Caution: Common Sense Planning Methods Can Be Hazardous to Your Corporate Health

Arthur M. Geoffrion University of California, Los Angeles
Tony J. Van Roy Katholieke Universiteit Leuven

A nontechnical presentation is given of the pitfalls awaiting those who attempt to use techniques based on common sense methods or heuristics for planning. The context of discussion is taken to be distribution planning, but the same pitfalls arise for all other planning contexts involving decisions of major consequence. *Ed.*

The central task of planning is the search for permissible decision alternatives that are best according to some specified performance measure(s). Most planning efforts today depend on common sense approaches for the generation and evaluation of decision alternatives. Sometimes these approaches are computerized, in which case they usually are called *heuristics*. The aim of this article is to show that such approaches, whether computerized or not, can be at once very plausible intuitively and yet grossly misleading. We argue that such approaches should be used with extreme care or, better still, abandoned in favor of a dependable approach as discussed in the final section.

Three simple examples are given to demonstrate typical pitfalls associated with common sense methods and heuristics. The particular context selected is distribution planning: questions about the proper number, size, and location of plants and distribution facilities; the proper assignment of customers to shipping points; which products to make or stock where; policies concerning inventory and transportation; and so on. Similar pitfalls await planners in other functional areas of the firm.

A few words are appropriate concerning the examples. They have been designed not so much for realism as to illustrate concisely the general points to be made concerning possible malfunctions of common sense methods. The reader should keep in mind that, while the examples are very small and hence susceptible to solution by brute force enumeration given sufficient perseverance, brute force is woefully inadequate in most realistic situations. Some formalization of common sense then becomes necessary to avoid the hopeless task of complete enumeration.

Approaches like those we shall describe are widely employed in practice. Experienced planners will perceive much that is tauntingly familiar in the simplified approaches and problems to follow.

First Example

A wholesaler marketing in four regions would like to distribute through two public warehouses to be selected from four candidate locations. Taking into account the storage and handling costs at each location, the common carrier freight rate from each location to each customer region, and the projected annual volumes, the total annual cost (in thousands of dollars) for each possible customer/warehouse assignment is given in Table 1.

Which two warehouses will result in the lowest total annual cost? A complicated procedure hardly seems necessary for such a simple question. One sensible approach (attributed to F. E. Grange by one book,[1] but invented and used independently by many others) works as follows.

Suppose all four warehouse sites are "open." What would be the least cost assignment of customers to warehouses? The obvious answer is to use the cheapest assignment for each customer. This requires

Table 1

Customer Region	Warehouse Sites A	B	C	D
1	110*	640	670	450
2	585	65*	590	115
3	165	200	125*	840
4	595	580	115	100*

Arthur M. Geoffrion is Professor of Management Science and Director of the Management Science/ Operations Management Study Center at the Graduate School of Management, University of California, Los Angeles. Holding the Bachelor of mechanical and Master of industrial engineering degrees from Cornell University and the Ph.D. degree in operations research from Stanford University, Dr. Geoffrion is the recipient with Glenn Graves of the 1976 NATO Systems Science Prize. He is a member of the Operations Research Society of America, the Institute of Management Sciences, the Association of Computing Machinery, the Mathematical Programming Society, and the National Council for Physical Distribution Management. In addition, he is currently coeditor of the Mathematical Programming & Networks Department of *Management Science*. Dr. Geoffrion is the author of *Perspectives on Optimization: A Collection of Expository Articles*.

Table 2

Customer Region	Warehouse Sites A	B	C
1	110*	640	670
2	585	65*	590
3	165	200	125*
4	595	580	115*

scanning for the smallest entry in each row, and results in best assignments as indicated by the asterisks in Table 1.

With the best 4-warehouse solution in hand, it is simple to find the best 3-warehouse solution. If site A is closed, customer 1 must be reassigned to its second choice, site D, at an additional cost of (450 − 110) = 340. No other reassignments are necessary because no other customers have site A as their first choice. In a similar manner, one can easily evaluate the added cost of closing each of the four warehouses one at a time:

closing site A adds a cost of (450 − 110) = 340
closing site B adds a cost of (115 − 65) = 50
closing site C adds a cost of (165 − 125) = 40
closing site D adds a cost of (115 − 100) = 15.

Since closing site D incurs the least added cost, the best 3-warehouse system is {A, B, C}. The assignment cost table for the best 3-warehouse system, with the best assignment for each customer indicated by an asterisk, is shown in Table 2.

A repetition of the same reasoning enables us to see which one of the three sites to drop so as to incur the least added cost. The calculations are:

closing A adds (640 − 110) = 530
closing B adds (585 − 65) = 520
closing C adds (165 − 125) + (580 − 115) = 505.

Hence C is dropped, leaving {A, B} as the best 2-warehouse subset of the best 3-warehouse system. This result is displayed in Table 3.

Thus we have applied Grange's eminently sensible method to the problem of finding the best 2-warehouse system for the example represented by Table 1. The method makes use of the convenient fact that, in the absence of any warehouse capacities or other complications, the jointly best assignment of customers to warehouses is always to assign each customer to its first choice (least cost) open site. The best 4-site system was obvious. The best 3-site system was determined easily from the best 4-site system. Finally, we found the best warehouse to drop from the 3-site system to obtain a 2-site system.

Looks good, doesn't it? Table 3 surely gives the best 2-site system, right? WRONG! Dead wrong. The best 2-site system is given by {A, D}, not {A, B}. It has a total annual cost of 490, considerably less than the total annual cost of 920 indicated by Table 3.

Common sense has let us down.

Second Example

A manufacturer has three plants, each capable of supplying one-third of the national demand, and three warehouses, each handling one-third of the national demand. Taking into account all transportation and warehousing costs and the volumes involved (but excluding production costs), the total annual cost in thousands of dollars for each possible plant/warehouse assignment is given in Table 4.

What set of plant/warehouse assignments will result in the lowest total annual cost? Remember that each plant can supply only one warehouse.

A plausible approach to this problem is to assign each warehouse *in turn* to the remain-

Table 3

Customer Region	Warehouses A	B
1	110*	640
2	585	65*
3	165*	200
4	595	580*

Tony J. Van Roy holds the M.A. degrees in electro-mechanical engineering and in industrial management from the Katholieke Universiteit Leuven. As a part of the Ph.D. program sponsored by CIM-Belgium, Mr. Van Roy spent 1978-1979 at the Graduate School of Management, the University of California, Los Angeles. He is experienced in production scheduling and distribution systems planning through several projects in industry. Mr. Van Roy's main interests are operations research and mathematical programming and their applications.

Table 4

Warehouse	Plants A	B	C
1	25	85	75
2	80	40	120
3	75	115	225

ing plant that can serve it least expensively. This results in the following set of assignments:

warehouse 1 served by plant A at a cost of	25
warehouse 2 served by plant B at a cost of	40
warehouse 3 served by plant C at a cost of	225
	290.

What would happen if, instead of assigning warehouses sequentially to plants, plants are assigned in turn to warehouses? It so happens that exactly the same set of assignments would result.

Is this the best set of assignments? No. The alert reader will have noticed that the high cost assignment of warehouse 3 to plant C could be avoided by interchanging the assignments of warehouses 2 and 3:

warehouse 1 served by plant A at a cost of	25
warehouse 2 served by plant C at a cost of	120
warehouse 3 served by plant B at a cost of	115
	260.

The initial approach was too naive in that it failed to anticipate the future consequences of the early assignments as they were made sequentially; the first two assignments (1 → A and 2 → B) forced a bad last assignment (3 → C).

A more sophisticated approach would be to make assignments not in some arbitrary order, but rather, in an order which anticipates the future consequences of current assignments by a "look ahead" calculation that takes account of second best choices. Consider again the assignment of warehouses to plants. If warehouse 1 does not get it's first choice (A), then it's second choice (C) is worse by 75 − 25 = **50**. Similarly, the second choice cost penalty for warehouse 2 is 80 − 40 = **40**, and for warehouse 3 it is 115 − 75 = **40**. Since warehouse 1 has the highest penalty for not getting its first choice, assign it first (to plant A). That leaves warehouses 2 and 3 for plants B and C. The second choice cost penalty is 120 − 40 = **80** for warehouse 2 and 225 − 115 = **110** for warehouse 3. Hence assign warehouse 3 next to plant B. That forces 2 → C. Happily, this more sophisticated way of selecting the order of assignment yields the improved solution given earlier with value 260.

Exactly the same "look ahead" method could be used to assign plants to warehouses instead of warehouses to plants. Again the improved solution with value 260 is generated.

If we merge the assignment of plants to warehouses and warehouses to plants with the second choice cost penalties guiding the sequence of assignments, we then have a still better approach known as "Vogel's Method" for problems of this general type.[2] Vogel's Method yields the same improved solution with value 260.

By now the evidence is pretty strong that the solution with value 260 is in fact the best possible. One way to check this belief is to see whether any pairwise interchange of assignments leads to a reduction in cost. There are three such interchanges:

interchange A and C	1 → C at a cost of	75
	2 → A at a cost of	80
	3 → B at a cost of	115
		270
interchange B and C	1 → A at a cost of	25
	2 → B at a cost of	40
	3 → C at a cost of	225
		290
interchange A and B	1 → B at a cost of	85
	2 → C at a cost of	120
	3 → A at a cost of	75
		280.

They are all losers.

Figure 1

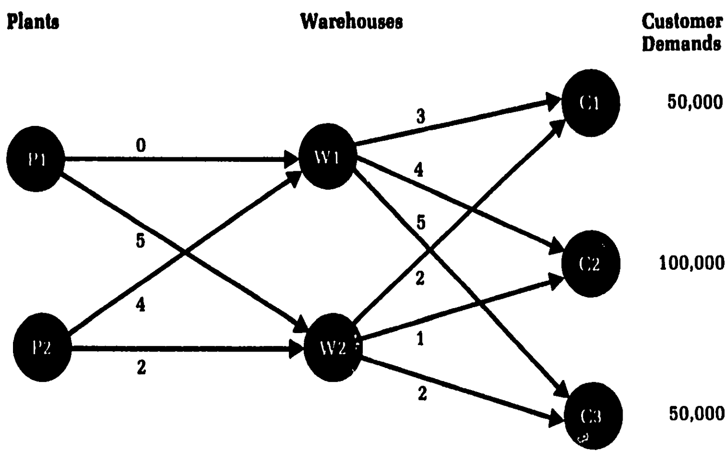

Is 260 in fact the lowest possible cost? NO! The lowest cost set of assignments is:

1 → C	at a cost of	75
2 → B	at a cost of	40
3 → A	at a cost of	75
		190.

Common Sense, even taken to the evolutionary culmination of Vogel's Method with pairwise interchange checks, is foiled again.

Third Example

Consider a company with two plants, two warehouses, and three customer groups. Figure 1 bears annotations showing freight rates ($/CWT) and annual demands (CWT). In addition, the annual capacity of plant P2 is limited to 60,000 CWT. Storage and handling charges at the two warehouses are identical and hence can be ignored. For a similar reason, unit manufacturing costs at the plants wash out. What should the annual transportation flows be through this system so as to minimize total annual transportation costs?

Clearly it would be naive to determine either the outbound or the inbound transportation flows in isolation from one another. For instance, serving C1 from W2 is cheaper than from W1 considering only outbound costs, but W1 seems preferable when one takes into account that the least inbound rate to W2 is $2/CWT worse than the least inbound rate to W1 (which happens to be a plant warehouse for P1).

The standard common sense approach to two-stage problems like the one posed is to specify the outbound flows first, but in such a way that the upstream consequences are taken into account. A natural way to do this is first to calculate the least total cost flow path (plant → warehouse → customer) for every customer, and to assign each customer to the warehouse on this path. This results in demands being placed on the warehouses. The inbound flows can then be specified to meet these demands in the marginally least cost manner. This general approach is used in several commercially available, proprietary distribution planning packages.[3]

Let us apply this approach to the problem at hand. The least cost path for C1 is P1-W1-C1 (it beats P2-W1-C1, P1-W2-C1, and P2-W2-C1). This suggests the assignment of C1 to W1. Similarly, the least cost path for C2 is P2-W2-C2, and for C3 it is P2-W2-C3. Thus the outbound flows will be:

50,000 CWT from W1 to C1 at OB cost	$150,000
100,000 CWT from W2 to C2 at OB cost	100,000
50,000 CWT from W2 to C3 at OB cost	100,000
Total Outbound Cost	$350,000.

The inbound flows must now be determined so as to fill the annual warehouse requirements of 50,000 CWT at W1 and 150,000 CWT at W2. The best inbound flow pattern obviously is (remember the capacity limitation on P2):

50,000 CWT from P1 to W1 at IB cost	$ 0
60,000 CWT from P2 to W2 at IB cost	120,000
90,000 CWT from P1 to W2 at IB cost	450,000
Total Inbound Cost	$570,000.

The total annual cost for this solution is **$920,000** ($350,000 for outbound plus

$570,000 for inbound). The solution seems to be very satisfactory because

— Every customer is assigned to the warehouse on its least cost path from plant-to-warehouse-to-customer;

— Every warehouse uses its least cost plant (up to the full available capacity).

Moreover, this solution is $200,000 better than the naive one of first specifying the outbound flows without regard for upstream costs.

By now your faith in common sense analysis must be sufficiently shaken to make you suspicious of the quality of this solution in spite of its intuitive appeal. Can you find the transportation flows which lead to a total annual cost of only **$740,000**?

How Can Common Sense Methods and Heuristics Fail So Badly?

Some of the reasons for the possible failure of even the most plausible methods can be gleaned from a study of the examples:

1. Facility location procedures which "drop" facilities one by one (as in Example 1) are apt to arrive at a poor answer because, although each drop may be very good or even best in isolation, there is no guarantee that the *combination* of drops will be good. Procedures which "add" facilities one by one have the same vulnerability. Unfortunately, it is usually out of the question to enumerate all possible combinations in practical problems (e.g., there are more than 155 million ways to select fifteen sites from among thirty candidates).
2. Procedures which specify decision choices sequentially can run into trouble because they fail to anticipate fully the future consequences of each choice as it is made. In Example 2 we saw that a sequential "greedy" assignment of warehouses to plants was not good because the first two assignments forced a very poor third one. A modified approach which "looked ahead" one step was better, but still did not solve the problem correctly. Unfortunately, it is usually totally impractical to look ahead far enough to be sure of making the best decision choices sequentially.
3. "Local improvement" procedures can stop short of the best solution because they lack the global perspective necessary to know when they have come to a dead end where a radical change is necessary to make further progress. An apt analogy would be a boat seeking to find the deepest spot in the Caribbean Sea using a sonar depth-finder; there is always a possibility that looking in an entirely different part of the Sea would produce spots still deeper than the result of following the local downhill gradient. The set of assignments with value 260 in Example 2 was such a dead end because no pairwise interchange of assignments could improve on it; a triple interchange was necessary to achieve an improvement, and for other examples a still higher order interchange is necessary.[4]
4. One part of a system can sometimes have a subtle influence on another part of the system. This is illustrated by Example 3, which shows that it is dangerous to deal with one distribution echelon (say, outbound) separately from another (inbound) even if an attempt is made to take account of the influence of one on the other.
5. Capacity limits of any kind are particularly troublesome. This is what spoiled the otherwise rational approach of Example 3.

In a nutshell, common sense approaches and heuristics can fail because they are *arbitrary*—arbitrary in the choice of a starting point, arbitrary in the sequence in which assignments or other decision choices are made, arbitrary in the resolution of ties, arbitrary in the choice of criteria for specifying the procedure, arbitrary in the level of effort expended to demonstrate that the "final" solution is in fact best or very nearly so.

The result is erratic and unpredictable behavior—good performance in some specific applications and bad in others. It is a dicey game.

If such approaches can fail so badly for highly simplified examples, it seems reasonable to expect that the kinds of failures possible with problems of realistic size and complexity can only be richer in variety and harder to spot.

Implications for Planning

Given the dicey nature of common sense methods and heuristics, does it necessarily follow that they should not be used as planning aids?

For instance, might they not luckily succeed in finding improved decision alternatives? The roll of the dice might indeed yield a substantially improved decision alternative. Suppose for the sake of argument that a heuristic were available which could find solutions with an average error of only 3%. Three percent sounds quite reassuring. But, consider that, at least in the area of distribution planning, most existing systems are already between 5% and 15% of the ideal. This estimate is based on firsthand experience with numerous real distribution planning problems and is in line with the experiences of others. Thus an average error of 3% would overlook somewhere between **20%** (3/15) and **60%** (3/5) of the possible gains from a planning study. From this perspective, a 3% error is not so small.

A much more profound weakness of heuristics in the context of planning arises from the *critical need to "solve" planning problems under several alternative sets of assumptions* or, in popular parlance, to ask "What if . . . ?" questions. The reasons for this need are several:

1. The proper responsibility of a planning team is not just to come up with good recommendations, but also to provide *convincing justification* for those recommendations. Merely arguing the sophistication of the planning methodology is not sufficient. Proper justification normally should include evidence concerning the sensitivity of the leading recommended alternatives to particular assumptions which are in doubt. Doubtful assumptions arise in all planning studies by the very nature of having to deal with the future: changing demand patterns, shifting inflation rates by cost category, etc. A comparative study of the performances of the leading decision alternatives under different possible futures is also in order.
2. Working simultaneously with more than one set of assumptions helps to reveal *why* certain decision alternatives are better than others. The managerial insights which arise in this way can easily be of more value in the long term than any of the specific recommendations arising from a planning effort.
3. The ability to compare solutions reliably under different sets of assumptions vastly enlarges the scope of planning issues which can be dealt with usefully. This idea has been elaborated upon at length elsewhere.[5]

If heuristic procedures are risky under a single set of assumptions, they are many times more so when results are to be compared under multiple sets of assumptions. Unfortunately, *taking differences between error-prone answers tends to magnify greatly the errors.*

This fact can be illustrated simply as follows. Imagine that you are given your choice of two rough cast gold ingots, each approximately one kilogram in weight. The true weights are A grams and B grams respectively, but these numbers are unknown to you. The only scale available has a random error range of ±15 grams. You weigh the ingots and obtain the readings $\hat{A}$ and $\hat{B}$. The difference, $\hat{A} - \hat{B}$, is your only estimate of which ingot is heavier. Now since $\hat{A} = A \pm 15$ grams and $\hat{B} = B \pm 15$ grams, $\hat{A} - \hat{B}$ will be $(A - B) \pm 30$ grams. Since $(A - B)$ is small, the ±30 gram error range will tend to swamp the true weight difference. Thus if A and B happen to be exactly identical, the scale could yield any conclusion from "A is 30 grams heavier than B" to "B is 30 grams heavier than A." If A is actually 20 grams heavier than B, the scale could yield any conclusion from "A is 50 grams heavier than B" to "B is 10 grams heavier than A."

Clearly any measuring instrument used to measure *differences* must have a small error range relative to the differences to be measured.

The third example discussed earlier is revealing in this regard. Let's ask: "What if the freight rate from W2 to C1 is decreased from $2/CWT to $0.5/CWT?" Surely this will result in a lower total annual cost. Yet reapplying the approach described for the original version of the third example yields a new solution with a total annual cost of $1,045,000. Compare this with the original cost of $920,000 produced by the same heuristic approach. In other words, the nonsensical conclusion emerges that *reducing* an outbound freight rate can result in *higher* total cost!

It is our repeated observation in practical planning studies that important total cost differences to be measured for "What if" questions often fall in the range of 0.1% to 1%. This implies the need for a technique whose error range is provably small relative to this already small range. The kind of heuristics in common use today seldom meet this criterion.

See the Appendix for a do-it-yourself exercise which simulates the kinds of frustrations arising in practice when trying to use heuristics for "What if" questions.

Is There a Better Way?

Yes, of course there is: *optimization.* Optimization does not fall prey to any of the pitfalls that can cause other methods to stop short of a best answer. It not only finds a best answer, but also *proves* that it is best. The weaknesses pointed out in the previous section are completely absent.

This is not to say that optimization is a universal panacea for the litany of ills recited earlier. It is not. There are drawbacks, some real and some imagined. One is that optimization may be difficult or even impossible to achieve for the planning problem at hand. Fortunately, great strides have been made in optimization technology over the past decade, and the prospects for continued progress are bright. Many planning problems have been brought within reach, including quite a few for which heuristics were the only practical option just a few years ago. Progress has been quite rapid, for instance, for distribution planning problems of the type which inspired the three numerical examples given earlier.[6]

Another potential drawback of optimization stems from its relatively mathematical nature; management may feel uncomfortable using a tool whose inner workings are not thoroughly understandable to them (heuristics at least have that virtue). This view, we submit, is based on a fallacy: it is not the technical mechanism of a decision aid that needs to be understood, but rather the *function* it performs. How ironic it is that heuristic techniques tend to have a *poorly* understood function (their solutions are of unknown quality), while optimization techniques have a very *well* understood function (their solutions are of precisely known quality). Would a manager select a TV set based on the simplicity of the electronics inside? More likely he would judge in terms of function (picture quality).

Still another drawback, in the minds of some, is the view that the available data may not be good enough to justify optimization. This view presumes that optimization is necessarily more expensive than heuristics. The surprising truth for many specialized planning problems is that optimization is considerably *less* expensive. And, even if optimization should be more expensive, the added cost may well be justified by the importance of the planning function and the comparatively large cost of data development necessary regardless of whether an optimizer or a heuristic is used. This view also conveniently ignores that management must analyze and plan on the basis of those same available data, and thus an optimal method can only be preferred as an analytical tool. None of the weaknesses of common sense methods and heuristics pointed out in the previous section of this article are mitigated by diminishing the quality of the data.

As a footnote to this discussion of the pros and cons of optimizing versus heuristic methods, a clarification is in order concerning hybrid methods. This class might be called "partial optimization" in that it applies an optimizing technique to some limited (relatively simple) aspect of the planning problem, and a common sense approach or heuristic to the remainder of the problem. An example would be when a heuristic is used to locate facilities and then, given the resulting set of locations, an optimizing technique like linear programming is used to find the best transportation flows. Empirical evidence shows that such an approach can yield quite poor results.[7] The Achilles heel of partial optimization is that it is also partially heuristic.

In conclusion, we wish to emphasize that the foregoing discussion does not imply that heuristics have no future or that they should necessarily be replaced by optimizers whenever possible. We ourselves cheerfully admit to publicly advocating and professionally applying heuristics in other problem contexts, and we also find them useful as starting techniques for optimization methods.

There are three main situations where heuristics must be taken seriously:

a. Applications where optimization is not yet practical under the current state-of-the-art of management science for a model of the desired degree of realism;

b. Applications where repeated decisions must be made, none of which are so major as to require extensive scrutiny; and

c. Applications where sophisticated heuristics are available with provably tight error bounds.

One may have little choice but to use a heuristic in situation (a), as the only alternative would be to resort to an oversimplified optimizable model. But the user must be on guard against possible poor results and must avoid, if at all possible, the dangerous temptation to make comparisons between "solutions" obtained with such a tool under different sets of assumptions. If such comparisons are unavoidable, management should exercise great caution in drawing any conclusions.

In situation (b), a heuristic need only pass the test of yielding better answers *on the average* than the leading alternative approach; over many applications, it is not the quality of any *one* answer that matters but rather the *average* quality of the answers. A good example which illustrates both (a) and (b) is vehicle scheduling. Optimization for such problems is usually prohibitively expensive, but techniques that are even 5% suboptimal on the average are preferable to more traditional methods that are (say) 10% suboptimal.

Situation (c) has only recently materialized for a few applications of a relatively simple nature. So far as the authors know, all of the heuristics used today for distribution planning, for instance, do not have any error bounds at all (!) and may yield arbitrarily bad solutions.

Since many planning problems fit neither situation (a) nor (b) nor (c), it follows that optimization rather than heuristics should be used in such cases.

To do otherwise could be hazardous to your corporate health.

Appendix

The Dicey Game of Heuristics

The use of heuristics is a dicey game. The simple do-it-yourself experiments given here will help bring to life some of the consequent practical difficulties. The reader can perform these experiments using any convenient source of "random" numbers.

Preliminaries. How large an error will be made by a heuristic depends on the "luck of the draw," and the odds are usually unknown. The odds have, however, been measured experimentally on occasion by researchers for specific techniques applied to specific classes of problems. The Table of Odds shown here was obtained empirically for a specific heuristic and a sample of fifty medium-sized problems similar to the first example except that the best number of warehouses was also to be determined (in the presence of equal fixed costs for all warehouses).[8]

The Table of Odds shows, for instance, that the odds of a perfect answer (0% error) are 34%, while the odds of 4½% error are 12%. The average error calculates out to be 2½%. To simulate the luck of the draw, enter the Table of Odds according to any random number between 00 and 99. For instance, a random number of 68 would yield an error of 3½%. There are 100 possible random numbers in all, and they have been placed in correspondence with each % error possibility in exact accordance with the odds (e.g., exactly 6 random numbers correspond to 3½% error).

Table of Odds

% Error	Odds (%)	Corresponding Random Numbers
0	34	00–33
½	6	34–39
1½	16	40–55
2½	8	56–63
3½	6	64–69
4½	12	70–81
5½	6	82–87
6½	6	88–93
7½	2	94, 95
8½	2	96, 97
12½	2	98, 99

First Experiment. Suppose that a shipper has an annual distribution bill of $10,000,000 and that the ideal distribution system has an annual cost of 5% less than this, or $9,500,000.

The average error of a heuristic with this odds table calculates out to be 2½%, which translates for this problem into an expected annual cost of $9,500,000 × 1.025 = $9,737,500 for the resulting approximate solution. But the actual approximate solution which comes out depends on the luck of the draw.

To simulate the luck of the draw four times, we construct four 2-digit, pseudo-random numbers from the last 8 digits of the first author's home telephone number:

(2 1 3) 3 9 4 - 4 6 5 5)
a b c d

The following results obtain.

Random Number	Corresponding % Error	Approximate Solution
33(a)	0	9,500,000
94(b)	7½	10,212,500
46(c)	1½	9,642,500
55(d)	1½	9,642,500

These results are typical in that they are drawn according to the Table of Odds for this problem. The user never knows what result will materialize.

Second Experiment. Suppose now that the same shipper wants to evaluate a change of transportation policy from policy P to policy Q. In truth, *policy Q is better than policy P* because the annual cost of an ideal distribution system under policy Q (instead of P) is

Table 6

Policy	Random Number	Corresponding % Error	Approximate Solution	Conclusion
P	38(a)	½	9,547,500	Q beats P by $247,500
Q	25(b)	0	9,300,000	

Table 7

Policy	Random Number	Corresponding % Error	Approximate Solution	Conclusion
P	15(c)	0	9,500,000	P beats Q by $218,500
Q	81(d)	4½	9,718,500	

$9,300,000 (instead of $9,500,000). But of course this truth is unknown.

We therefore present the question "*What if transportation policy P is changed to Q?*" to the heuristic. The first author's office, rather than home, telephone number will be used this time:

(2 1 3) 8 2 5 - 1 5 8 1.
a b c d

Table 6 shows the results obtaining when a and b are used. If c and d are used, on the other hand, the results are as shown in Table 7.

The correct conclusion was reached in the first draw (although the savings were overestimated by $47,500), while a dramatically wrong conclusion was reached in the second draw.

This is typical of the instability of common sense methods and heuristics when used to address "What if" questions.

Reader's Experiment

The same shipper wants to evaluate Marketing's proposal to raise the service level parameter in the highly competitive Northeastern region from an in-stock rate of 93% to 95%. Suppose that, in truth, this would raise annual distribution costs from $9,500,000 to $9,675,000.

Twice simulate the application of a heuristic to the question: *"What if the service level is increased from 93% to 95%?"*

Fill in the following two tables.

First Draw

Service Level	Random Number	Corresponding % Error	Approximate Solution	Conclusion
93%				
95%				

Second Draw

Service Level	Random Number	Corresponding % Error	Approximate Solution	Conclusion
93%				
95%				

How do your conclusions match the actual cost difference of $175,000? Do you feel comfortable using an approximate procedure with the previously given Table of Odds on "What if" questions like this?

The authors gratefully acknowledge the constructive comments of R. F. Powers.

References

1
See R. E. D. Woolsey and H. S. Swanson, *A Quick and Dirty Manual* (New York: Harper & Row, 1975), pp. 150–151.

2
See N. Reinfeld and W. Vogel, *Mathematical Programming* (Englewood Cliffs, NJ: Prentice-Hall, 1958).

3
For example, DPM (see F. H. Mossman, P. Bankit, and O. K. Helferich, *Logistics Systems Analysis*, University Press of America, 1977) and LOGISTEK (see H. N. Shycon, "The Computer as a Tool in Planning the Distribution System." Paper distributed at the 1978 Annual NCPDM Conference, Chicago, October 1978).

4
The number of possible interchanges grows more rapidly with the order of the interchange than intuition suggests. For example, an assignment of size 10 (e.g., assign 10 warehouses to 10 plants as in Example 2) has 45 pairwise interchanges, 600 triple interchanges, and 4,830 quadruple interchanges.

5
A. M. Geoffrion and R. F. Powers, "Facility Location Analysis Is Just the Beginning," *Interfaces*, in press.

6
For example, see:
A. M. Geoffrion, "A Guide to Computer-Assisted Methods for Distribution Systems Planning," *Sloan Management Review*, Winter 1975, pp. 17–41;
A. M. Geoffrion, G. Graves, and S. Lee, "Strategic Distribution System Planning: A Status Report," in *Studies in Operations Management*, ed. A. Hax (New York: North-Holland, 1978).

7
For example, see J. Krarup and P. M. Pruzan, "Selected Families of Discrete Location Problems," (Institute of Datalogy, University of Copenhagen, Report No. 77/7, June 1977), pp. 48–50.

8
See A. S. Manne, "Plant Location under Economies-of-Scale—Decentralization and Computation," *Management Science* 11 (November 1964): 213–235.
The computational approach is heuristic in that it adds as well as drops warehouses one at a time in quest of a least cost configuration. The approach also falls under the "partial optimization" category because, for each trial configuration of warehouses, the optimal flows are computed.