Better distribution planning with computer models

Models are too important to be left to management scientists alone

Arthur M. Geoffrion

Much progress has been made in recent years in the area of computer-based models for distribution system planning. The weakest link is no longer in the tools themselves, but rather in the way these new tools are used. The author conveys a sense of what is now possible and offers some practical suggestions to management. He illustrates many of his general points with a successful study in which he participated at Hunt-Wesson Foods, Inc.

The author is professor of operations research at UCLA's Graduate School of Management. Mr. Geoffrion is also a departmental editor of *Management Science*, an associate editor of *Mathematical Programming*, and a council member of the Institute of Management Sciences. He has served as a consultant to several government agencies and industrial companies, most recently concerning optimization techniques for the design and distribution of production systems.

Computer-based models for distribution system planning have been around for many years. They have undergone a steady evolutionary development and have been advocated relentlessly by the management science community. Yet an honest appraisal of current practice would show that management does far more distribution system planning *without* the benefit of computer models than with them, even in large firms with adequate management science staffs. Why should this be? There seem to be two primary reasons:

- The conceptual design and computational capabilities of such models have often fallen short of what is required for a truly useful decision support tool.
- The task of developing the data called for by such models has often been formidable.

Fortunately, developments in recent years have been reducing difficulties encountered in the past. Experience has taught much about the proper conceptual design of such models; technological advances in management science have vastly improved the computational ability of models to perform as intended; and the data-base movement is making it more economical to develop the data necessary for comprehensive distribution system planning models.

My purpose is to offer some insights that will make it easier for managers to assess the advantages of these developments. For the sake of illustration, I

Author's note: This study was partially supported by a National Science Foundation grant. I also wish to acknowledge the contribution of Steven M. Niino, Hunt-Wesson Foods, Inc., Glenn W. Graves, UCLA, and Shao-Ju Lee, California State University at Northbridge.

shall refer to a successful modeling study carried out at Hunt-Wesson Foods, Inc.

Building the model

Hunt-Wesson Foods, Inc. produces tomato products, cooking oil, matches, puddings, shortening, and many other products at 14 locations (Wesson refineries, Hunt canneries, and copackers). At the start of the study, it distributed nationally through 12 distribution centers. Annual sales were in the vicinity of $450 million and growing fairly steadily. Transportation was by common rail and by both common and contract truck carriers, with extensive use of the storage-in-transit privilege for large rail-supplied customers. The company's policy was to service each customer from a single distribution center for all products. This was done as a convenience for the customer, to simplify accounting, communication, and marketing functions, to permit economies of scale in delivery, and to achieve greater consistency in order cycle times.

A few years ago, the company decided to undertake a planning study because it faced pressing distribution-center expansion and relocation issues. Management recognized that these issues, while seemingly regional in character, were in fact interwoven with the company's entire national distribution system design. It decided to employ a computer-based method that would not only resolve the immediate questions in their proper national perspective, but would also comprehensively rebalance all distribution center locations, assignment of customers to distribution centers, and aggregate annual product flows through the system.

Developing the computer-based model and its associated data required a coordinated team effort which called upon personnel in accounting, data processing, marketing, operations research, production, traffic, and other functional specialties. It took more than one man-year of work over more than six calendar months to construct the model. A simplified summary of the facts and structural assumptions used to build the model is given in the ruled insert on page 97.

The primary outcome of the study was that five changes were recommended in distribution center locations (the movement of existing facilities to different cities and the opening of new facilities). The company implemented the three most urgent changes at the earliest opportunity, as well as improvements in assigning customers to distribution centers. Realizable annual cost savings were estimated to be in the low seven figures.

Basic questions

Solving the basic distribution system planning problem includes answering the following questions for one or more target years, perhaps 1–5 years hence:

- How many distribution centers should there be?
- In which cities should they be located?
- What size should each distribution center be, and which products should it carry?
- Which distribution center(s) or plant(s) should service each customer?
- How should each plant's output of each product be allocated among distribution centers or customers?
- What should the annual transportation flows be throughout the system?
- For a given level of customer service, how do the costs of the best distribution system compare with a projection of the current system to the target period?

Many other related questions may be important, depending on the concerns of the company conducting the study. Manufacturing managers may be concerned about plant capacity planning, financial managers about investment feasibility studies and asset utilization, and marketing managers about product or customer profitability analysis. Distribution management sits in the middle of all this and may be thinking about rebalancing shipping-mode selection policies, inventory stocking policies and customer service levels, consolidation programs, sharing facilities between divisions of the same company or with other companies, and so on.

There is, of course, no such thing as a "final" solution to these problems. Changes in plant capacity, the addition and deletion of product lines, changes

in the economics of warehousing and transportation, shifting markets, and competitors' actions all imply a continuing need to monitor and readjust a distribution system. It is therefore essential to recognize that, at the conclusion of a computer-based planning study, one has a collection of computer programs, data sources, organizational procedures, and know-how which can be used over and over again. They stand as a capital asset, ready to repeat variants of the planning study under revised conditions on short notice at relatively small cost.

The role of the computer

It is important to distinguish between the three basic roles which the computer can play in supporting a distribution system planning project:

1
Preparing data and forecasts pertinent to the planning period.
2
Evaluating particular candidate plans in terms of cost and other quantifiable operating characteristics such as customer service levels.
3
Finding a plan that is best by some quantifiable criterion.

Let's look at each of these roles.

Preparing data
In serious planning studies, large amounts of current data (costs, demands, etc.) must be summarized and projected forward in time. The need for a computer here is inescapable because standard, available data summaries are usually inadequate.

Quite often it is necessary to formulate *conditional* projections or forecasts based on alternative assumptions about the future. Uncertainty about the future is, after all, one of the realities of the planning process. In addition to external influences such as the state of the national economy, management must also take into account possible shifts in internal business policies. A good way to analyze these uncertainties and to clarify the assumptions underlying them is to develop systematic cost/demand/policy *scenarios.*

Each scenario should consist of a complete set of cost, demand, and policy data along with a description of the conditions under which these data are plausible. As many scenarios can be developed as management feels necessary.

Evaluating plans
The term *candidate plan* refers to a possible set of fully specified answers to the planning questions I posed earlier. The second role of the computer is to assist in "evaluating" alternate candidate plans, one at a time, by calculating their costs and other quantifiable results under a given scenario. This is an obvious task for a computer, but one that necessarily stops short of actual choice. Only the planning team is able to judge the relative merits of alternate plans under different scenarios in light of considerations that lie outside the scope of computer evaluation.

Except for the most superficial evaluation of such plans, a computer model typically must calculate in detail the effects on finance, transportation, manufacturing, marketing, and perhaps other business functions. A properly designed model impartially integrates these effects and presents an evaluation in accordance with overall corporate objectives; at the same time, it does not submerge the performance measures that are of interest to those responsible for the individual business functions.

Finding "best" plans
The third and technically most difficult role for the computer is to determine candidate plans that are *best* by some quantifiable criterion (usually cost subject to customer service level requirements). This must always be done within the context of some stipulated scenario. The resulting plans will be examined in light of alternate scenarios and factors outside the scope of the computer model. Some candidate plans will be discarded altogether; others kept and possibly modified.

Using the computer in this way, with different criteria and scenarios, provides a flexible tool and largely relieves the planning team from the tedium of having to guess good plans in full detail; of course, it may still do so whenever such an exercise appears worthwhile. As I shall discuss later, this role also permits both "what-if" questions to be answered and alternative policies to be compared with more precision than would otherwise be possible.

Progress in management science

The ability of a company to employ a computer successfully in its first two roles depends on (a) avail-

ability of data and adequate data processing capability within the firm, and (b) availability of qualified staff personnel to work with senior management. Generally speaking, and without meaning to belittle the very real practical difficulties, the necessary management science technology has been available for many years.

Successful use of the computer in the third role, on the other hand, has been seriously limited by inadequate management science technology except for relatively simple distribution planning problems.

The basic difficulty has been that adequate computational efficiency could be achieved only by imposing stringent assumptions that severaly limited modeling flexibility. This is not to say that companies did not *try* to use a computer to find least cost plans. Many did. The result was usually either an exorbitant computing bill or surrender to expedient but less than fully satisfactory computer programs of an ad hoc nature.

Within the last few years, however, management science has been able to overcome its technical difficulties and now has the ability to achieve *both* a high degree of modeling flexibility and a high degree of computational efficiency. The Hunt-Wesson study provides a nice illustration of this capability. The model was large and complex by prevailing standards, yet a true least-cost plan for a given scenario could be calculated at least 100 times faster than was previously possible.[1] The model's realism and efficiency combined to make it a highly effective decision support tool.

How can a computer-based distribution planning model be used to maximum advantage once it is built? It pains me as a management scientist to admit it, but there can be no doubt about the validity of the conventional wisdom that management science models are too important to be left to management scientists alone. Senior management must actively oversee most of the computer runs made once the model is built.

Using the model

There are three general categories of computer runs with which management need be concerned:

1
Validation runs.
2
Runs designed to develop insights into the company's operations.
3
Runs designed to answer specific questions ("What-if" runs, etc.).

The distinctions drawn here are not intended to be sharp; a single run often serves more than one of these purposes. Nevertheless, I shall discuss the three categories separately.

Validation

The structural assumptions of a model (e.g., the 13 items listed for the Hunt-Wesson model), the numerical data called for by the structural assumptions, and the computer software which manipulates the data all must be validated. The combination of these three distinct entities comprise what in common parlance is called "the model."

Validation of a model should be approached at two levels. At the first level, one tries to verify that the data put into computer memory are consistent with what was *intended*, and that the computer software manipulates the data in a mathematically correct way in accordance with the *intended* structural assumptions. This is a purely technical validation task in which management need not be involved, but emphasis on the word "intended" signals the incompleteness of verification at this level. At a higher level, management needs to determine the extent to which the intentions were actually on the mark.

Management-level validation is concerned with how well the behavior of the model matches the behavior of the organization's system. It is good to begin with a skeptical attitude that views the model as being "on probation" until proven trustworthy. This trust must be earned by solid evidence. What kind of evidence? The traditional approach is to run a case that compares the calculations of the model with history for some previous year. This is done by fixing facility locations at those of a previous year and by using historical cost and demand data from that year. Any discrepancies between history and the

1. See Arthur M. Geoffrion and Glenn W. Graves, "Multicommodity Distribution System Design by Benders Decomposition," *Management Science*, January 1974, p. 822. For a more readable orientation to current management science technology, see my article, "A Guide to Computer-Assisted Methods for Distribution Systems Planning," *Sloan Management Review*, Winter 1975.

model's product flows and total costs should be explainable on the basis of 20-20 hindsight.

This is a valuable check, but it is only the beginning. Additional runs should be made to ascertain that the model behaves reasonably and in accordance with management's understanding of the company's distribution system. Other than testing against historical results, the two most important ways of going about this are to (a) run various highly restricted cases in which the model's scope is so tightly constrained that management can predict what the results should be by inspection or by manual analysis, and (b) run several versions of the same case in which the data are changed in a systematic way for which the impact is at least qualitatively predictable (e.g., increasing all distribution center fixed costs should reduce the number of distribution centers used in the least cost solution).

Validation testing usually leads to adjustments in the data, and sometimes in the structural assumptions as well. As the model stabilizes and proves itself reliable in performing the relatively transparent tasks thus far assigned to it, managerial confidence will rise to the point where a gradual transition can be made to the next and far more exciting model use.

Developing insights

The chief benefit of a model should be its ability to help deepen managers' insights into their system. Among the most important insights to be gained are those that answer these questions:

- How do the major cost categories change as the number of distribution centers changes, and thereby determine the bowl-shaped total cost function that identifies the least-cost number of facilities?
- How do changes in the company's policy regarding customer service influence total distribution system costs?
- How sensitive is the most appropriate system design to the cost and environmental assumptions in which there is significant uncertainty or likelihood of change?

Additional insights may be desirable depending on the particular application at hand.

No sure guidelines have yet been devised to help with the development of such insights. The main reason is that a computer-based model is merely a numbers-in/numbers-out device which tells *what* but not *why*. The "why" must be coaxed out by a comparative synthesis of the results of many runs, each designed to study some variation of a previously run case. The synthesis comes partly from a collation of all these results and partly from knowledge about the system quite separate from the model, namely the combined experience of the management team involved in the project. A recent development which can facilitate this process is the use of highly simplified pencil-and-paper "auxiliary" models designed to explain the general solution behavior (but not the details) of the full model. I believe that this approach could point the way to a general methodology for guiding management's use of a computer model in quest of insights.

If there is any key to success at this phase of a distribution system planning study, it is to constantly ask *why* the model gives the results it does. This attitude has been captured aptly by the numerical analyst R.W. Hamming in his motto: *The purpose of computing is insight, not numbers.*

Answering specific questions

A common mistake is to virtually ignore the development of insights and to pass directly from validation to using the model the way in which it is so obviously well suited: providing specific numerical results. The pitfall here is that genuine understanding is a prerequisite for intelligent evaluation of numerical results produced by what is, after all, only a simplified representation of the company's distribution system.

If validation and development of insights have been thorough, this final phase will be extremely valuable to the management team as it hammers out its conclusions. Many of the computer runs needed will be of the "what-if" type. For instance:

- What if a certain uneconomical distribution center is forced to remain open as desired by a district marketing manager or a certain large customer?
- What if the relative economics of transportation modes should change substantially because of increasing fuel costs or because of federal regulatory policy changes?

Summary of the model's facts and assumptions

1.
The many hundreds of Hunt-Wesson's different products are aggregated into 17 product groups. Nearly all products in a given group are similar with respect to the plants which can make them, the production technology by which they are made, and their gross-to-net shipping weight ratios. All quantities are measured in gross shipping hundredweight (cwt.).
Examples of the product groups are: bottled cooking oil, packaged shortenings, ketchup, tomato sauces, puddings, and matches.

2.
The production capacity (cwt./yr.) of each of the company's 14 plants is given for each product group.
Not every plant could make every group – in fact, only about 50 plant-product group combinations were possible.

3.
Distribution center locations to be evaluated by the model are limited to a list of 45 cities composed of current locations, major metropolitan demand concentrations, the locations of major competitors' distribution centers, and other gateway cities.

4.
The size of each distribution center chosen by the model must be between stipulated lower and upper limits expressed in terms of total annual throughput volume.
The lower limit was due to the need to accommodate rail shipments of about eight boxcars, which permits important economies of scale.

5.
The many thousands of individual customers are aggregated into 121 customer zones defined in terms of zip codes.
The zones consist of a key city with a significant demand plus (usually) adjacent areas with relatively small associated demand. Each zone is sufficiently small in geographical area so that the best distribution center for the key city is also the best one for the other customers in the zone.

6.
The annual demand forecasted for each product group in each customer zone must be met.
The demand forecasts were based on historical summaries prepared by computer from magnetic tape archives of shipment invoices. These tapes also included most of the information required to develop a realistic transportation cost structure.

7.
Each customer must continue to be serviced from a single full-line distribution center for all products.

8.
No distribution center location is allowed to service a customer zone unless it is sufficiently close to the zone that normal delivery time would be less than a specified maximum number of days.
This maximum is the model's primary control over the level of customer service.

9.
Differential production costs ($/cwt.) are given when there is a significant difference between unit costs at different plants for a given product group.
Not to be overlooked among the reasons for differential production costs is the influence of plant location when it is significant in determining delivered raw material costs (as it was here in the case of refinery products).

10.
The costs of each distribution center are expressed as a "fixed" charge if a location is selected for use by the model plus a "variable" cost ($/cwt.) applied to the total annual throughput volume over all product groups.
These costs included all handling, storage, and other operating expenses, inventory carrying charges, and amortized capital costs associated with expanding an existing distribution center or acquiring a new one. A standard throughput cost curve like the hypothetical one shown in the figure was developed based on historical analysis of the current distribution centers' costs. This curve was then modified for each of the 45 locations to reflect geographical differences. A straight-line approximation proved adequate, although a more accurate broken-line approximation would have been used if necessary to portray economies and diseconomies of scale over the range of interest.

11.
For each product group and possible plant-to-distribution center transportation link, a unit freight rate ($/cwt.) is specified.
A judgment was made in each case whether truckload or rail carload lots were the most appropriate mode of supply.

12.
For each product group and possible distribution center-to-customer zone link, a total annual freight cost ($) is specified.
It made sense to specify an annual cost for a link, since each customer had to receive all of a product group from a given distribution center if he received any at all. Most of the delivery costs were generated by computer from specially constructed approximation tables, taking into account the proportions of mode (LTL, TL, CL) and shipment size likely to prevail in each case.

13.
Transit rail rates are used whenever they are advantageous.
The storage-in-transit rate applicable to some product groups is figured as though the shipment moves directly from plant to customer with a small charge for stopping over at the intermediate distribution center (so long as this stop is not too far off the direct line), instead of as the sum of the regular plant-to-distribution center and distribution center-to-customer rates.

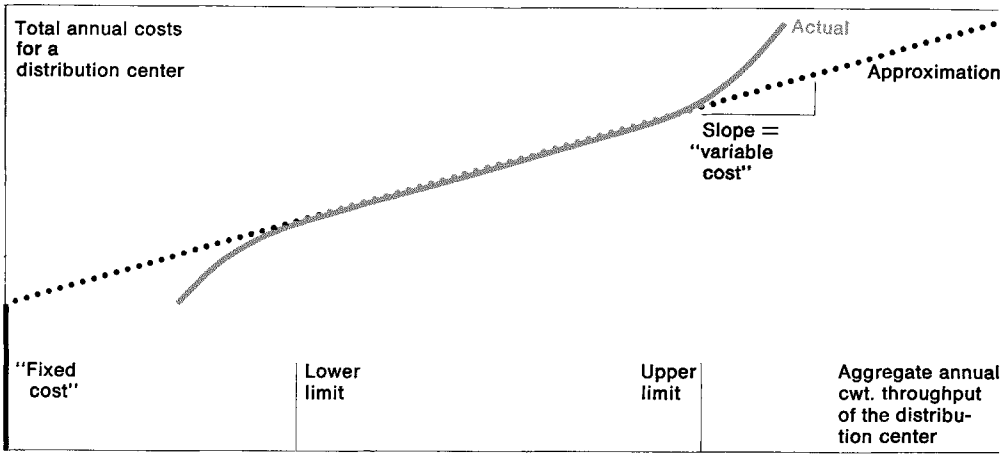

□
What if a strike should knock out a key facility?
□
What if a cooperative arrangement could be worked out with another firm to share portions of the distribution system?
□
What if a certain regional product line goes national?
□
What if certain products are dropped from certain warehouses?
□
What if division X is allowed to have its own distribution network independent of the corporate network?

Other runs may address the effects of expected demand pattern shifts over time or the effects of changing assumptions made in the model. The objective is to obtain a feeling for the "robustness" of the results yielded by the model. In the case of Hunt-Wesson, for instance, a mature market position and relatively stable business environment caused the best system configuration to be quite stable over the planning horizon. On the other hand, it was found that the lower throughput limits specified for the distribution centers (see item 4 in the ruled insert) had a surprisingly noticeable influence.

Balancing objectives

Still other runs may enable management to strike the best balance between cost and noncost performance measures.

Consider, for instance, the balance between total distribution cost and customer service level. It is known that improving customer service almost always increases costs at an increasing rate, but it is not possible for the model to choose a proper customer service level because no one knows how to quantify accurately the influence of customer service level on future sales. The model can, however, do the next best thing. By varying the model parameter that determines the level of customer service, a series of runs can generate a trade-off curve which gives the least total distribution cost for various possible levels of customer service. Deciding where the firm should operate on this curve poses the customer service issue as crisply as one could hope.

Addressing related issues

Plant capacity planning is just one of many secondary distribution-related topics that can be addressed effectively by using a distribution planning model. What would be the effect on the best facility configuration and total distribution cost if a plant's production capacity for a certain product group is expanded or reduced? If a product group could be made at a plant which could not make it before? If a new plant is built in a certain area?

The study team should be alert to opportunities to derive such additional benefits from a model.

Setting priorities

Among the last runs to be done are those addressing implementation priorities for various desirable changes. It is appropriate to study the impact of these changes individually and jointly.

Generally those which prove to be responsible for greater improvements in overall system performance should be put into effect before those which yield smaller improvements. The model may show that some changes are of such small value that it would be wise to postpone their implementation or at least to hold them in abeyance until it is possible to restudy them using data that has been updated. Implementation priority studies are necessary because of practical limits on capital expenditures and organizational changes.

In the Hunt-Wesson study, this analysis focused almost entirely on the impact of distribution-center location changes on total system costs, since this was the most important determinant of capital requirements and organizational dislocations. The most economical reconfiguration required six location changes. These six changes were examined individually and in various combinations. In each case, it was necessary to reoptimize all service areas and transportation flows so that the revised network of locations would be used as efficiently as possible.

Two changes emerged as important enough to warrant immediate implementation. A third change was considered less beneficial but was also adopted because it gave additional warehouse space in a geographical region where more space was particularly needed. Two other changes were recommended for reconsideration at a future time, and the sixth change was so marginal that it was dropped.

2. See Robert J. Atkins and Robert H. Shriver, "New Approach to Facilities Location," HBR May-June 1968, p. 107 and Harvey N. Shycon and Richard B Maffei, "Simulation—Tool for Better Distribution," HBR November-December 1960, p. 65.

Making it happen

There are a few points for management to keep in mind when initiating a computer-based planning study of the type I have described, particularly when shopping for suitable management science support from internal staff or external sources. It is not necessary to nose very deeply into questions of technology assessment provided management insists on three assurances.

Plain English: The first is that a plain-language statement of the current version of the model will be available at all times. Mathematics and computer jargon are unnecessary to communicate the structural assumptions from which a model is fabricated, as my 13-point outline of the Hunt Wesson model illustrates. A plain-language statement is a prerequisite for a manager to interpret the results. It also gives some protection against technical legerdemain and exaggerated claims of model "realism."

Inexpensive and easy to use: The second assurance needed is that the computer software should be inexpensive enough to run and easy enough to use that dozens of major runs can be made, if necessary, within the time and resources available. The need for numerous runs is plain from my discussion of the many ways in which a model should or can be used. Some of the older technologies require several days of preparation between runs and/or computer costs in excess of $1,000 per run.

Optimizing capability: Third, assurance should be obtained that the computer program will have a true optimizing capability. That is, it should be able to sift through the multitude of alternate plans and find one which is best according to an appropriate criterion (usually minimum total cost to achieve desired customer service levels for a given scenario). This is in contrast to a program that cannot be guaranteed to find the best choice—either because it has no mechanism for generating alternate candidate plans, or does have such a mechanism but of such a design that the search is likely to be incomplete. Perhaps the term "extremizing" would be better than "optimizing," since the latter unfortunately connotes that the computer's choice should somehow also be the "correct" one to adopt.

An optimizing capability allows far more useful work to be done in each computer run than would otherwise be possible. Without it, the project team must carry an extra burden of drudgery that can only result in less complete validation, fewer insights, and less extensive use of the model to address specific issues of concern.

Not only does an optimizing capability enhance the value of most individual runs, but it also provides the opportunity to make valid *comparisons* between the results of different runs. This is an extremely important consideration because the conclusions reached by a planning project typically rely far more heavily on comparisons between computer runs than on runs considered individually. With "quasi-optimizing" programs, such as so-called cost calculators or simulators fitted with heuristics, one never knows whether different results are due to different inputs or to the vagaries of the computer program.

Substantial progress

Five years ago, it would not have been realistic for management to call for the second and third of these assurances. That it may do so today is a consequence of the substantial progress that has taken place in the field of management science since the days of earlier influential contributions reported in these pages.[2]

In addition to progress in management science technology, there has also been steady progress in the availability of data sources. The rapid development of corporate data bases is of great importance to distribution system planning. More companies are coming to treat their operating data as an asset rather than a nuisance, and to organize it in a manner that permits more convenient access than in the past. Data sources available outside the company are also improving. Specifically, freight rates—by far the biggest data development chore in most planning studies—can now be purchased en masse from at least two large commercial data bases. As more data can be obtained from reliable machine-readable files, the costs of data collection drop dramatically and computer-based distribution planning studies become more attractive.

Hence advances in data sources and in management science during the last decade have substantially reduced the costs of building distribution planning models at a realistic level of detail.