RATE-REVENUE-COST OF SERVICE SIMULATION OF A NATURAL GAS UTILITY

D. Jeffrey Blumenthal

On-Line Decisions, Inc.

ABSTRACT

A major midwestern utility was faced with the dilemma of increased demands for natural gas energy while the supplies of natural gas were diminishing. Questions on how to meet seasonal or off-peak demands, how to utilize underground storage, and how much new supply would be required to eliminate waiting lists for residential customers were typical of the problems confronting this utility. In addition, the cost of service due to inflation and pipeline cost adjustments indicated that this company should perhaps request a general rate increase and required substantial calculations in the PGA (Purchase Gas Adjustment) area.

A set of linked models was built for this utility to answer these questions. The models considered the areas of rate, revenue, cost of gas, PGA, and income. Each model could be run independently of the others; however, all could be linked to provide an integrated planning system for the utility. The models were implemented in a conversational mode on a time-shared computer and allowed the model users to interactively vary rate structures, pipeline contract terms and alternative sources of supply over a five-year planning horizon.

INTRODUCTION

Historically, natural gas has been considered as an inexpensive source of energy. In addition, it is virtually nonpolluting. As a consequence, the demand for natural gas energy has reached all-time highs. The traditional peak demand period for natural gas has been the winter space heating season. However, because of the non-polluting aspects of natural gas energy, electric power generating companies have placed off season or summer demands on natural gas energy. Many companies have constructed under-ground storage facilities for servicing the peak winter heating season demands. Most natural gas utilities, therefore, have a variety of alternatives on how or to whom they should sell natural gas. The supply of natural gas, on the other hand, has remained relatively constant for the last several years, resulting in an apparent shortage of natural gas in the past several winter heating seasons. New supplies or potential new sources of natural gas appear to be considerably more expensive than existing ones. New processes, such as coal gasification, represent a potential energy source. However, their impact is still several years off. A proposed pipeline from the north slope of Alaska would represent a considerable capital investment on the part of a pipeline company. Liquefied natural gas from the Middle East has been suggested for alleviating some of the demand requirements on the Eastern Seaboard. No doubt a combination of all of these alternatives is possible. Investor-owned public utilities are regulated either by the Federal Power Commission or by a statewide organization. Servicing an area rapidly growing in population, a major midwest natural gas utility was faced with finding supplies of new, more expensive gas. The company was also involved in major investments in underground storage facilities. With increased investments in plant and equipment, would a rate increase be justified? How fast would the company recoup additional costs of the more expensive gas? What new supplies would be required to eliminate waiting lists and satisfy the off-peak demands of power generating companies? The answers to these questions were not readily available. Moreover, in seeking answers, the company uncovered an additional problem--

that of varying internal forecasts. Marketing based its figures on one set of assumptions, while Operations worked on yet another. The Treasurer's group at headquarters was assigned to reconcile these forecasts into a unified operations plan. In spite of earlier failure at modeling, the company felt that simulation models could be useful in consolidating these forecasts and developing a unified operating plan which would answer questions concerning alternative supply sources, sales to customers on waiting lists, and off-peak demand sales.

OVERVIEW OF THE MODEL SYSTEM

Most utilities normally function in four operational areas. These are:

1. Rates and revenues
2. Operating costs
3. Construction and capital budgeting
4. Financing

A fifth area, that of consolidation, is also necessary. Ogden (Reference 1) discussed the problem of modeling these particular areas and illustrated the types of questions and data required.

The needs for the company under discussion here, however, and for most operating gas utilities today, is in the areas of rates and revenues and the operating costs. These areas require additional detail to answer "what if" questions, necessary for computer simulation. Consequently, a top level planning system was developed for this utility illustrated in Figure 1, and consisting of five separate operating modules.

The basic design was developed during a one-month initial Phase I study, consisting of interviews with key executives within the organization. Data availability and traditional forecasting methods were also studied at this time. The result of Phase I was the five module system to simulate the rate and cost of service area for this utility by months over a time horizon of five years.

The Rate Module isolated the smallest definable element of revenue to the company. Customers were classified by the rate schedule utilized, their particular revenue class, such as residential space heating, and by billing method. The primary purpose of the Rate Module was to determine the actual therms sold and the base dollars of revenue generated in each of these elements.

The Revenue Module took the therms sold from the Rate Module and determined the therm sendout--the actual number of therms distributed in a particular month. The sendout is related, but not equivalent to the therms actually billed. In addition, the revenue model added various adjustments such as purchased gas adjustment (PGA) and various state and municipal taxes to the base dollar revenue.

The Cost of Gas Module was primarily concerned with pipeline contract purchases and various injections and withdrawals from the company's own underground storage. The company utilized several methods of inventorying storage gas. The model valued inventory to arrive at a total gas distributed figure. The model, of course, evaluated precise pipeline purchases and the cost of those purchases to be used by the PGA model. New, as well as existing sources, were considered.

The PGA module duplicated the actual PGA calculations. The cost of purchases for twelve of the past thirteen months was compared to a base cost. The difference was amortized by formula to a rate per therm. Six major PGA rates were calculated in this module. These rates were then fed back to the Revenue Module for calculation of total dollar revenue.

The Income Module served to consolidate the revenue and cost dollars to arrive at a simple income statement. Some of the outputs were an earnings statement and a return on the rate base.

TECHNIQUES USED IN MODEL CONSTRUCTION

As mentioned earlier, one of the basic outputs of the Rate Model was the forecast of therm consumption. For certain elements, that is, rate-revenue class-bill basis, only a small number of customers were involved and here forecasting therms sold was done by simple extrapolation of past data. In a preponderance of cases, however, a large number of customers would be represented by a particular rate element. Here, therm use was forecast by breaking out the average use per customer and multiplying by the number of customers.

This breakout of the number of customers was extremely useful in simulating a waiting list of various types of service. The model user could vary the number of customers to suit the availability of gas.

Forecasting the average use per customer, on the other hand, was more difficult. In the majority of cases, the amount of usage by a particular customer is variable both with the respect to time. In fact, the largest class of natural gas users were those whose usage depend heavily on weather, i.e., space heating customers.

Figure 2 illustrates the relationship between daily therm usage and outside temperature for a typical residential space heating customer. The figure illustrates the relatively nonlinear behavior of usage versus temperature. A simple transformation of data allows the curve to be broken into two approximately linear sections. This is done by defining what is known as the degree day. By performing a linear regression of usage against degree days, one can determine the base use as the intercept, and the slope as the use per degree day. Figure 2 illustrates this regression line.

In actual fact, however, the problem is not quite that simple. As can be seen in Figure 2, breaking the curve into two linear segments does not match up over the entire range. The basic departure of the regression line from actual usage occurs at both ends of the temperature curve.

To isolate temporal effects from the basic weather effects, the linear regression of average usage per customer as a function of degree days was performed on a twelve month rolling basis. A regression was performed for twelve months points, then the oldest term was dropped out and a new term added. By repeatedly doing the regression, one could take the resulting figures for base use and use per degree day and extrapolate trends in their growth. The model allowed the user to assume various growth rates in both the base use and the use per degree day.

To calculate the base dollars of revenue was perhaps the most complicated of the forecasting formulas used in the model system. A typical rate schedule as illustrated by Figure 3, explains the techniques used for the majority of cases which depicts a space heating rate.

Shown in figure 3 is the base dollars of revenue to a particular customer, based on the total amount of therms that customer uses within a particular month. The figure illustrates that the rate schedule consists of a series of piece-wise linear segments, each one known as a block. Since a straight line can be represented by two points, an intercept and a slope, it is possible to relate the dollars of base revenue as a function of therm within any particular block by such two numbers. If all customers utilized the same amount of gas, the base dollars of revenue could be simply calculated by taking the revenue per customer times the number of customers, where the revenue per customer could be computed from a fixed plus a variable component, depending on usage multiplied by the number of therms. This, of course, can be easily recognized as a simple linear regression. Indeed, regression is often utilized for calculating basic revenue as a function of average therms used. However, all customers don't utilize the same amount of gas. Figure 4 illustrates bill density, basically a distribution of the number of customers utilizing a specific number of therms. The particular shape of the curve, of course, varies from month to month and depends a great deal on the type of service offered. Bill density, or bill frequency analysis, as it is sometimes called, is performed by most utilities. Consequently, curves of the type in Figure 4, are often readily available. Once such a curve has been defined, it is possible to calculate the base dollar revenue. If

(3.2) $f(t)dt$ = number of customers using t to t + dt therms then, if we let r(t) represent the revenue from the rate structure, (3.3) r(t) = base revenue/customer using t therms

then the revenue can be calculated as

(3.4) $\text{Revenue} = \int_0^{\infty} r(t)\; f(t)\; dt$

While very attractive from a theoretical point of view, Equation 3.4 is not very useful from a practical point of view. For one thing, an integral of infinite range is quite hard to simulate.

One's first impulse, is to limit the range of the integration. However, virtually any limitation usually excludes the one oddball customer extremely far out on the usage curve, and it is precisely this customer generating high use, who represents large dollar revenues.

To get around this problem, an alternative formulation can be developed. This is done by considering the percentage of the total number of therms which are sold in a particular block. Since all but the last block end at a finite number of therms, it is possible to calculate the total number of therms used in all, except the last block. The remainder, of course, is then allocated to the final block so that the infinite integral never has to be evaluated. To calculate the number of therms used in a particular block, one defines what is known as an ogive. The ogive, or commodity distribution, is defined by

(3.5) $G(x) = \int_0^{x} t f(t)dt + x\int_x^{\infty} f(t)dt$

G(x) represents the total number of therms sold in the blocks 0 to X therms. It consists, basically, of two parts, the first part containing all those therms sold to customers whose usage terminated in the block 0 to X. This is represented as the integral from 0 to X of a particular therm level times the number of customers in that level integrated from O to X. For customers whose bills terminate beyond the block O to X, their total usage is simply X times the number of customers whose bills fall outside of that range. Although the second integral appears to be of infinite range, it should be recognized that the number of customers whose bills do not terminate between O and X therms is equal to the total number of customers less those whose bills do, in fact, terminate between O and X. Consequently, one can avoid performing the infinite integration. The ogive corresponding to the bill density illustrated by Figure 4, is shown in Figure 5. The use of ogives is discussed in most introductory texts on rate making fundamentals. However, a particularly good treatment, one which is correct in mathematical terms, is given in Reference 2. There are several normalizations which can be performed on the bill density and ogive distribution. The first of these normalizations is to express the number of customers or number of therms; that is, representing the functions f(t) and G(x) as a number ranging between O and 1. This normalization makes the bill density and ogive curves correspond roughly to probability densities and distributions. A second normalization used is to translate the therms used by a particular customer and relate them to the average for the entire therm usage. As an example, a customer utilizing 150 therms in a particular month, where the average customer utilized 75 therms, would be normalized to a therm usage of 150 divided by 75, or 2. A rather surprising result of this normalization is to remove the effects of weather from the ogive and bill density. The utility under study constructed distributions for each rate and revenue class on a monthly basis. The advantage of using an ogive is obviously that one can examine both changes in a particular rate in a particular block or the changing of individual block sizes. Such analyses are essential for any rate case preparation. The model constructed for this utility could, of course, examine various rate changes and prepare the basic data needed for rate analysis. However, ogives prepared on past data are only history. If one can assume that basic distribution of customer usage remains constant, then utilizing an experimentally observed ogive may be fine for the forecasting problem. However, where it is expected that the density will change its shape; that is, skew one way or the other, depending on the types of customers brought on, then it is necessary to use techniques other than the experimentally observed ogive to calculate the revenue. A technique utilized in this model was to take a curve such as illustrated in Figure 4 and represent it by a mathematically defined probability density function. For instance, Figure 4 was found to be similar to the gamma-1 probability density function (Reference 3). The gama density is a large family of two parameter distributions. Since the bill density is normalized to unit mean, a two parameter distribution will have one remaining parameter, the variance, for adjustment. By changing the variance, the gamma distribution could be skewed to the right or left and vary its shape over a considerable range.

The company's existing bill frequency analysis program was modified to plot curves similar to Figures 3 and 4, as well as giving the normalized ogive at 100 selected points and calculating what value of the variance to use in the gamma-1 density function. In addition to the gamma-1 density, a pareto distribution and several single parameter density functions were also made available to the program, including exponential, uniform and the triangular density functions.

With this added flexibility, a model could be used not only to analyze changes in the rates and the blocking size, but also examine the effect of various distribution changes in the types of customers utilizing a particular rate. Despite the relatively large number of assumptions and smoothings of data, beginning with the degree day-average use relationships and culminating in the use of the gamma function to express customer density, the overall error, in both therms forecasted and basic dollars revenue for the basic revenue elements, did not exceed one half of one percent in total.

The therms used by each of the revenue elements served as a basic input to the revenue model, whose first function was to calculate actual therm sendout. For customers billed bimonthly, a bill would represent sendout extending back into the past two months. If one assumes that customers are billed uniformly throughout the month, one arrives at what is commonly known as the 25-50-25 rule, indicating that, of the therms billed in a particular month, 25% of them were sent out in that month, 50% from the previous month, and 25% in the month prior to that. For monthly billing basis customers, the rule is 50-50. Certain customers' usage could be identified exactly in the month in which it was sent out. These are typically major power generation companies whose meter is read at the end of a particular month, every month, and the sendout corresponds exactly with billed therms.

The cost of gas module was used to evaluate pipeline purchases and cost out the dollar value of service for a particular accounting period. It also provided a forecast for pipeline purchases to be used in the purchased gas adjustment calculation. At the present time, all pipeline purchases are done on contract, using what is known as a two-part rate schedule. The cost to the utility of a particular quantity of gas is broken down into demand and commodity charges. The demand charge is based on the maximum allowable daily draw of gas from a particular pipeline, whereas commodity charge is a variable charge, depending on the number of units withdrawn from the pipeline. Stiff penalties are also included in the rate. These penalty charges prevent a particular utility from withdrawing a greater quantity of gas than that specified in the demand charge.

The net effect of the two-part pricing forces the utility to withdraw virtually all of the demand quantity gas. Doing this achieves the lowest cost per unit figure. The ratio of actual pipeline withdrawals over the contract demand quantity is known as the load factor. Most utilities operate at a load factor approaching very nearly 100%. By utilizing underground storage facilities, it is possible for companies to maintain a virtual 100% load factor, selling what gas they can to customers and pumping the remainder in or out of storage, as the case may be.

The cost of gas module handled two factors involved with the underground storage. The first, that of inventory, was accounted on a layer by layer basis, using the LIFO method. The model could duplicate the accounting relationship required to perform the LIFO evaluation. However, company-owned underground storage was lumped into one massive underground pool, whereas, in fact, the company, itself, utilized six separate underground storage facilities. The model would take the sendout coming from the revenue section, compare it with pipeline purchases, and compute a net injection or withdrawal figure for underground storage. This figure was compared with guidelines used to establish maximum injection and withdrawal rates from underground storage. Needless to say, however, since all storage fields were lumped into a single storage field, considerable judgment was allowed on the part of the model operator on whether or not such injection or withdrawals were indeed realistic or even physically possible.

One additional feature of the model was its ability to perform the energy and pressure adjustment factors. Most pipeline withdrawals and injection rates were computed in units of thousands of cubic feet of gas purchased, withdrawn, or injected in storage. Gas, however, is normally sold on a therm or energy content basis.

Consequently, the model would adjust gas pressure to normal athmospheric pressure and adjust the energy content based on the number of therms per thousand cubic feet. Although of not great significance, the energy content of gas has varied a few percent during the past several years. Needless to say, the computer simulation model performed this annoying calculation without very much trouble. The purchased gas adjustment model performed an involved calculation used to adjust the cost of gas above a given established cost. If, for example, natural gas from the pipeline cost \$.07 per therm, and \$.05 per therm was the established base price, then the utility was allowed to charge a basic \$.02 per therm PGA or purchased gas adjustment. The PGA allows the public utility to keep pace with pipeline price increases. Historically, most pipelines and utilities have been heavily regulated so that prices, themselves, change only infrequently. However, pipelines have an adjustment factor similar to those in effect by the public utility. Hence, changes in price from the pipeline can occur almost daily and a month does not go by where there is no change in the PGA. The actual PGA calculation, itself, takes the established purchases from the cost of gas module and assumes a 100% factor to the pipeline. That is to say, the consumer of the gas does not become penalized if the utility does not utilize 100% of its available demand gas.

Month to month variations are smoothed out by considering twelve of the past thirteen months. The cost of this period is averaged and compared to the base rate figure. The difference then goes into the PGA rate per therm to be sent to the revenue model. To handle the effect of pipeline increases occuring in midmonth, a spreading factor was used to establish a modified PGA rate. Six separate PGA rates were computed by this model and passed to the Revenue Model. In addition to calculating sendout, the Revenue Model added the previously calculated PGA rate times the number of therms sold to the base dollars of revenue. Because the Revenue Model considered therms billed, it was necessary to allocate, using a spreading formula, the PGA rate over the past several months to compute an average PGA rate for a particular rate and revenue class.

Taxes, such as municipal service taxes, proportional to revenue dollars, were also computed in the revenue model. The revenue dollar figure was then passed over to the Income Model.

The basic purpose of the income module, simplest in this system, was to consolidate the revenue dollars and the cost dollars, adding other factors to get an income and earnings per share and a return on the rate base. The income model served primarily as a report generator.

Income taxes, per share earnings, and return on the rate base were virtually all the calculations performed in this particular model. Many of the factors were left as inputs. Major areas not included in the income model were the capital budgeting, capital expenditure area and the financing area.

IMPLEMENTATION AND VALIDATION

The system of models was implemented in Fortran using the GPOS package of On-Line Decisions, Inc. and runs on a time-shared computer in a highly conversational manner. There were several reasons for choosing this type of an operating environment: the scientific-algebraic nature of the revenue, bill frequency calculations, as well as the availability of a large number of subroutines and subprograms for data referencing in GPOS. Though Fortran appeared to be the best language, programming was not a key factor. Since they were based on an accounting system, the models were relatively simple in terms of discreet events. In accounting, closing the books occurs once per accounting period. The time horizon of the model was five years by months. Consequently, the models, themselves, will run 60 times, once for each month of the five-year planning horizon. Fortran subprograms on the On-Line Decisions' Operating System took care of the time variation in the data.

Time sharing was chosen for two basic reasons: to insure availability of the operating system, and to heighten the degree of interaction required to run a particular simulation model.

The models had very few decision rules programmed into them since they were not optimizing models. The project's goal was to allow middle management to actively interrogate the models to answer "what if" questions. In this instance, the model builders and model users were different individuals. Because of their highly interactive nature, the models were easy to build, but hard to run. The person who ran the models was required to interpret the results and decide on an appropriate course of action, i.e., modify the input data appropriately. Because of the dichotomy between builders and users, considerable effort was put into designing the appropriate interface for the nontechnical user.

The GPOS package handled most of the conversational programming within the model system. To validate the model, it was decided that two years of actual data would be placed into the system and the results compared with actual results. In the PGA area, it was noted that arithmetic errors had been made in certain instances. Once detected, accounting would input reconciling items to offset these previous calculation errors.

Needless to say, to model the randomness of human error making was a difficult task. Provisions were made, however, to include reconciling items in many of the key areas. All told, the process of validation took approximately twice as long as the total programming and implementation phases.

ACCEPTANCE AND USE

After model validation, the model data collectors had to switch hats. Instead of concerning themselves with data collecting and analysis, they now had to consider forecasting and formulation of alternative strategies to the model.

To help gain an understanding of the key and critical relationships, a sensitivity analysis was run over most of the variables within the model system. This analysis involved placing small changes in the input variables and noting the effect on key model outputs.

The GPOS system had this sensitivity capability already programmed within it. After initial forecasts and alternatives had been run through the models, a shift in emphasis in the models began to be observed. For example, since the PGA module duplicated the hand calculations used in the PGA calculations, the module began to be employed by the people within that section to check their own calculations.

In addition, although the model did answer the questions concerning the rate-revenue-alternative source of supply questions, the need for other areas soon became apparent.

While top management had initiated the project, they were not involved in the model construction and validation phase. With the introduction of the working model, "what if" questions and alternative strategies were initially slow in coming. However, usage of the system has averaged approximately 30 hours per month on a connect hour basis. The users fall into approximately three categories - at the top level, the senior financial officer; at the middle management level, the assistant treasurer; at the staff level, within the financial department, various planner analysts.

CONCLUSION

From the time the concept of simulation modeling had met with initial acceptance to the time when the completed modeling system was accepted for use, approximately five months had elapsed. Four men worked almost continously on the project.

The first month of the project was spent performing a feasibility study, specifically indentifying the key areas for modeling, the people who would be involved in the modeling project and the types of "what if" questions that needed to be answered.

Phase II required four months to complete with approximately half of this time spent in technical specification of the modeling system. The specification period went through existing forecasting methods, analyzed data availability techniques to be used as to their accuracy and validity, organized the way in which data would be input to the model, and the report formats coming from the modeling system. The programming phase of the modeling project required about three weeks to complete and validation, almost six weeks to complete. The whole exercise of validation was viewed as a training course for the planner analysts involved in the collection and usage. At the end of the four month period, a course was run to review with middle management and to present to top management the techniques and results of the previous Phase II work.

One year has lapsed since the model was accepted by this utility. The model answers a variety of "what if" questions almost daily. The personnel within the assistant treasurer's staff are becoming known as "keepers of the model," and this staff is being given more and more responsibility in determining and accepting strategies in meeting future operational planning.

Although the model has been in use for an entire year, the model is not frozen. . .a favorable indication. Indeed, a planning model should be dynamic and adjust to changing planning conditions. The true benefits of this simulation modeling system are just now becoming evident. Coal gasification plants are being constructed and arctic pipeline contracts are being negotiated. With the model, utility officers can examine how these new sources of natural gas energy will effect the company's operation and insure the stockholders an adequate return on their investment.

LIST OF ILLUSTRATIONS

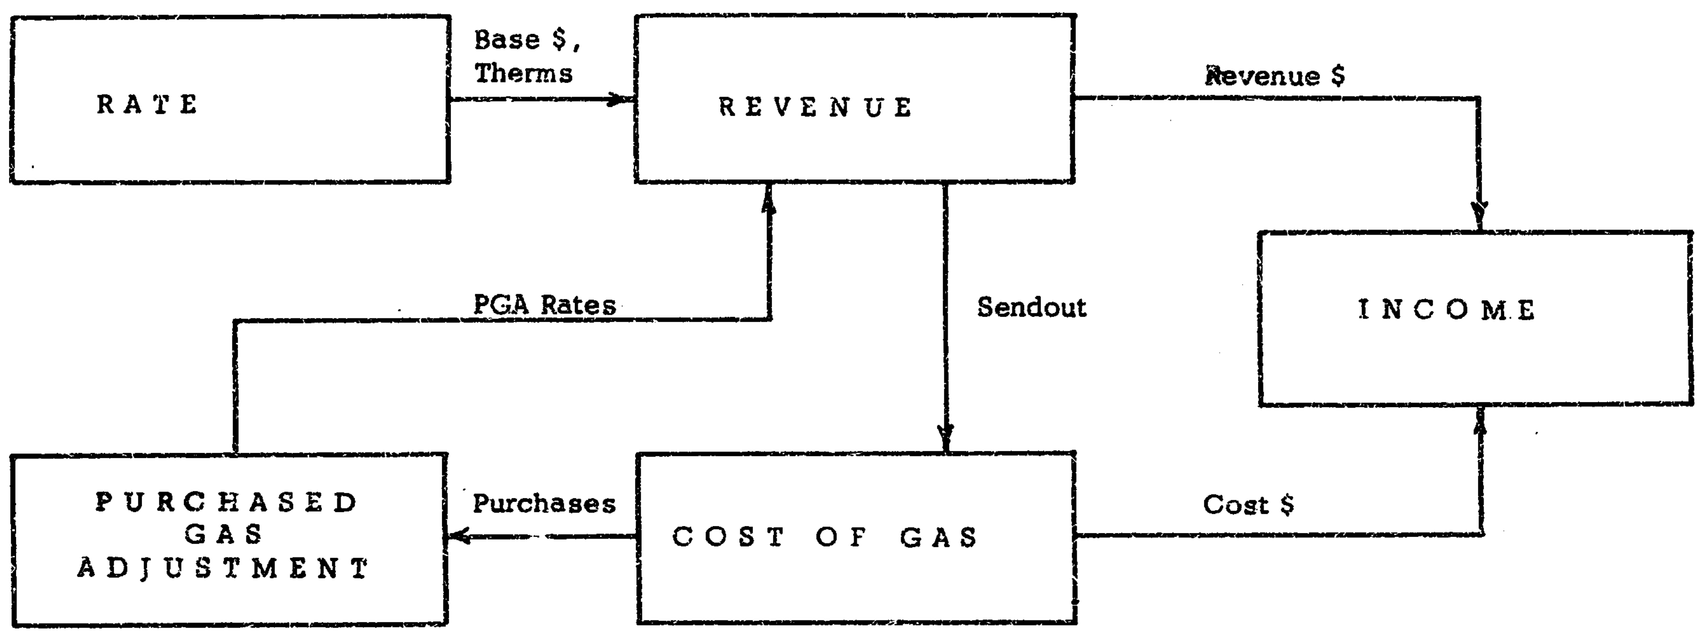
RATE
Base $,
Therms
REVENUE
Revenue $
INCOME
PGA Rates
Sendout
PURCHASED
GAS
ADJUSTMENT
Purchases
COST OF GAS
Cost $

FIGURE 1

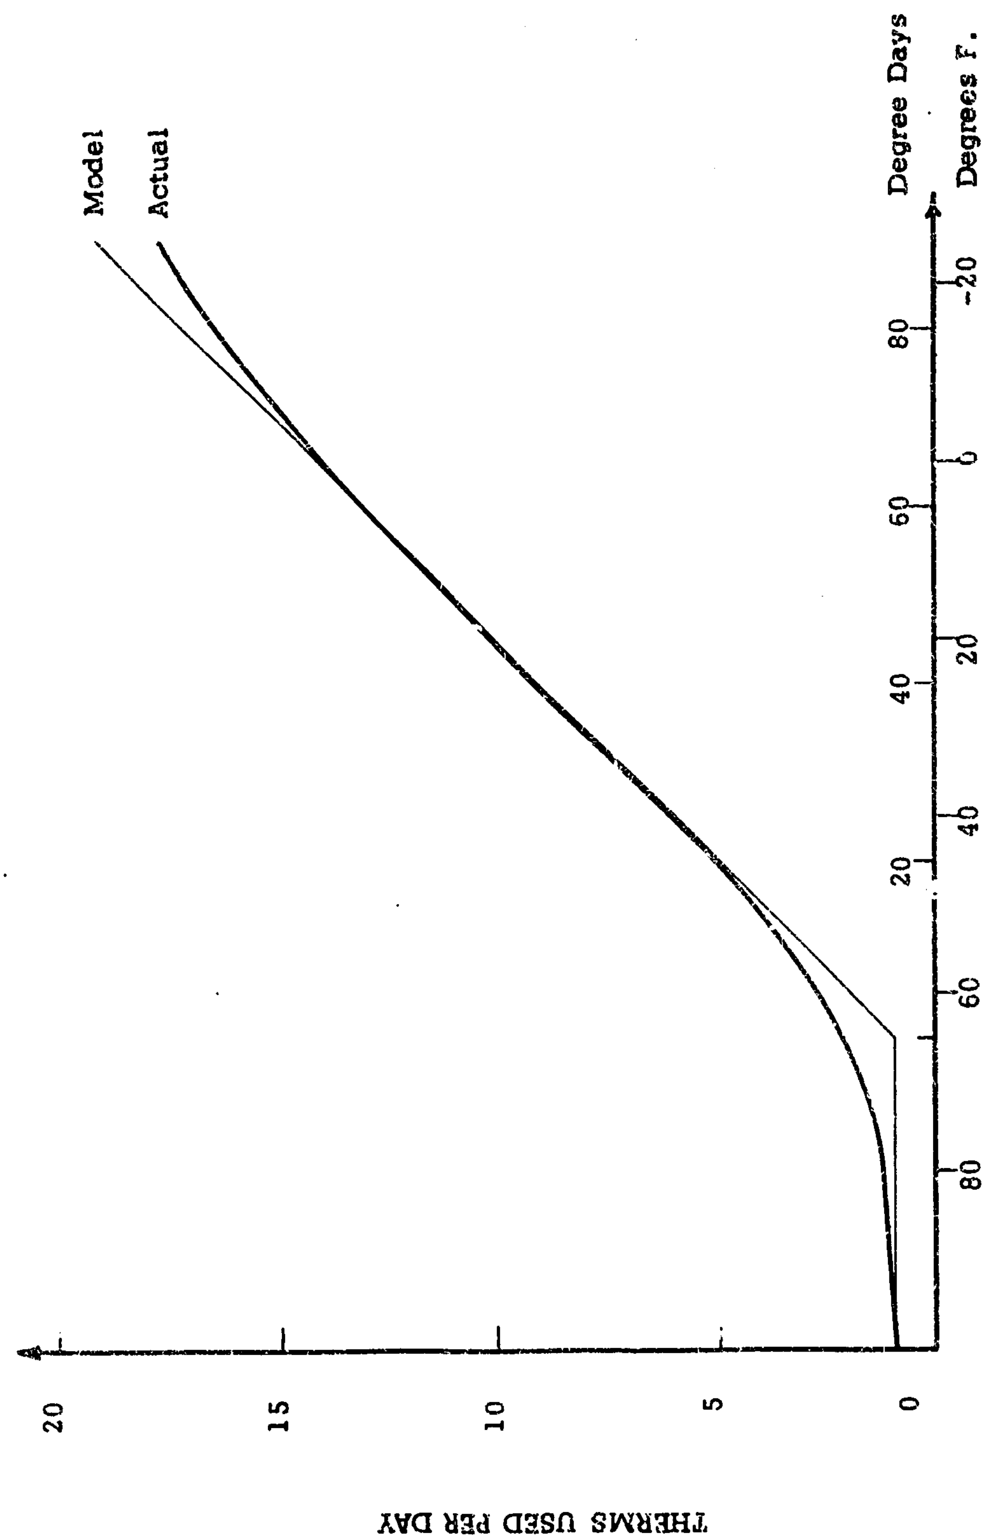

Model
Actual
Degree Days
Degrees F.
THERMS USED PER DAY
0
5
10
15
20
20
40
60
80
80
60
40
20
0
-20

FIGURE 2

FIGURE 3

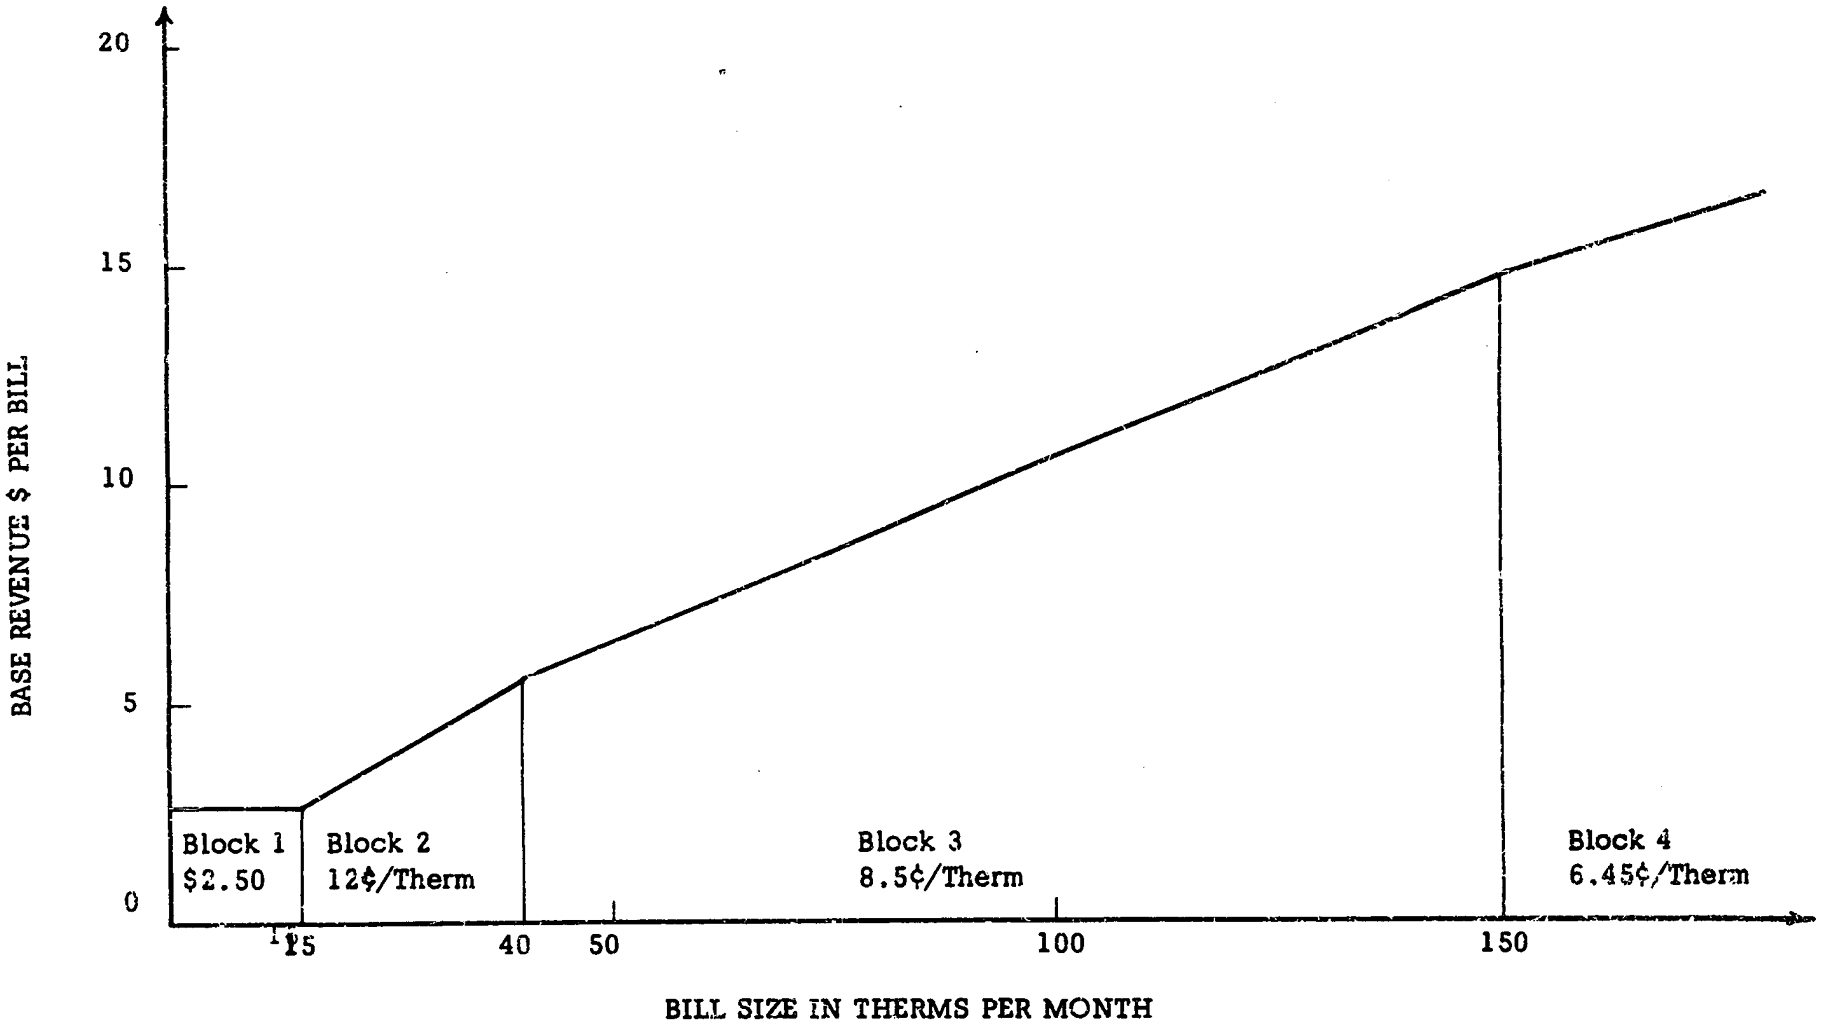
BASE REVENUE $ PER BILL
20
15
10
5
0
Block 1
$2.50
Block 2
12¢/Therm
Block 3
8.5¢/Therm
Block 4
6.45¢/Therm
15
40
50
100
150
BILL SIZE IN THERMS PER MONTH

FIGURE 4

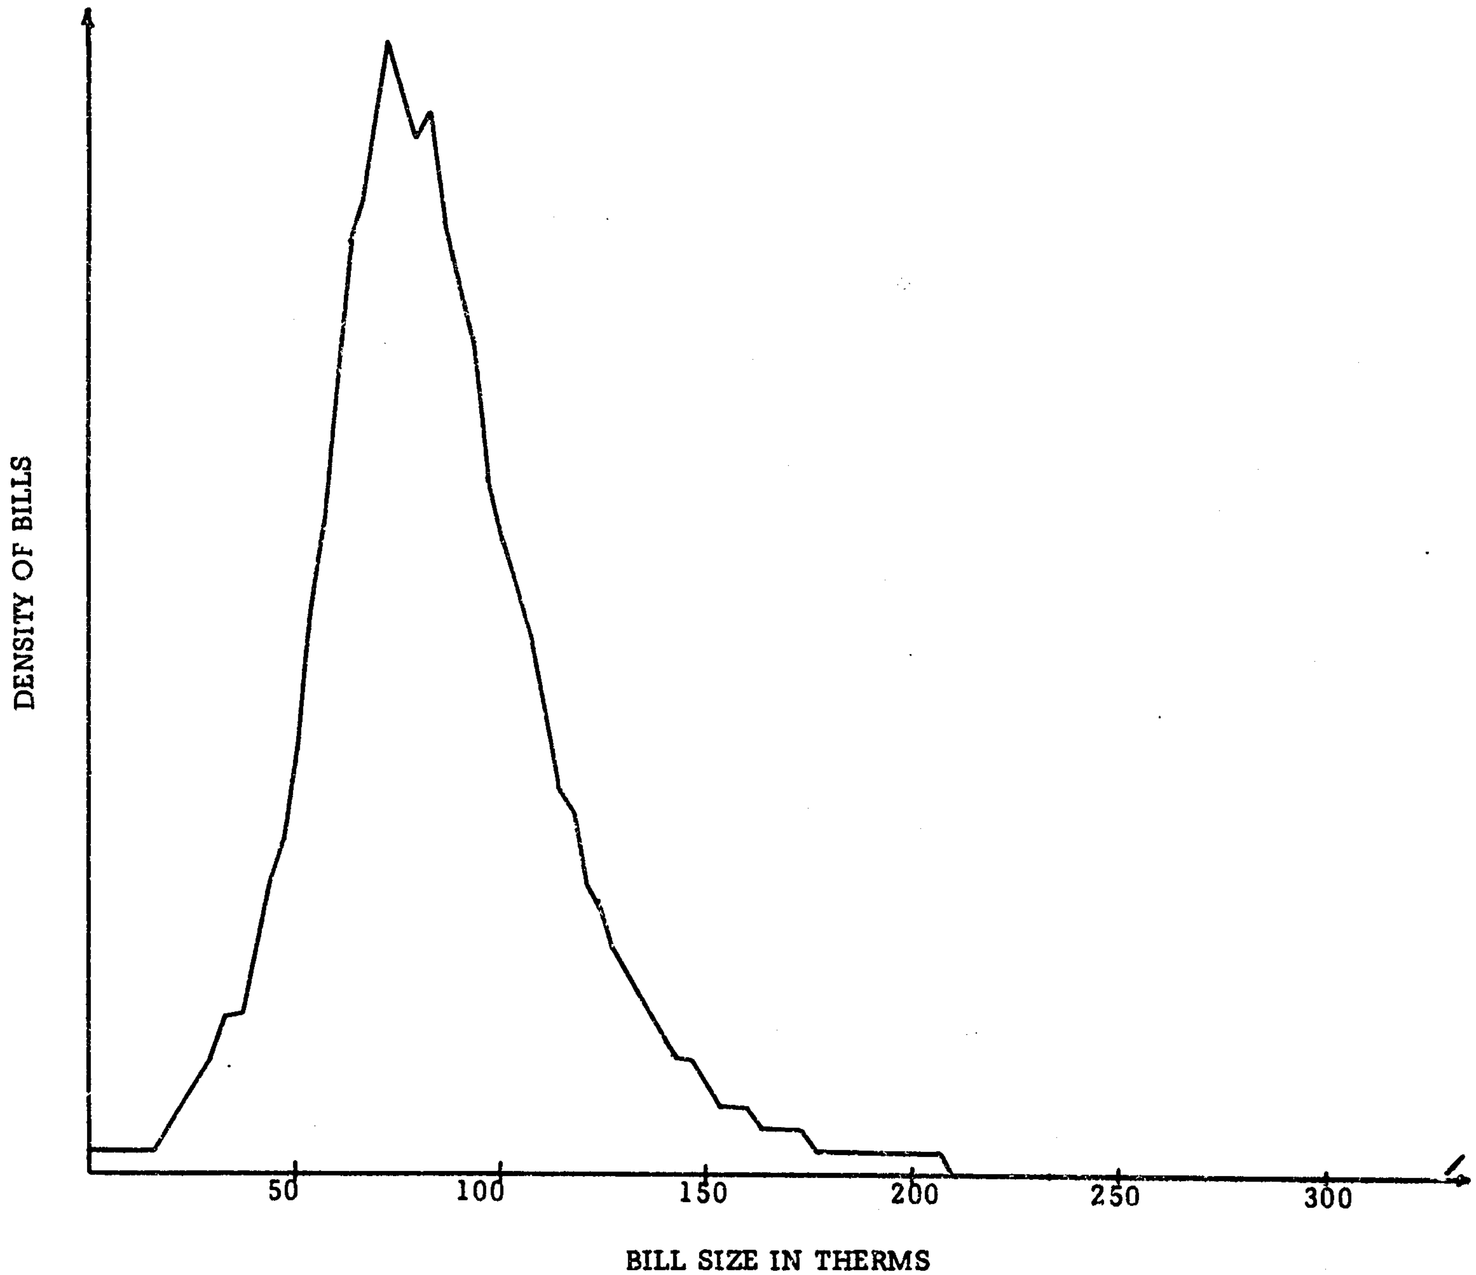

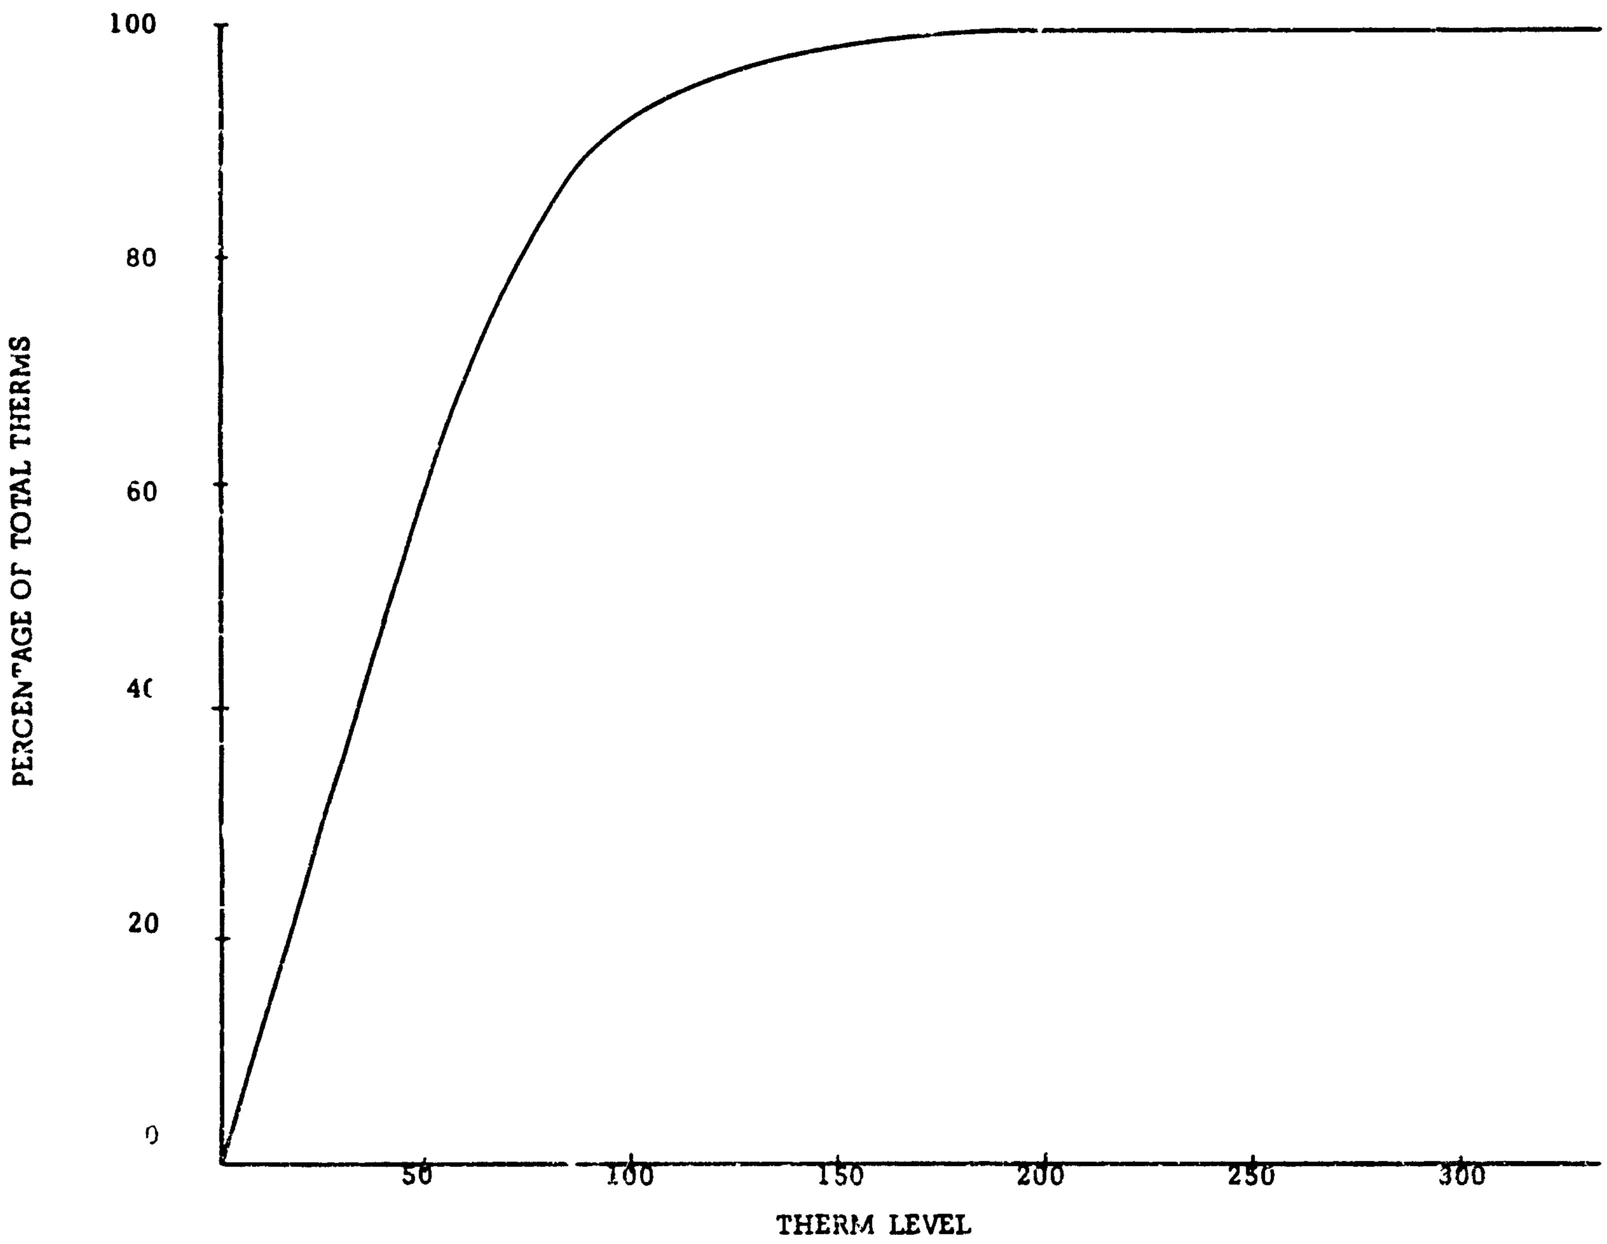
PERCENTAGE OF TOTAL THERMS
100
80
60
40
20
0
50
100
150
200
250
300
THERM LEVEL

IGURE 5

Reference 1 — Ogden, J., "What Happens If - A Planning System for Utilities," Public Utilities Fortnightly, March 16, 1972

Reference 2 — Liittschwager, J. M., "Mathematical Models for Public Utility Rate Revisions," Proceedings of the Conference on Public Utility Valuation and Rate Making Process, Iowa State University, 1969.

Reference 3 — Raiffa, H. and Schlaiffer, R., Applied Statistical Decision Theory, Harvard University, Boston, Massachusetts, 1961, pp. 225-226.