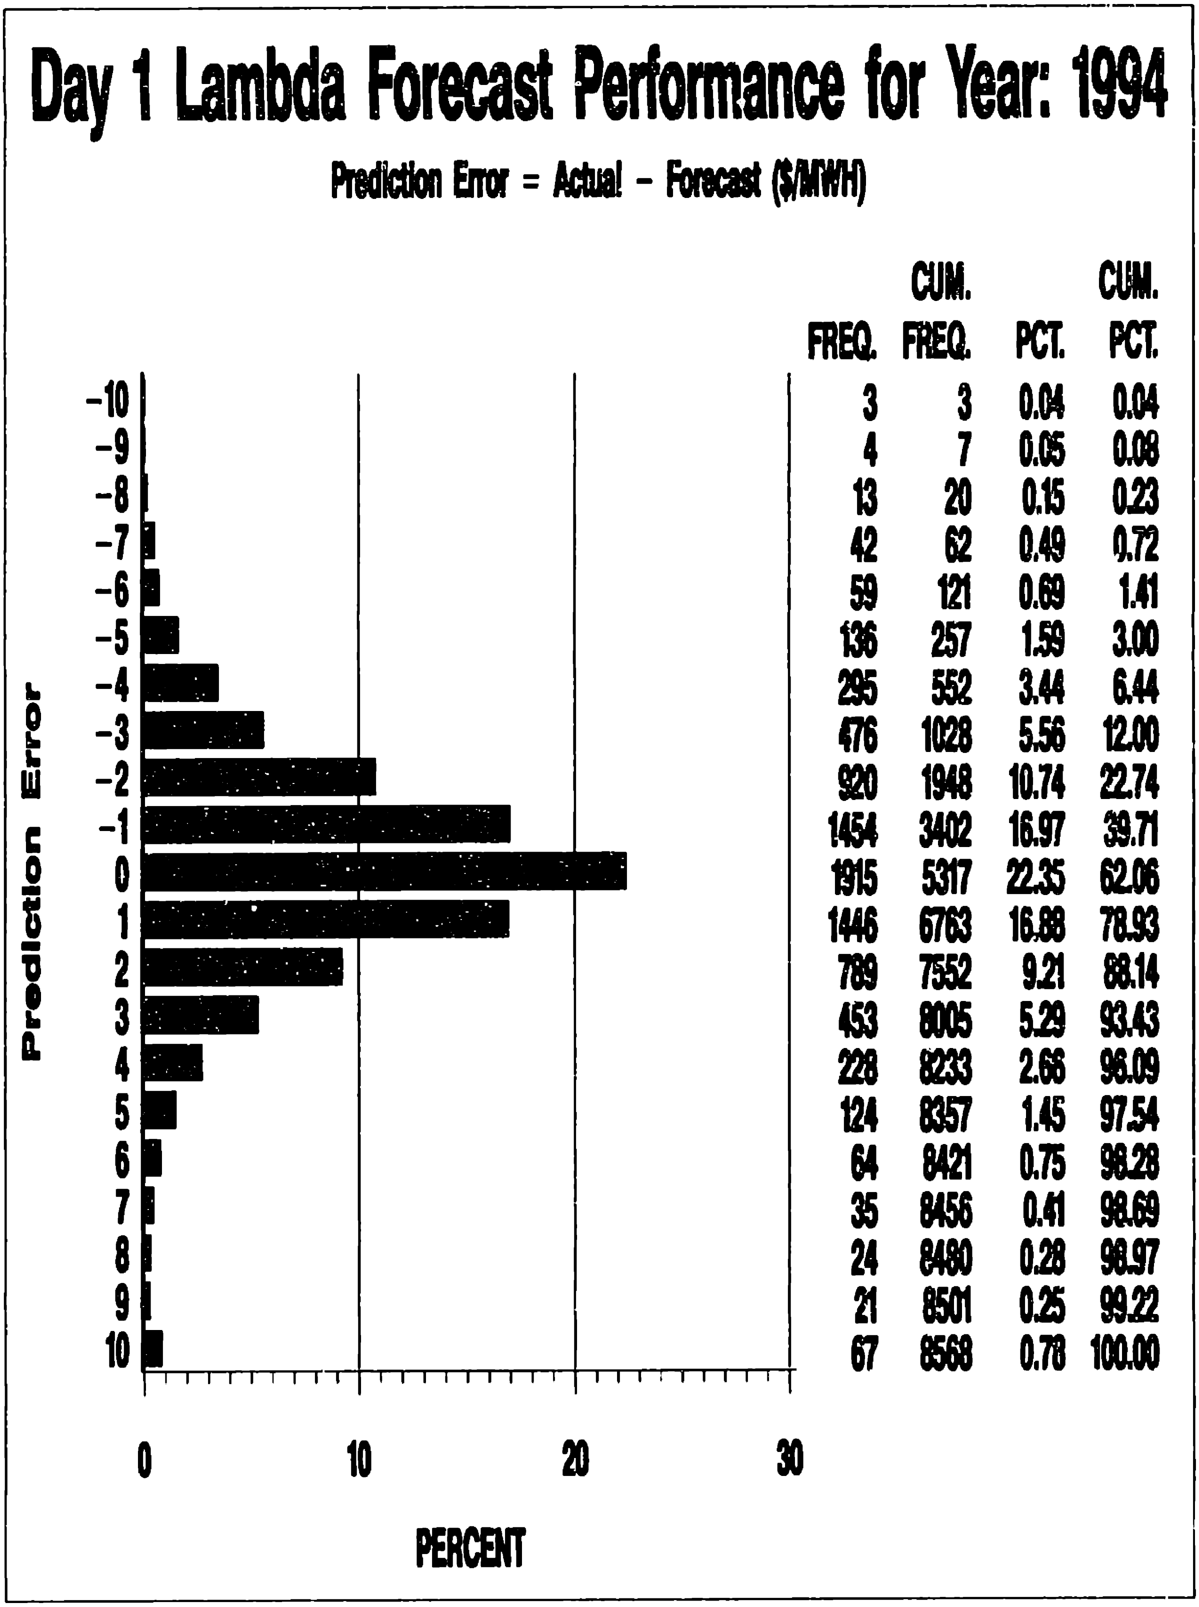

Prediction Error	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-10	3	3	0.04	0.04
-9	4	7	0.05	0.08
-8	13	20	0.15	0.23
-7	42	62	0.49	0.72
-6	59	121	0.69	1.41
-5	136	257	1.59	3.00
-4	295	552	3.44	6.44
-3	476	1028	5.56	12.00
-2	920	1948	10.74	22.74
-1	1454	3402	16.97	39.71
0	1915	5317	22.35	62.06
1	1446	6763	16.88	78.93
2	789	7552	9.21	88.14
3	453	8005	5.29	93.43
4	228	8233	2.66	96.09
5	124	8357	1.45	97.54
6	64	8421	0.75	98.28
7	35	8456	0.41	98.69
8	24	8480	0.28	98.97
9	21	8501	0.25	99.22
10	67	8568	0.78	100.00

Figure 2-13: Distribution of Prediction Errors For Day-1

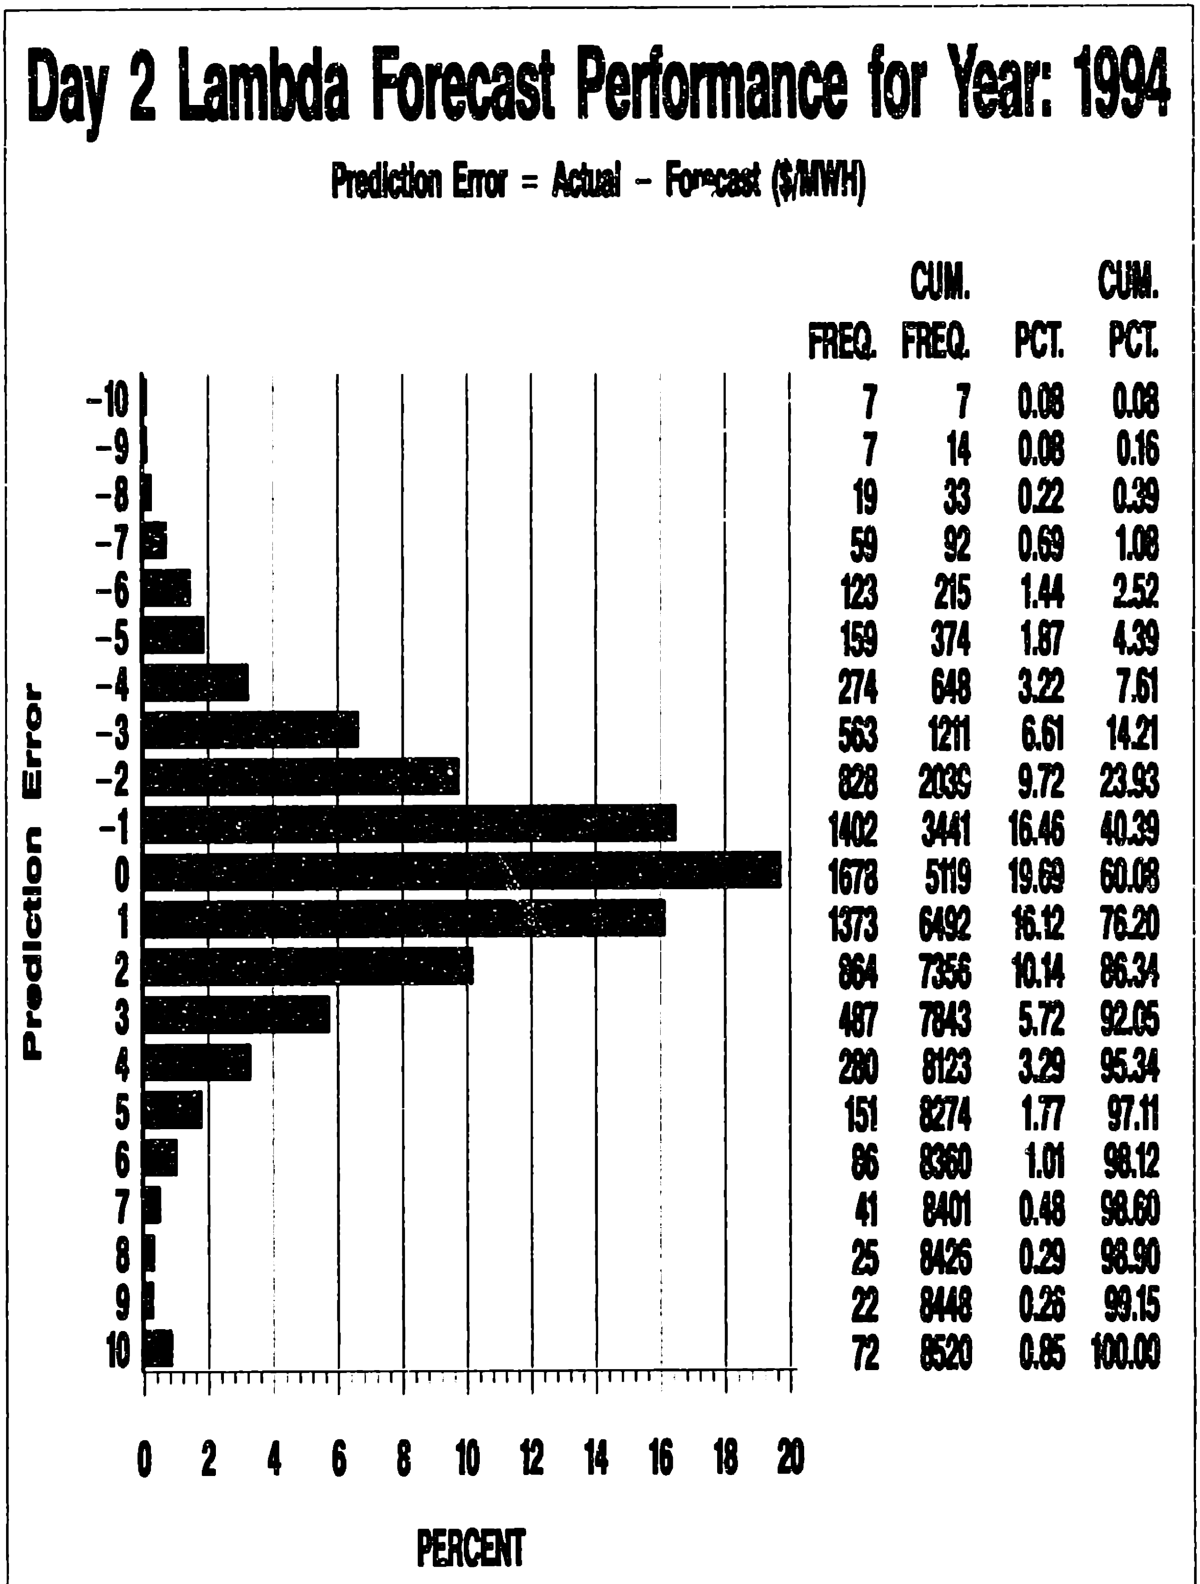

FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
7	7	0.08	0.08
7	14	0.08	0.16
19	33	0.22	0.39
59	92	0.69	1.08
123	215	1.44	2.52
159	374	1.87	4.39
274	648	3.22	7.61
563	1211	6.61	14.21
828	2039	9.72	23.93
1402	3441	16.46	40.39
1678	5119	19.69	60.08
1373	6492	16.12	76.20
864	7356	10.14	86.34
487	7843	5.72	92.05
280	8123	3.29	95.34
151	8274	1.77	97.11
86	8360	1.01	98.12
41	8401	0.48	98.60
25	8426	0.29	98.90
22	8448	0.26	99.15
72	8520	0.85	100.00

Figure 2-14: Distribution of Prediction Errors For Day-2

Chapter 3

Generator Dispatch and Fuel Exposure Management with a Single Dispatch Price

Fuel exposure results when there is a discrepancy between the gas nominated for burns ahead of time and the actual usage during the day. Then Sections 1.2 and 1.3 discussed the implications of the bid based unit dispatch system and the possibility that only one bid price may be allowed per generator. With a single dispatch price, any resulting fuel exposure has to be off-set by transacting in the intra day gas market unlike in the current system[1] with additional options to adjust the generator dispatch to match fuel availability.

In this chapter, the forecasts generated by the λ model are used to project the hourly dispatch of a gas unit. The generator dispatch algorithm which uses only one dispatch price per generator is developed in Section 3.1. Section 3.4 immediately builds on the analysis to estimate the amount of day ahead fuel usage and develops a fuel exposure management strategy.

The bigger goal here is to maximize the overall profitability of the BTU converter. Managing the fuel exposure alone with out consideration to the economic value of us-

[1]The current strategies and the resulting performance in managing fuel exposure are reviewed in Section 1.4.

ing the gas available at hand may result in lower profits. A comparison is to be made between the expected payoffs in the electricity and gas markets. The analysis presented in Section 3.2 suggests that there may have been profit making opportunities that were not exploited by the current methods of managing fuel exposure. The final section in this chapter discusses the problems with various data available for testing, benchmarking the strategy, and limitations of the strategy itself.

3.1 Economic Dispatch of a generator

Forecasts of hourly λ are used to project unit commitment based on the heuristic "commit when pool λ is greater than the marginal cost of generation[2]." Of course various operational constraints specified in the unit's NX-12A[3] need to be met. The unit commitment algorithm takes into account the various operational constraints and identifies the most profitable way to run the generating unit.

Recall that λ (\$/MWH) represents the market value of electricity at this moment in time while the quoted dispatch price (\$/MBtu) indicates the value of gas at hand. The dispatch price multiplied by the unit's heat rate (MBtu/MWH) is the dispatch cost (\$/MWH) for the unit; the spark spread (\$/MWH) is the difference between λ and dispatch cost. When ever the spark spread is positive, the BTU converter is making a profit and the above heuristic would commit the unit for generation.

These considerations are incorporated into the economic dispatch[4] algorithm through the equations given below. Unit characteristics such as block heat rates, operational constraints like minimum run/shut-down times, etc. are input to the algorithm along with the hourly $\hat{\lambda}$ to project the optimal hourly dispatch for the unit.

- at all times, the unit to be dispatched is assumed to be either generating (in on-state) or not generating (in off-state).

[2]The assumption is that hourly pool λ is independent of the dispatch of the single unit under consideration.

[3]See Section 1.5 for definition.

[4]**ED** for short

- the unit earnings are calculated at the market clearing price or λ at that instant. The earnings in the power market are then $\Sigma\lambda * \text{MWH}$.

- at each hour, dispatch cost[5] is compared against hourly $\hat{\lambda}$ to decide either to continue in the current state or change state i.e. keep running when ($\hat{\lambda}$ - dispatch cost) ≥ 0 or continue in off-state when ($\hat{\lambda}$ - dispatch cost) < 0.

- when a change of state is advised at any hour because it is economical, first the minimum run/down times are checked. If the run/down time requirements are satisfied, the economics of changing state is evaluated over the *minimum* run/down time. The unit will be committed if the profit/loss function is positive for that period or else the unit is recommended for shut down.

$$\text{profit/loss}^6 = \Sigma_{j=1}^{3}(\hat{\lambda}_{t+j} - \text{dispcost}_{t+j})$$

- the algorithm assumes no start-up cost for unit dispatch. This cost is recovered at the end from unit earnings for all the start-ups incurred. Typical cost per start-up includes a fuel expense of ≈400 MBtu and a fixed $ amount of ≈ $5,000.

- when recommended for shut down or going to off-state, the losses incurred to keep the unit operational at minimum-load until the next on-state are compared with start-up cost. The unit is marked for shut down only when the start-up cost is less than the accumulated losses of unit running at minimum load; otherwise the unit is committed for minimum load operation.

- the unit reaches full load level in no time i.e. unit ramp up/down times[7] are ignored for simplicity.

- Summer capacity rating in effect from June 1 to September 30 and a higher Winter capacity rating rest of the time.

---

[5]dispatch cost ($/MWH) = FLAHR (MBtu/MWH) * Fuel Price ($/MBtu).
[6]Here, the minimum run time = minimum down time = 4 hours
[7]typical ramp up/down times are 6 MW/minute

The affects of simplifying assumptions on unit commitment are easy to understand and is the topic of discussion for Section 3.5.2. Albeit some of the ideal assumptions, the ED algorithm is realistic and becomes a powerful tool for fuel nomination and performance evaluation. First the algorithm is used to simulate dispatch of a gas unit ex-post using actual λ. This back-testing accomplishes two important tasks. One, it establishes the accuracy of the algorithm. Two, on the assumption that the ED simulation approximates economic optimum, it provides a basis to evaluate the performance of current strategies and determine any potential lost profits or missed arbitrage opportunities. Any competitive firm would want to determine whether such opportunity existed and how best to exploit it. These issues are discussed in Section 3.2.

The day ahead unit commitment can be projected when the ED algorithm is run with predicted $\hat{\lambda}$ as input. The forecasting power of the spot energy price model and unit commitment algorithm combination is established in Section 3.3. The ex-ante forecast of unit commitment enables to estimate the fuel usage and hence ways to manage fuel exposure.

3.2 Economic performance of current strategies

The ED algorithm is run using the actual hourly λ to project hourly dispatch of a gas unit over 357 days of 1994. Figure 3-1 is a plot of discrepancy between the actual and simulated ED.

The upper half or negative values result when the unit is committed by the algorithm but there is no operation in actual (or unit generating at a lower load) during that time. The lower half or positive values result when the unit is not committed in the simulated dispatch but there is actual generation.

The error extremities at 150,165, -150 and -165 bars in the Figure 3-1, accounting for approximately 15% of the total error are more interesting. The error bars at -165 and -150 result when the unit is fully committed, up to the Winter or a lower Summer rating, in the simulated dispatch where as there is no corresponding operation in actual during that time. The -105 error bar (≈8%) is due to unit commitment at minimum load in the simulated dispatch, while the unit is shut down in actual. Similarly, the -60, and -45 error bars result when the unit is fully committed in the simulated dispatch when the actual operation is at minimum load.

The error of the -15 MW bar in the chart (≈11%) can be explained by actual generator operation at full load, to the level permitted by the ambient temperature, below the rated capacity. The small errors, between the extremities and zero-error bar, are largely due to the deviation of simulation output (discrete with no ramp-times) from real output which is continuous anywhere from 0 to full-load. The model does switch the rated capacity to a Winter rating during Winter, and de-rate during Summer; further sensitivity to ambient temperature changes are ignored.

Most of the discrepancy in Figure 3-1 is due to undercommitment and might indicate missed profit opportunities. In particular, the generator tends to be under-committed by over 235,000 MWH. Table 3.1 presents the cash flows, and start-ups accumulated over the 8352 hours of operation. The cash flow calculations are based on the following assumptions:

- Actual hourly λ is taken as the market price of electricity for revenue calculations. Revenue = $\Sigma(\lambda * \text{MWH})$.
- Cost of generation or fuel cost = $\Sigma(\text{FLAHR}$[8] $* \$/\text{MBtu} * \text{MWH})$.
- The recovery cost per start-up is approximated to \$5,000 + (400 MBtu * 2.25 \$/MBtu) = \$5,900

Table 3.1 is indicating that the actual dispatch may be ignoring opportunities for further profit from what appears to be a tendency not to dispatch the unit as much as would be warranted by the observed spark spread. Any profit maximizing firm would be interested to find out if the actual operation is sub-optimal compared to the simulated ED.

[8]Full Load Average Heat Rate in MBtu/MWH

Dispatch	Revenue($)	Cost($)	Start-ups	Start-up($)	Profit($)
Simulated	15,371,608	13,157,245	51	300,900	1,913,463
Actual	10,640,407	9,025,738	117	690,300	924,369
Simulated - Actual	4,731,201	4,131,507	-67	-389,400	989,094

Table 3.1: Cash Flows for Actual and Simulated Unit Dispatch

Figure 1-4 in Section 1.4 represents the fuel exposure as a result of the ad hoc methods now used to deal with the problem. The current analysis is pointing out that simply trying to match gas flows alone does not necessarily result in the best overall economics. The gas nomination methods should systematically consider the estimated or observed spark spread.

There are arguments to debate if the potential lost profits are result of pure sub-optimality or mostly noise and non-existent. The reasons in favor of sub-optimal performance are given first succeeded by those favoring unreal profits. Unfortunately, several of the arguments are unsubstantiated since relevant data necessary to quantify the issues are not available; a discussion on these data limitations is differed until Section 3.5.1.

The current system allows the generator owner to change the dispatch price to a higher or lower price intra day. This option allows one to control the unit dispatch in order to reduce the fuel exposure and take advantage of any arbitrage opportunities in the intra day gas and electricity markets. From figure 1-4 and relevant discussion in Section 1.4, the electric company enjoyed an embedded option to carry on relatively large gas flow imbalances. When the pipeline operators are lax to allow daily imbalances, even with unfavorable intra day market conditions, there is less need to force dispatch to match gas at hand. Thus more often than not, the dispatch price is switched in actual to take advantage of arbitrage opportunities and realize a higher profit. The simulated dispatch, on the other hand, does not allow such option to switch dispatch price. Thus one can expect the simulated dispatch would be less profitable than the actual instead of the other way round.

Another peculiarity of the discrepancies in dispatch is that they tend to occur in bunches. This effect would be more striking when looking at the daily error plots as in Figure 3-2 generated by summing up the hourly data of Figure 3-1. For a unit of ≈170 MW capacity, a discrepancy of 3,500 MW is roughly 20 hours of operation in one day. It is hard to brush away such deviations as due to simplification of physical constraints such as ignoring ramp-times in the ED algorithm. It could be insightful to study the causes for such behavior since the implications can suggest improvements in real-time dispatch operation.

Finally, the period over which start-up cost is recovered (or amortized) can greatly influence the unit commitment. When using ED algorithm to simulate unit dispatch, the start-up costs are deducted in the end as shown in Table 3.1. Since shorter periods[9] over which start-up costs are recovered would reduce the number of hours of operation (and start-ups), there can be a reduction in the magnitude of undercommitment. However, changing the cost recovery criteria is likely to increase potential profits while reducing disparity in the energy generated.

Turning to counter arguments, one can propose that the simulated dispatch is done ex-post and hence superior performance is not a surprise. Irrespective of the applicability of this argument[10] here, it is valuable to identify the reasons for such superior performance in the back tested strategy. An argument can be made that in real-time, one does not have the perfect foresight to trade-off start-up cost Vs. losses incurred for minimum load operation. The ED algorithm is run with out such a trade-off and the results still indicate undercommitment by over 112,000 MWH and corresponding potential lost profits amount to $270,720.

Another reason for undercommitment could be that in actual operation, the unit loading level is dependent on the ambient temperature and at times may run below its rated capacity. From figure 3-1, most of this discrepancy is indicated by the -15 MW bar in the chart – undernomination by ≈14,000 MWH of energy. That still leaves the ≈235,000 MWH of undercommitted energy unattributed to any simplifying assumptions.

[9]The current NEPOOL policy regarding the horizon of interest is given in Section 3.5.2.
[10]via. the popular stock phrase "past performance is no guarantee for future results."

The aggregation of all real time lambdas during the hour to one number may smooth out minute variations in the generator output due to unit ramp-times, temperature sensitivities etc. Certain events[11] in the actual operation may not be fully reflected by the hourly λ which can cause the actual dispatch to deviate from optimal ED. Since the real system operation can only near simulation efficiency, some of the potential profits can be considered as cost of real business operation.

Overall, there are strong reasons to conclude that the ED algorithm approximates the economic optimum and thus provide a more tractable basis for comparison. In particular, the generator tends to be undercommitted by about 235,000 MWH and the missed profit opportunities amount to about $1 Million.

3.3 Projecting the Unit Commitment

The unit commitment algorithm is used to project the hourly dispatch of a gas unit over 357 days of 1994. Figure 3-3 shows the distribution of error between the optimal (simulated) and the forecasted hourly ED for day ahead hours t+25 to t+48. The plot highlights the day ahead forecasting power of the application ≈83% as given by the 0 error bar.

Other error bars are due to forecast errors for unit operation at full-load, minimum-load or no-load. These are explained earlier in Section 3.2. The assumption of zero ramp-times and operation of a single unit in isolation from the pool imparts a highly modal structure to the error distribution.

Figure 3-4 shows the distribution of error between the day ahead actual and forecasted hourly ED. From the analysis presented in Section 3.2 and by comparison with Figure 3-1, some of the error in Figure 3-4 can be attributed to reasons, other than forecasting accuracy, such as sub-optimal actual performance.

One can also look at forecast performance for a closer time period than day ahead. Figures 3-5 shows the distribution of error between the optimal (simulated) and forecasted operation for up to 24 hours ahead (hours t+1 to t+24). As expected, the

[11] such as OP#4 discussed in Section 2.4

prediction power improved to ≈85%. However, going into hours much close to the actual, importance of hourly generation data greatly increases. Thus lacking any generation data input to the model, looking ahead just few hours from the actual may not be reliable. A discussion of such issues was presented in Section 2.4.

The following Section 3.4 expands the analysis from unit commitment to forecasting day ahead gas usage. Also implications for the gas nomination process and fuel exposure impact are discussed and improvement over the current ad hoc methods is demonstrated.

3.4 Forecasting Day Ahead Gas Usage

Once the day ahead unit operation is forecasted as above, it is straightforward to estimate the amount of day ahead gas usage for nomination. The amount of fuel burned in any hour is ≈AHR * MWH where AHR is the average heat rate[12] in MBtu/MWH for that level of loading.

Plots of error distribution between the *hourly* optimal gas usage and forecasted gas usage will be identical to Plots 3-3 or 3-4 with a change of scale from MW to MBtu. Since the inter day gas market operates more on a daily as opposed to hourly basis, the daily gas usage is of more interest. This daily usage can be aggregated from the hourly data. Plot 3-6 shows the distribution of error between the optimal (simulated) and predicted gas usage for day ahead[13].

If Day-2 forecast is the quantity nominated for day ahead usage, then Figure 3-6 represents the fuel exposure for the whole year 1994. By comparison with Figure 1-4 in Chapter 1, there is not significantly greater fuel exposure here than the final result of more adaptive control methods now used to deal with this problem. With out the option to switch dispatch prices intra day, all the fuel exposure has to be off-set by transacting in the intra day gas market or by utilizing the embedded options to cover resulting gas imbalances.

[12]The block heat rates are reported in the unit's NX-12A.
[13]The plots are shown for electric day instead of 8 a.m. to 8 a.m. gas day. For actual implementation, 24 hour running totals can be used to match the gas day.

Month	Prediction Error (MBtu)
February	7,698
March	5,728
April	69,350
May	68,314
June	-23,822
July	-52,879
August	-111,896
September	-9,373
October	-5,584
November	-42,998

Table 3.2: Prediction Error by Month

The day ahead prediction error is between ±5% and ±10% of the initial day ahead nomination quantity roughly 60% and 75% of the time respectively. Specifically, ≈17% of the time, more than 5% of the initial nominated quantity is burned by the (simulated) optimal dispatch, and ≈22% of the time, more than 5% of the initial nominated quantity is left unused. Figure 3-7 shows the distribution of these percentages.

The prediction error or the daily under/over nomination error sequence showed no correlation with day-to-day changes in the dispatch price. However, this does not eliminate the possibility of correlation with convenience spreads (or the inter to intra day price change). Table 3.2 presents the aggregated prediction error by month and is indicating some autocorrelation among the nomination error. The is confirmed by the estimated ACC of ≈30%. The following Section 3.4.1 quantifies the impact due to such exposures.

Section 1.5.2 formulated the gas market structure in Table 1.3. If actual intra day transaction data are available, one can readily calculate the convenience spread and find out exactly if each day is a penalty or a discount day. Lacking such actual data, one can consider various scenarios to grasp the possible range of fuel exposure impact by making assumptions about the bid/ask convenience spread.

3.4.1 Assessing the Impact of Fuel Exposure

The problem is tackled as three scenarios. The first scenario estimates the possible bounds for fuel exposure impact by using statistics estimated from the available actual intra day transaction data. Then recognizing the forward correlation between today's intra day and tomorrow's inter day prices, the second and third scenarios make a more realistic assessment of the fuel exposure impact.

Actual intra day transaction data i.e. the quantity and price of intra day gas are obtained for a total of 52 days over a six month period. A volume weighted price index is constructed from the individual transactions on each day. This intra day index has a mean = \$3.51/MBtu and a σ = \$2.3/MBtu. Of more interest are the characteristics of convenience spreads. The convenience spread is calculated as the difference between the day ahead gas index for a pipeline close to the plant and the intra day index; convenience spread = intra day index - pipeline gas index. This inter to intra day price change series has a mean = \$0.299/MBtu and a σ = \$1.219/MBtu. While the volume weighted average spread is = \$0.119/MBtu.

One can also estimate the bid and ask spreads by separating the intra day transaction data into purchases and sales. Such calculation resulted in a volume weighted bid spread[14] = \$-0.243/MBtu while the volume weighted ask(ed) spread[15] is = \$0.366/MBtu. Two observations can be made from these assymetric spreads. The negative ask(ed) spread indicates that on average there are more discount days than penalty days due to overnomination. While the positive bid spread is indicative of more penalty days than discount days on average due to undernomination. The assymetry in magnitude (\$0.24 Vs. \$0.36) can be due to reasons like price impact.

Using the above data it is possible to estimate bounds for fuel exposure impact. The worst case possibility is when all of the off-setting transactions are against the firm i.e. all days are penalty days. The best possibility then, is when all of the off-setting transactions are in favor i.e. all days are discount days.

Accumulated from the daily prediction errors, there is a total undernomination of

[14]bid spread = bid price for gas - pipeline gas index
[15]ask(ed) spread = ask(ed)price for gas - pipeline gas index

Market Structure	Bid Spread ($)	Ask(ed) Spread ($)	Fuel Exposure due to Undernomination ($)	Fuel Exposure due to Overnomination ($)
Symmetric	0.15	0.15	67,948	97,637
Symmetric	0.25	0.25	113,264	162,729
Asymmetric	0.10	0.20	45,305	130,183
Asymmetric	0.35	0.25	158,570	162,729

Table 3.3: Fuel Exposure Impact when All Gas Flow Imbalances are Off-set through Intra Day Trading

453,056 MBtu and overnomination of 650,915 MBtu of gas over the year 1994. If the convenience spread is $0.10/MBtu[16], and say the bid and ask spreads are symmetric then the resulting cash flows due to the fuel exposure are presented in the first row of Table 3.3. The cash flows for other convenience spread structures are also presented in the table. In the worst case, a cash flow equal to the sum of last two columns for each row can be incurred as cost due to the intra day market transactions. In the best case, the same total cash flow can be additional profit. Realistically, the potential cash flow due to the exposure can be anywhere between these two extremes.

Sometimes, the pipeline companies extend embedded options that allow over/under drawing of gas up to a certain percentage of the initial nominated quantity. If 10% slack is allowed, then the total under and overnominations accumulated from the daily prediction errors are 374,657 and 559,228 MBtu respectively over the year 1994. The impact of this fuel exposure is estimated in Table 3.4.

In order to estimate a more likely impact that lies between the above bounds, one has to be able to nail down each day as either a penalty or discount day. This is possible with the assumption that today's intra day price can strongly influence tomorrow's inter day price. Since the inter day dispatch prices are available for all of the days in the test year, it is possible to determine, ex-post, the leading change in the dispatch prices. In mathematical formula, the lead change is $L(\text{disprice}) =$

[16] $0.10 is between the volume weighted average and raw spreads

Market Structure	Bid Spread ($)	Ask(ed) Spread ($)	Fuel Exposure due to Undernomination ($)	Fuel Exposure due to Overnomination ($)
Symmetric	0.15	0.15	56,198	83,884
Symmetric	0.25	0.25	93,664	139,807
Asymmetric	0.10	0.20	37,465	111,845
Asymmetric	0.35	0.25	131,129	139,807

Table 3.4: Fuel Exposure Impact when Embedded Option Covers Up to 10% of Slack in Gas Flow Imbalances

($\text{disprice}_{t+1} - \text{disprice}_t$), where "disprice" is the dispatch price for that day, and $L(.)$ is the lead operator. Assuming perfect correlation, the intra day price is given by inter day dispatch price + [sign(L(disprice)) * convenience spread], where sign(.) is the sign operator which is negative for negative arguments and positive for positive arguments.

In words, while the magnitude by which the intra day price moves relative to the dispatch price is given by the convenience spread, the actual direction of movement is determined by the sign of the lead change in the dispatch prices. The convenience spread is positive (or the intra day price went up relative to the dispatch price for that day) when the lead change in the dispatch price is positive. Likewise, the convenience spread is negative (or the intra day price went down relative to the dispatch price for that day) when the lead change in the dispatch price is negative.

For the second scenario, the convenience spread is held constant as before but the direction of movement is determined by the sign of lead change. The resulting impact of fuel exposure is calculated in Table 3.5. In the last row of the table, the bid and ask spreads are set to levels observed with the actual intra day data and the resulting fuel exposure impact is a loss of $24,614. This small loss indicates the presence of several arbitrage or discount days.

In the final scenario, the convenience spread is not held constant but allowed to vary in tandem with the lead change in dispatch prices. So, intra day price = disprice + L(disprice). In other words, intra day price is set equal to lead(disprice)

Market Structure	Bid Spread ($)	Ask(ed) Spread ($)	Fuel Exposure due to Undernomination ($)	Fuel Exposure due to Overnomination ($)
Symmetric	0.15	0.15	-9,485	-1,488
Symmetric	0.25	0.25	-15,809	-2,481
Asymmetric	0.10	0.20	-6,323	-1,985
Asymmetric	0.35	0.25	-22,133.41	-2,481

Table 3.5: Fuel Exposure Impact when Intra Day Price = disprice + [sign(L(disprice)) * convenience spread]

or the next day's dispatch price. Under this scenario, the intra day price series have a mean = \$2.49/MBtu and σ = \$0.91/MBtu, same as the dispatch price. The fuel exposure impact due to undernomination adds up to -\$37,997, while that due to overnomination is -\$4,171, a total loss of \$42,168. Again, the resulting fuel exposure impact by using the λ forecaster and unit commitment algorithm combination is quite small! However, the higher levels of volatility observed with the actual intra day prices can hypothetically place the fuel exposure impact anywhere with in the bounds established in the first sceanario. At the same time, invoking the embedded options to absorb small imbalances would further reduce the fuel exposure impact.

3.5 Conclusions

The combination of the λ model and unit commitment algorithm performed well. The analysis suggests the current ad hoc manner of managing fuel exposure may have resulted in actual dispatch that ignored opportunities for further profit. The potential lost profits of $\approx$\$1 Million from Table 3.1, may *as well be considered to represent the fuel exposure impact due to current practices. Then, quite a significant improvement is demonstrated that almost all of these potential profits sans the fuel exposure impact of few thousands[17] is recoverable by employing the strategies and

[17]as estimated in the previous Section 3.4.1

tools developed here. This improvement is possible by methodically considering the opportunity costs and convenience spreads associated with the commodity markets the BTU converter is dealing with.

The limitations with various data or lack of it are mentioned in various parts of the chapter. Section 3.5.1 summarizes the data problems in one place. Several of the implications of the assumptions made for the ED algorithm are already highlighted in Section 3.2. These and other new issues are discussed in Section 3.5.2 in more detail and possible causes are identified. Then Section 3.5.3 discusses probable reasons for gas usage forecast errors.

3.5.1 Limitations of Unit Dispatch and Gas Data

The current system allows multiple dispatch prices to be quoted for a generator and there is an option to switch among the quoted prices intra day. When the dispatch price is changed to a lower or higher price, the unit dispatch and hence the generator output are affected. Although the hourly generation data are recorded, no record of the intra day change in the dispatch price is kept. Of course, the actual amount of gas burnt is a function of generator operation. Hence, the generation data and burns data are not appropriate for comparison with single dispatch price method.

Additionally, the inter day gas nomination data are not recorded separately, the total gas delivered includes the amount of gas nominated day ahead (inter day) plus any other intra day transactions made. Since the goal here is to improve the day ahead predictions, the gas data as collected is not useful unless the two data are separated. Also, keeping the intra day data separate would allow calculation of convenience premiums.

3.5.2 Implications of Simplifying Assumptions for the Economic Dispatch Algorithm

The algorithm assumes instant change in the loading level of a generator i.e. the generator is either not loaded (off-state) or can reach any of the pre-defined loading blocks in no time. The newer units typically have response rates of several MW/minute and can reach full-load capacity in less than a half hour. While the assumption is not too simplifying at hourly level, zero ramp-times and block loading level imparts discrete states for the generator output while the actual output is more continuous.

Another reason that contributes to the finite state output, hence the highly modal error structure in plots such as Figure 3-3, has to do with operating a single unit in isolation from the pool. With only one unit under consideration, the next cheapest incremental block of capacity available is going to be from the one unit itself. Whereas with several units in the pool, the next block of capacity to be committed will be the cheapest block across all the units in the pool. Though it sounds serious, it is not an over simplification unless there are several similar units in the pool with incremental heat rates (and fuel price) close enough to cause one unit to run below full-load capacity. Thus we will see the unit will be either on at full-load (where it is cheapest) or off at zero or minimum-load when trading-off with start-up cost.

Some of the simplifying assumptions, like start-up cost recovery or amortization, can be relaxed. Just like start-up cost is compared to minimum-load operation when recommended for shut-down, the recovery of start-up cost is critical when going to on-state. The length of time period over which start-up cost needs to be recovered can influence the economics of when to turn the unit on. As an extreme example, if the policy is to recover all the start-up costs with in the first hour of operation, then a 170 MW unit with a start-up cost of $5,900 would be committed for generation whenever the spark spread is greater than $35 + dispatch cost. There are probably a few on-peak hours that satisfy such a constraint. While start-up cost amortization would reduce the number of hours of operation (and start-ups) and hence undercommitment discrepancy, further analysis is warranted to find out the best recovery period that maximizes net profit. Currently, NEPOOL performs unit commitment on a day by day basis using their Optimized Daily Forecast (ODF) algorithm. NEPOOL's ODF accepts pumped storage and hydro schedules and looks at 24 hour periods (one day-ahead) for thermal unit commitment.

Further improvement in the ED projections could be possible by improving the forecast accuracy of hourly λ itself. Absolute prediction accuracy (as shown by Plots 3-5 and 3-3) close to 100% could be achieved by making the λ model robust to mean and variance changes in hourly λ. The heteroscedastic nature of λ causes the errors to occur in bunches and, if not taken care of, could potentially throw an aggregated forecast (like daily totals calculated from hourly data) far off. This is the reason for decrease in the prediction power when comparing the hourly unit commitment plot 3-3 with the daily gas usage plot 3-6. These issues are also discussed in Section 2.4.

3.5.3 Limitations for Forecasting Gas Usage

The hourly plots of pool load, actual and predicted λ are scrutinized for periods including those days when the day ahead gas usage forecast error is extreme. One of the most plausible reason for the extreme errors seems to be lacking any generation input to the model. Actions such as peak shaving with hydro and/or forced outages often alter the ED conditions substantially thus causing the generator to run much different than anticipated. A major capacity loss intra day, such as a nuclear reactor outage, can can double or triple the hourly λ transiently until a cheaper unit picks up the slack. It can also raise the daily average λ by by as much as \$2/MWH until the base capacity is restored. Events such as this would result in day-ahead prediction errors that persist for at least two days in a row.

At least one extreme error is found due to an incorrect data point (outlier) in the hourly pool load data. Although, the reason for forecast error is now obvious, this does suggest the need to incorporate routine checks for outliers in the data before input to the model.

Error = mw_rl - mw_ed

mw_rl: actual hourly generation

mw_ed: optimum (simulated) hourly generation

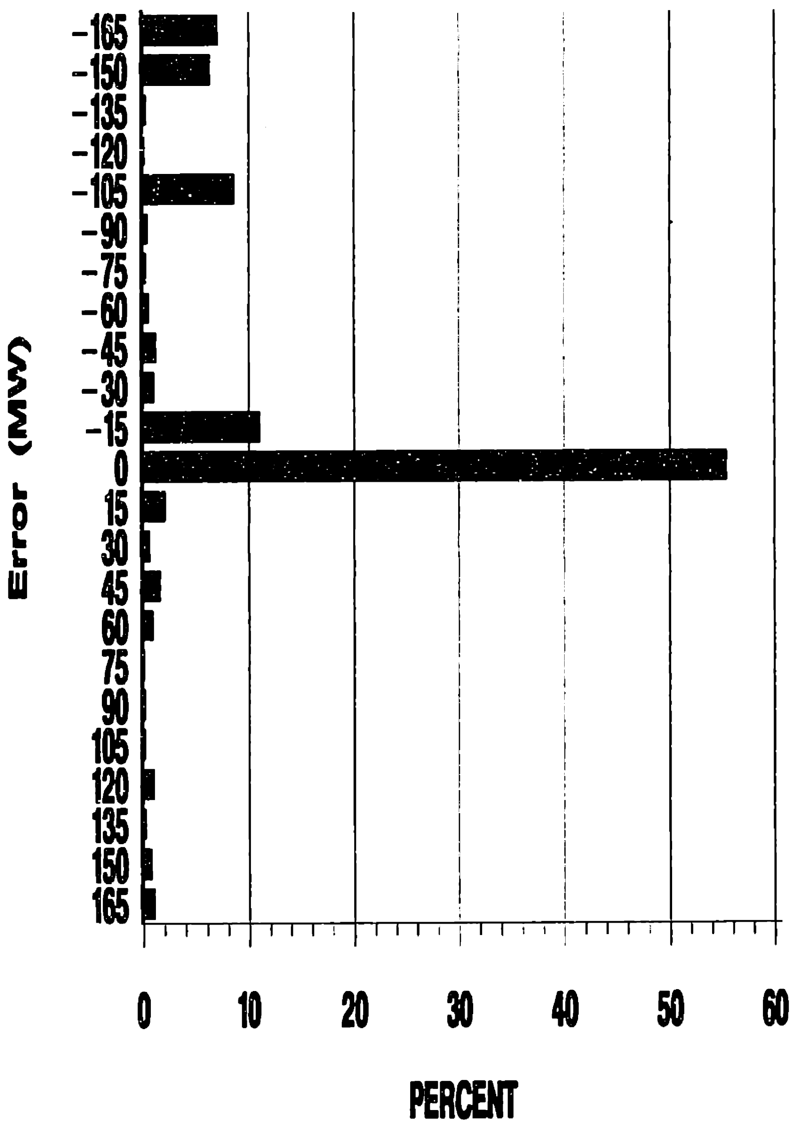

FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
584	584	7.01	7.01
525	1109	6.30	13.32
19	1128	0.23	13.54
9	1137	0.11	13.65
716	1853	8.60	22.25
33	1886	0.40	22.65
20	1906	0.24	22.89
42	1948	0.50	23.39
97	2045	1.16	24.56
83	2128	1.00	25.55
915	3043	10.99	36.54
4610	7653	55.36	91.89
168	7821	2.02	93.91
47	7868	0.56	94.48
131	7999	1.57	96.05
70	8069	0.84	96.89
5	8074	0.06	96.95
10	8084	0.12	97.07
10	8094	0.12	97.19
77	8171	0.92	98.11
16	8187	0.19	98.31
57	8244	0.68	98.99
84	8328	1.01	100.00

Figure 3-1: Distribution of (Actual - Optimum) Hourly Generation

Error = mw_rl - mw_ed

mw_rl: actual daily generation

mw_ed: optimum (simulated) daily generation

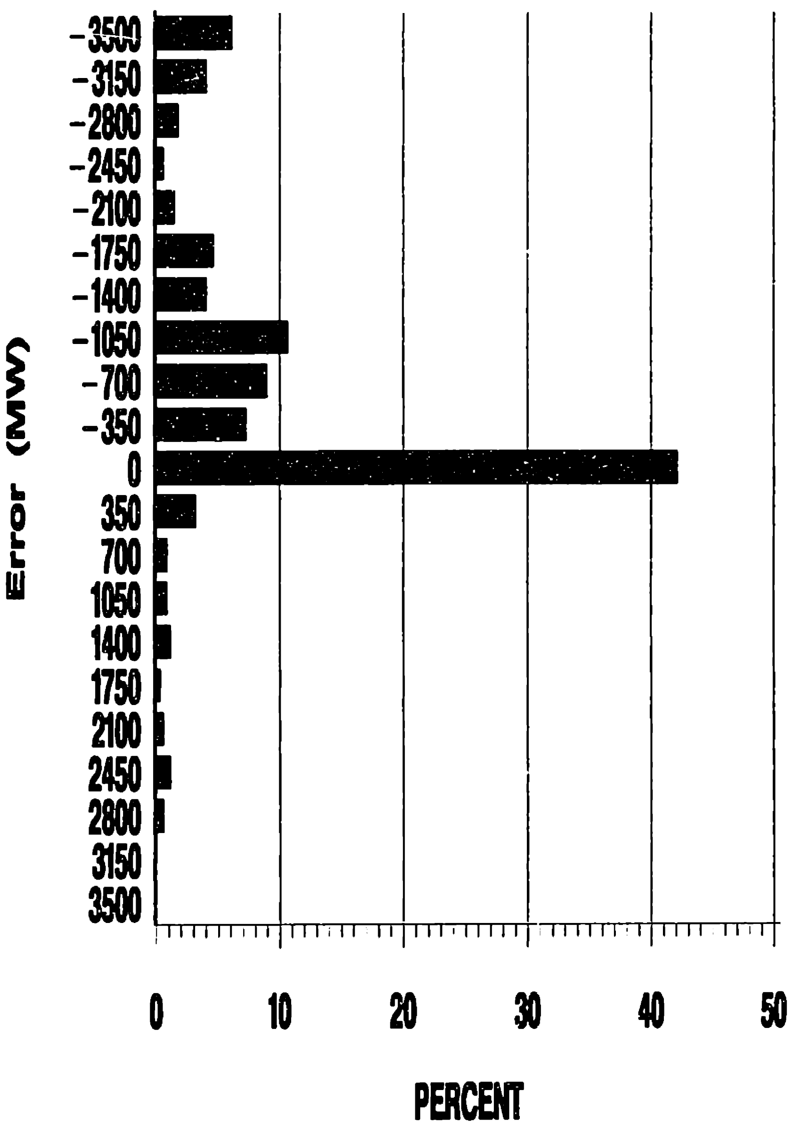

Error (MW)	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-3500	21	21	6.05	6.05
-3150	14	35	4.03	10.09
-2800	6	41	1.73	11.82
-2450	2	43	0.58	12.39
-2100	5	48	1.44	13.83
-1750	16	64	4.61	18.44
-1400	14	78	4.03	22.48
-1050	37	115	10.66	33.14
-700	31	146	8.93	42.07
-350	25	171	7.20	49.28
0	146	317	42.07	91.35
350	11	328	3.17	94.52
700	3	331	0.86	95.39
1050	3	334	0.86	96.25
1400	4	338	1.15	97.41
1750	1	339	0.29	97.69
2100	2	341	0.58	98.27
2450	4	345	1.15	99.42
2800	2	347	0.58	100.00
3150	0	347	0.00	100.00
3500	0	347	0.00	100.00

Figure 3-2: Distribution of (Actual - Optimum) Daily Generation

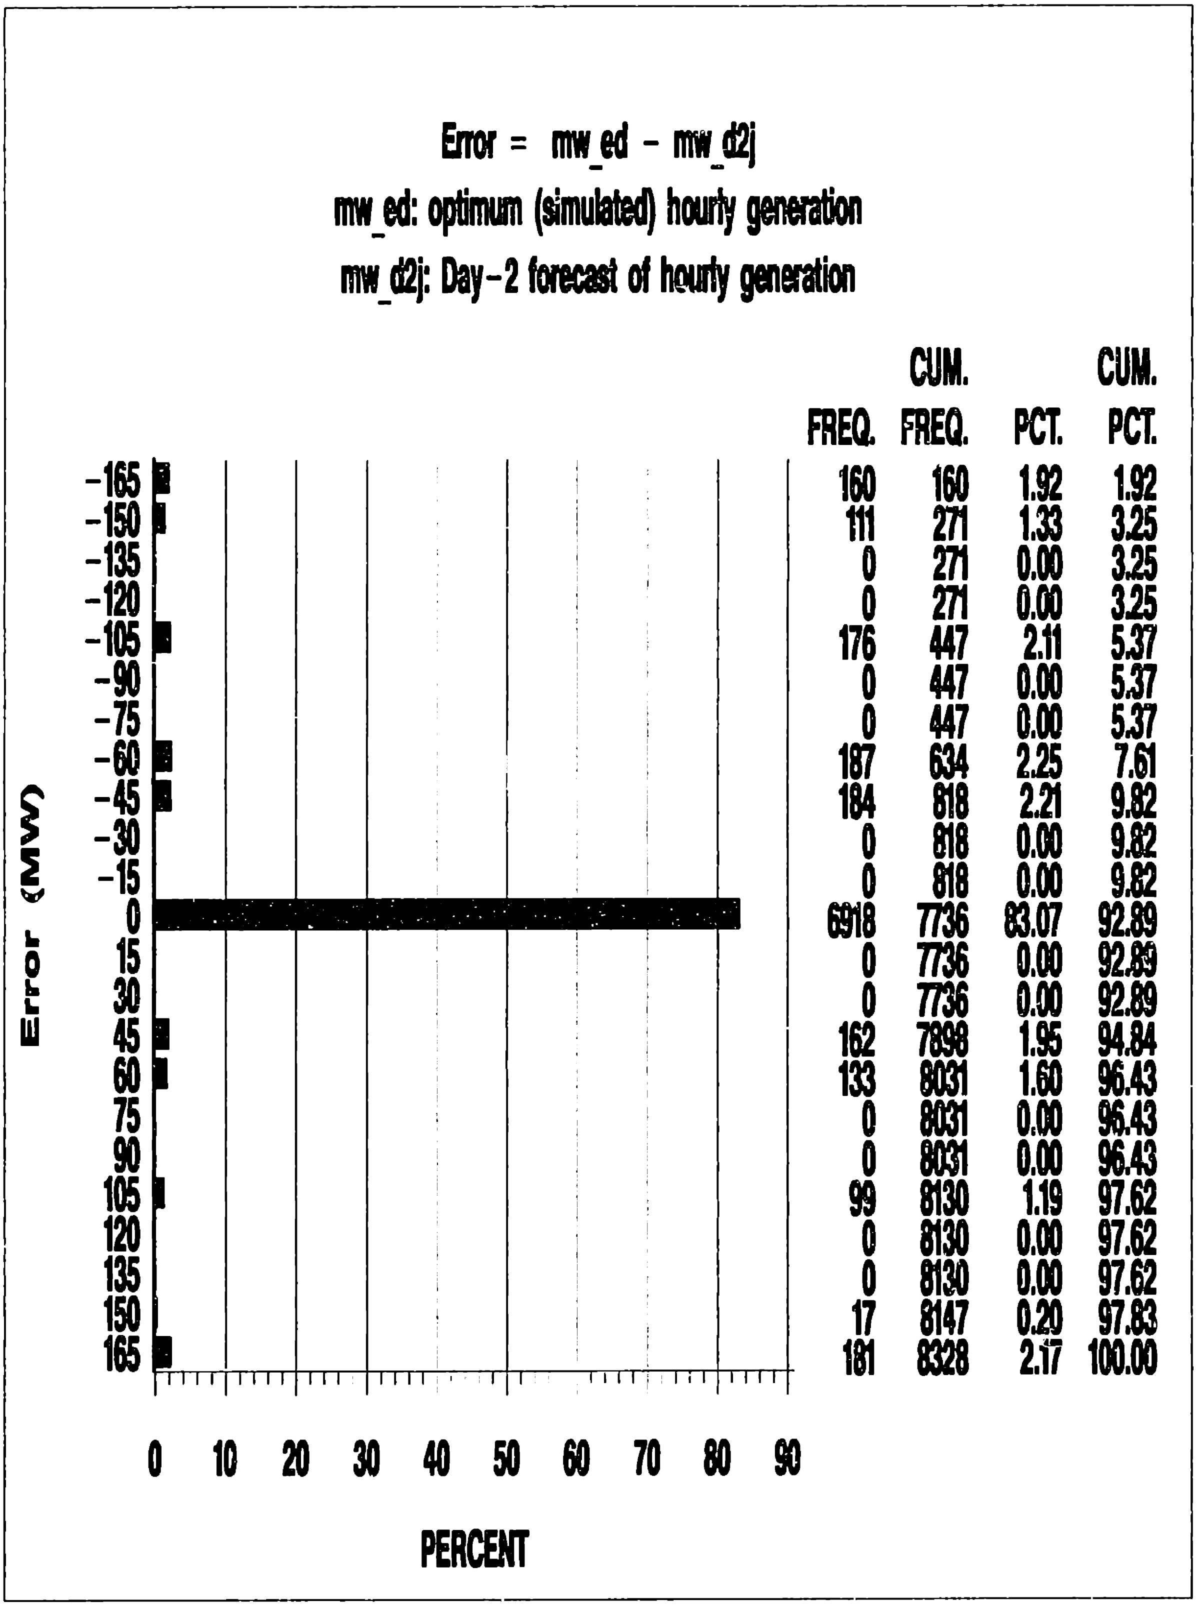

Error (MW)	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-165	160	160	1.92	1.92
-150	111	271	1.33	3.25
-135	0	271	0.00	3.25
-120	0	271	0.00	3.25
-105	176	447	2.11	5.37
-90	0	447	0.00	5.37
-75	0	447	0.00	5.37
-60	187	634	2.25	7.61
-45	184	818	2.21	9.82
-30	0	818	0.00	9.82
-15	0	818	0.00	9.82
0	6918	7736	83.07	92.89
15	0	7736	0.00	92.89
30	0	7736	0.00	92.89
45	162	7898	1.95	94.84
60	133	8031	1.60	96.43
75	0	8031	0.00	96.43
90	0	8031	0.00	96.43
105	99	8130	1.19	97.62
120	0	8130	0.00	97.62
135	0	8130	0.00	97.62
150	17	8147	0.20	97.83
165	181	8328	2.17	100.00

Figure 3-3: Distribution of Prediction Error for Day Ahead Generation

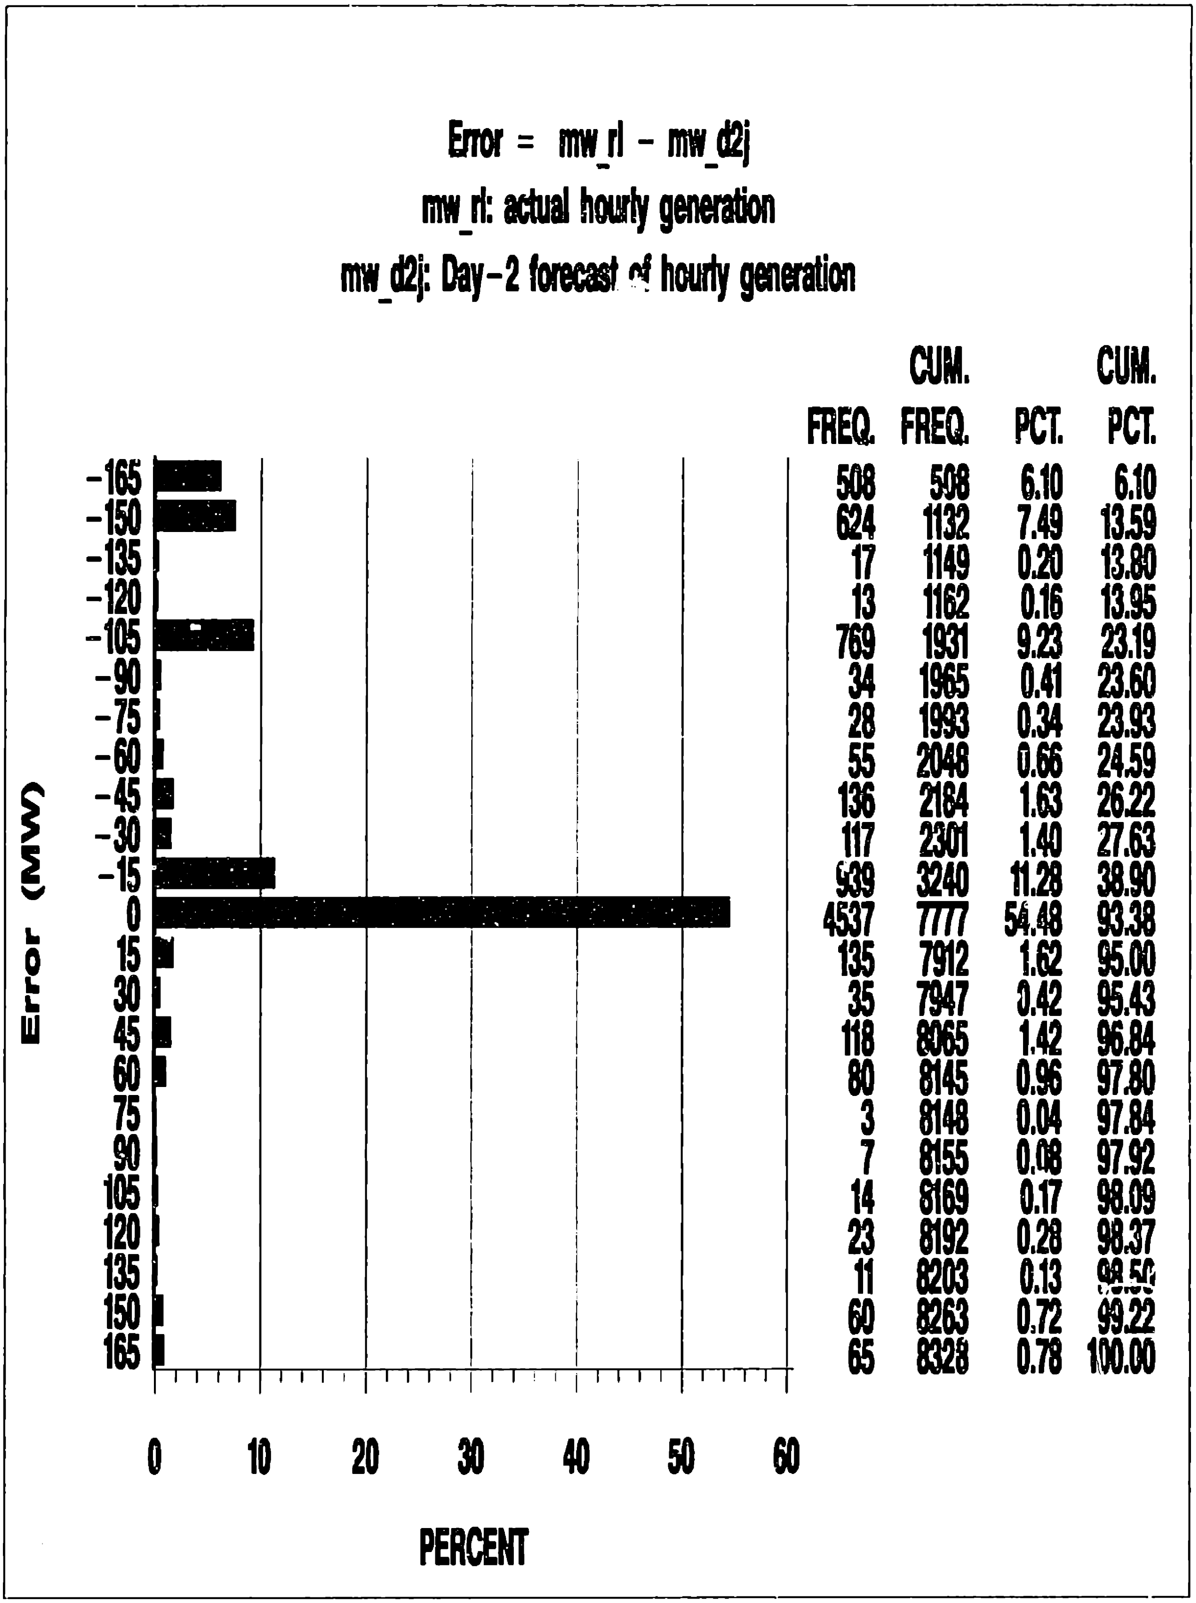

Figure 3-4: Distribution of Prediction Error for Day Ahead When Compared to Actual Generation

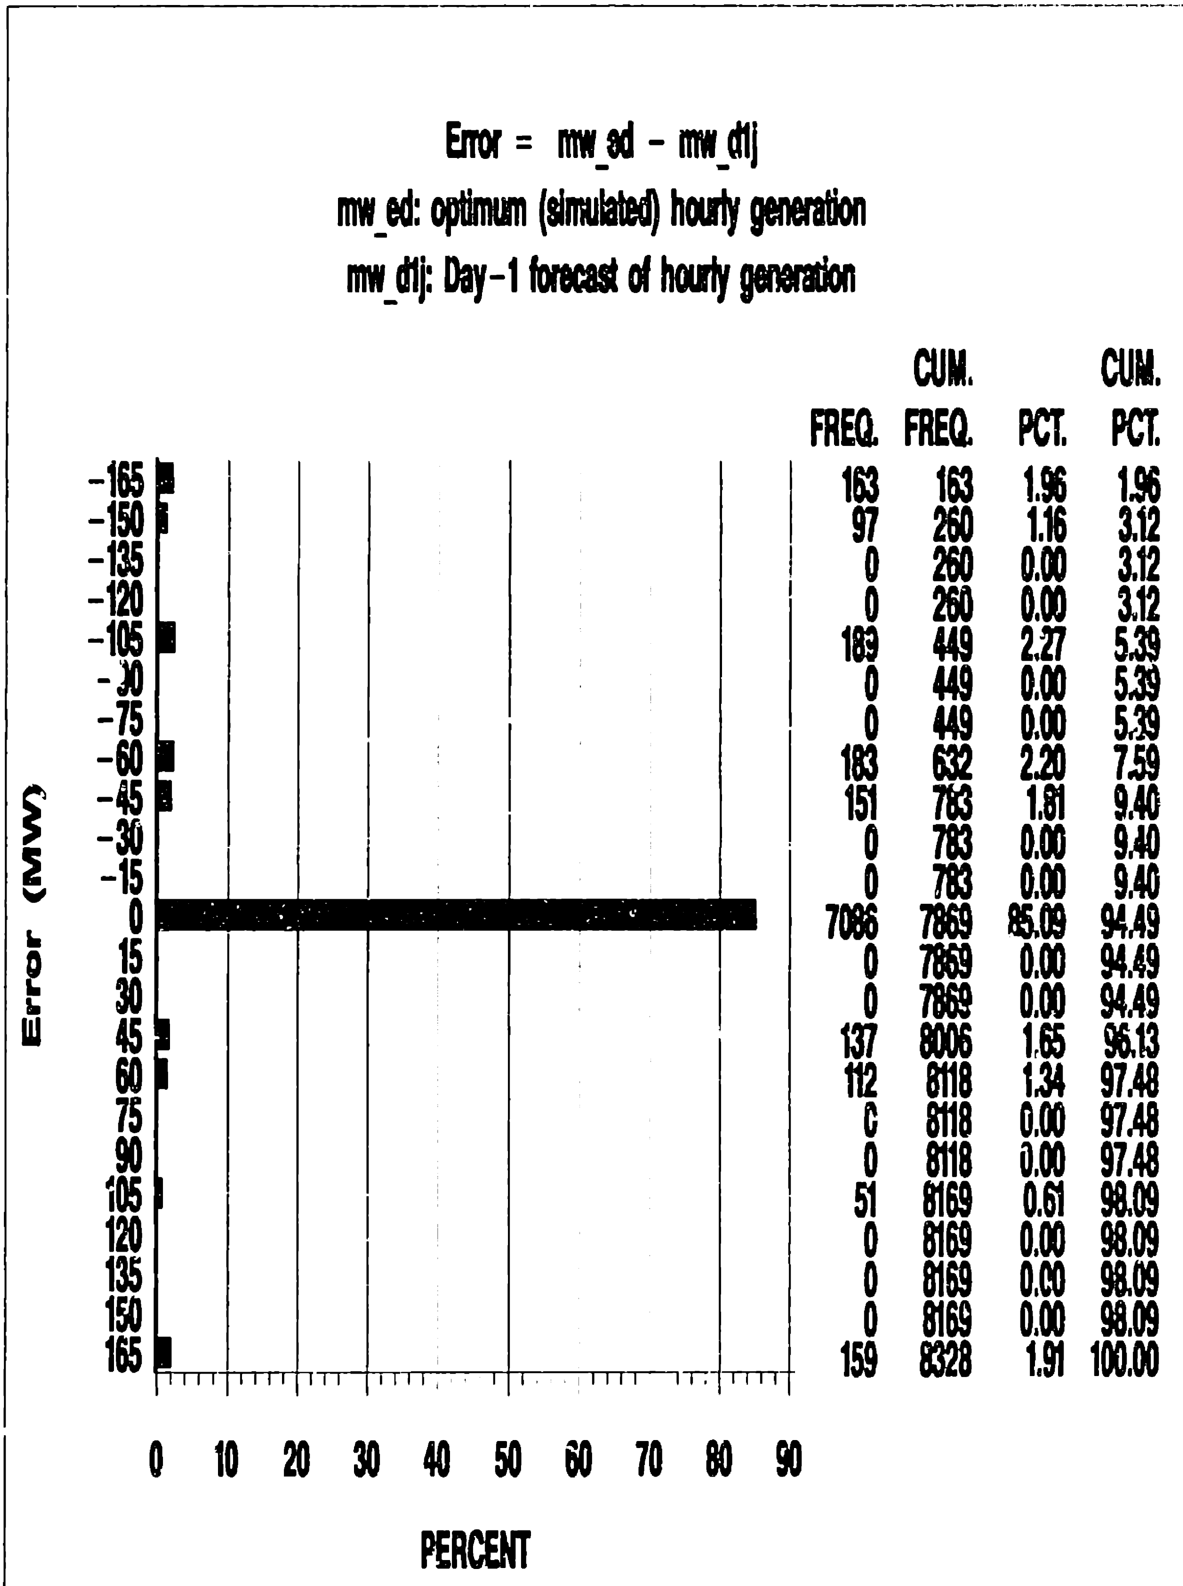

Error (MW)	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-165	163	163	1.96	1.96
-150	97	260	1.16	3.12
-135	0	260	0.00	3.12
-120	0	260	0.00	3.12
-105	189	449	2.27	5.39
-90	0	449	0.00	5.39
-75	0	449	0.00	5.39
-60	183	632	2.20	7.59
-45	151	783	1.81	9.40
-30	0	783	0.00	9.40
-15	0	783	0.00	9.40
0	7086	7869	85.09	94.49
15	0	7869	0.00	94.49
30	0	7869	0.00	94.49
45	137	8006	1.65	96.13
60	112	8118	1.34	97.48
75	0	8118	0.00	97.48
90	0	8118	0.00	97.48
105	51	8169	0.61	98.09
120	0	8169	0.00	98.09
135	0	8169	0.00	98.09
150	0	8169	0.00	98.09
165	159	8328	1.91	100.00

Figure 3-5: Distribution of Prediction Error for Generation up to 24 Hours Ahead

Error = gs_ed - gs_d2j

gs_ed: gas used in optimal generation (electric day)

gs_d2j: Day-2 forecast of gas usage (electric day)

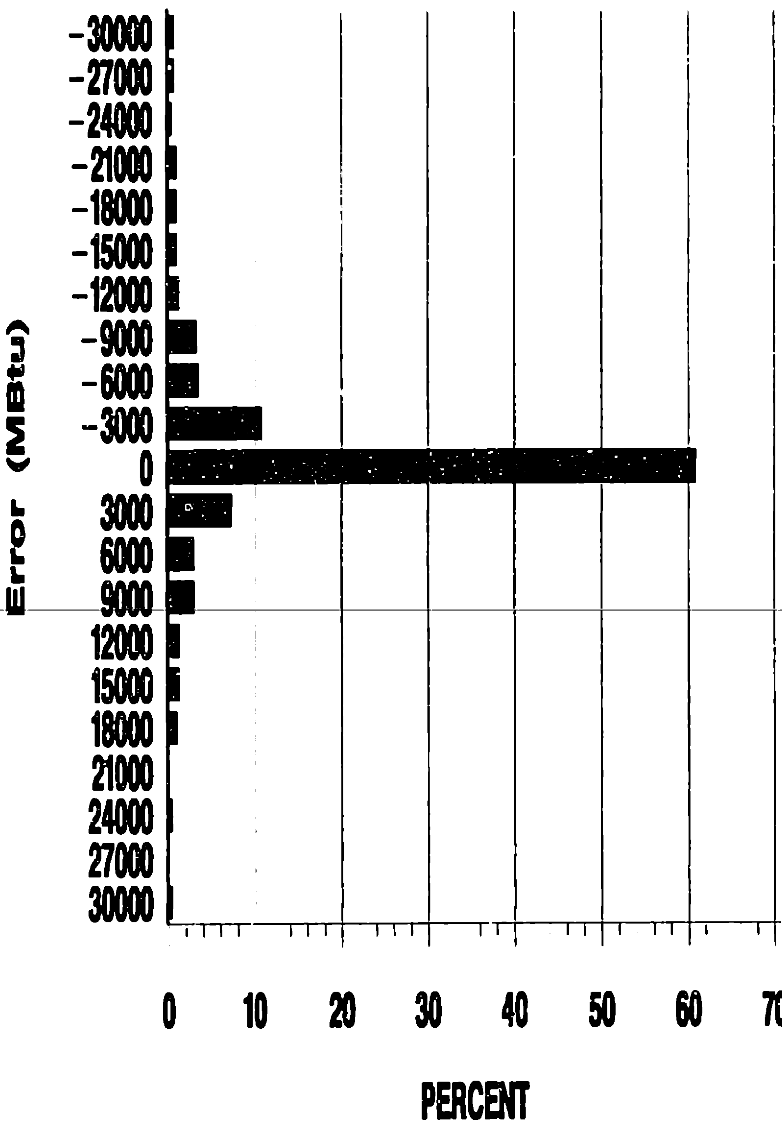

Error (MBtu)	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-30000	2	2	0.58	0.58
-27000	2	4	0.58	1.15
-24000	1	5	0.29	1.44
-21000	3	8	0.86	2.31
-18000	3	11	0.86	3.17
-15000	3	14	0.86	4.03
-12000	4	18	1.15	5.19
-9000	11	29	3.17	8.36
-6000	12	41	3.46	11.82
-3000	37	78	10.66	22.48
0	211	289	60.81	83.29
3000	25	314	7.20	90.49
6000	10	324	2.88	93.37
9000	10	334	2.88	96.25
12000	4	338	1.15	97.41
15000	4	342	1.15	98.56
18000	3	345	0.86	99.42
21000	0	345	0.00	99.42
24000	1	346	0.29	99.71
27000	0	346	0.00	99.71
30000	1	347	0.29	100.00

Figure 3-6: Distribution of Prediction Error for Day Ahead Gas Usage

Error = (gs_ed - gs_d2j)/gs_d2j

gs_ed: gas used in optimal generation (electric day)

gs_d2j: Day-2 forecast of gas usage (electric day)

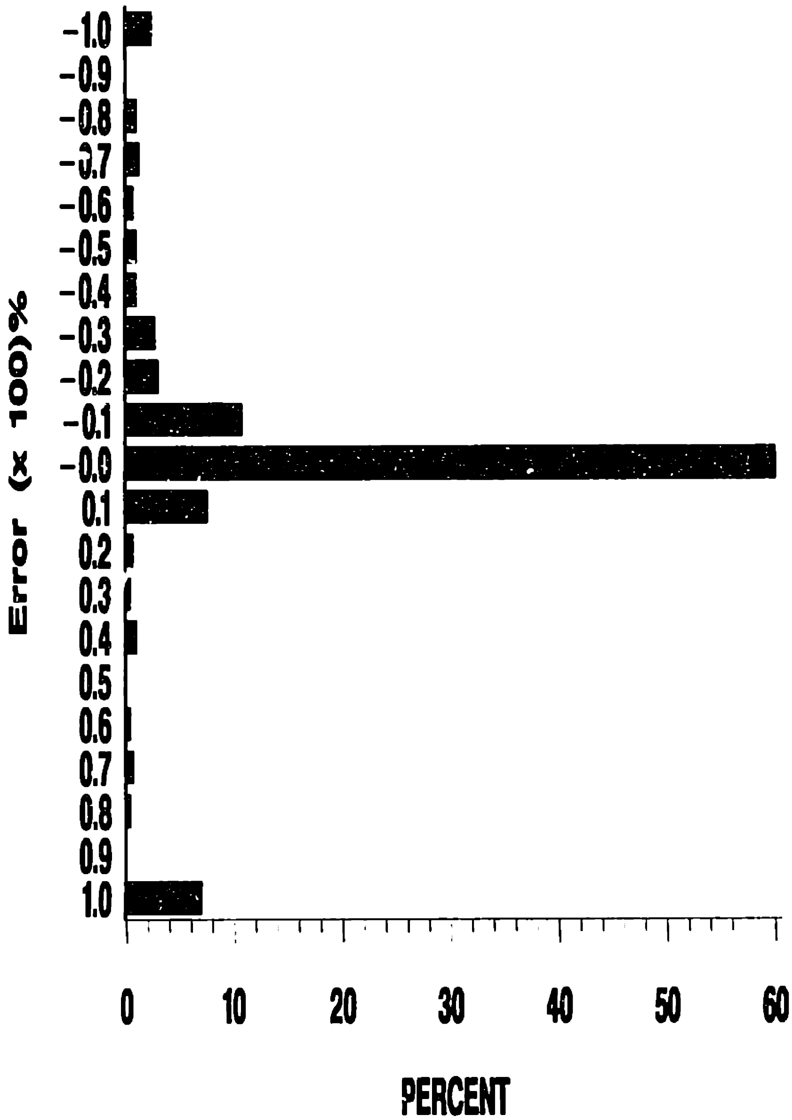

Error (x 100)%	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
-1.0	8	8	2.31	2.31
-0.9	0	8	0.00	2.31
-0.8	3	11	0.86	3.17
-0.7	4	15	1.15	4.32
-0.6	2	17	0.58	4.90
-0.5	3	20	0.86	5.76
-0.4	3	23	0.86	6.63
-0.3	9	32	2.59	9.22
-0.2	10	42	2.88	12.10
-0.1	37	79	10.66	22.77
-0.0	208	287	59.94	82.71
0.1	26	313	7.49	90.20
0.2	2	315	0.58	90.78
0.3	1	316	0.29	91.07
0.4	3	319	0.66	91.93
0.5	0	319	0.00	91.93
0.6	1	320	0.29	92.22
0.7	2	322	0.58	92.80
0.8	1	323	0.29	93.08
0.9	0	323	0.00	93.08
1.0	24	347	6.92	100.00

Figure 3-7: Distribution of Prediction Error as a % of Day Ahead Nomination

Appendix A

Tables of Statistics for Various Distribuitons

Moments	
Number of observations	8,759
Mean	0.0004
Standard Deviation	1.887
Skewness	0.737
Kurtosis	18.19
Quantiles (Def=5)	
Maximum	25.34
Median	-0.07
Minimum	-17.14
Mode	0

Table A.1: Relevant statistics of $\Delta\lambda = \lambda_t - \lambda_{t-1}$

Moments		
Forecast Period	Day-1	Day-2
Number of observations	8,568	8,520
Mean	0.064	0.078
Standard Deviation	2.710	2.927
Skewness	1.996	1.70
Kurtosis	15.822	12.866
Probability the distribution is Normal	0.089	0.077
Quantiles (Def=5)		
Maximum	35.30	34.610
Median	-0.048	-0.022
Minimum	-12.224	-10.749
Mode	-12.224	-10.749

Table A.2: Statistics of Interest for *One-Stage* Model

Moments		
Forecast Period	Day-1	Day-2
Number of observations	8,544	8,472
Mean	0.045	0.019
Standard Deviation	2.734	2.848
Skewness	1.853	1.334
Kurtosis	14.823	10.474
Probability the distribution is Normal	0.087	0.068
Quantiles (Def=5)		
Maximum	35.032	34.983
Median	-0.044	-0.043
Minimum	-12.309	-10.725
Mode	-12.309	-10.725

Table A.3: Statistics of Interest for *Two-Stage* Model

Moments	
Number of observations	8,328
Mean	-28.324
Standard Deviation	67.171
Skewness	-0.599
Kurtosis	0.720
Probability the distribution is Normal	0.301
Quantiles (Def=5)	
Maximum	171
Median	-0.08
Minimum	-171
Mode	0

Table A.4: Relevant Statistics for Distribution of (Actual - Optimum) Hourly Generation

Moments	
Number of observations	347
Mean	-680
Standard Deviation	1,236
Skewness	-0.707
Kurtosis	0.855
Probability the distribution is Normal	0.861
Quantiles (Def=5)	
Maximum	2,749
Median	-164
Minimum	-3,834
Mode	0

Table A.5: Relevant Statistics for Distribution of (Actual - Optimum) Daily Generation

Moments	
Number of observations	8,328
Mean	-2.807
Standard Deviation	46.514
Skewness	-0.216
Kurtosis	7.176
Probability the distribution is Normal	0.426
Quantiles (Def=5)	
Maximum	171
Median	-0.0
Minimum	-171
Mode	0

Table A.6: Relevant Statistics for Distribution of Prediction Error for Day Ahead Generation

Moments	
Number of observations	347
Mean	-570
Standard Deviation	6,441
Skewness	-0.929
Kurtosis	7.794
Probability the distribution is Normal	0.77
Quantiles (Def=5)	
Maximum	28,696
Median	0
Minimum	-34,370
Mode	0

Table A.7: Relevant Statistics for Distribution of Prediction Error for Day Ahead Gas Usage

Bibliography

[1] Terence C. Mills "Time series techniques for economists" Cambridge University Press 1990.

[2] Donald B. Percival, and Andrew T. Walden "Spectral Analysis for Physical Applications" Cambridge University Press 1993.

[3] James D. Hamilton "Time Series Analysis" Princeton University Press 1994.

[4] "SAS/ETS Software Usage and Reference" Sas Institute Inc.

[5] Robert Engle, et al., "Short-Run Forecasts of Electricity Loads and Peaks"