LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

ADAPTIVE CONTROL SYSTEMS IN MARKETING

342-68

John D. C. Little

June 1968

AUG 21 1968
DEWEY LIBRARY

Paper given at American Marketing Association Meeting, June 17-19, 1968, Philadelphia, Pennsylvania

This work has been supported in part by a grant from the Marketing Science Institute and by Project MAC, an MIT research program sponsored by the Advanced Research Projects Agency, Department of Defense, under the Office of Naval Research Contract No. Nonr-4102(01).

HD28
.M414
no. 342-68

Dewey

MASS. INST. TECH.
SEP 10 1975
LIBRARIES

RECEIVED
SEP 4 1968
M. I. T. LIBRARIES

ABSTRACT

In conducting marketing operations, a company must adapt its actions to changing market conditions. The problems involved are discussed in general and then a specific adaptive system is studied. The system is concerned with setting promotional spending in a situation where company sales depend on promotion, but the relationship changes with time. The adaptive system works as follows: Information about sales response is collected by performing a field experiment. The information is used to update a sales response model. Promotion rate is then chosen to maximize expected profit. The cycle is repeated In designing the experiment, sample size is chosen to minimize the cost of imperfect information plus the cost of experimentation. A numerical example is studied analytically and by simulation. The adaptive system is found to work better than various other policies. In a sensitivity analysis, an adaptive system derived for one underlying model of the market is found to perform well even when certain other models actually apply.

About the author: John D. C. Little is Professor of Management, Head of the Mathematical Methods Group, and Head of the Marketing Group at the Sloan School of Management, Massachusetts Institute of Technology.

S 25935

1. INTRODUCTION

Today I want to talk about an idea that is obviously right. It is this: A company should learn from its experience in an organized way.

Consider, for example, the job of the marketing manager. Either formally or informally he must go through a certain cycle of operations. First, he must assemble marketing information. For consumer goods this means finding out about consumers, their needs, attitudes, and buying habits. It means finding out about retailers, what they carry, their prices, their allocation of shelf space, and their display practices. The marketing manager must further determine what the competitors are doing in product, price, and promotion. He must gather data on communications activity such as advertising and sales effort. He must obtain internal company information on shipments, cost of goods sold, marketing costs, production capacities and the like.

Second, the marketing manager must use the information to modify his conception of the market. How will the market respond to a given new product? a new type of promotion? an increase in advertising? a cut back in trade allowances? Either explicitly or implicitly the manager must have a feeling for how the market works. He uses the information he has collected to develop or revise that feeling.

Third, the manager uses his conception of how the market works to determine marketing actions. He develops the marketing mix for a new product or modifies that of an old one. He adjusts budgets, changes packaging, rejects old promotions and introduces new ones. He reallocates sales effort among his products. Such changes are made under the assumption

GENERAL BOOKBINDING CO.
QUALITY CONTROL MARK

that they are improvements. In turn, such an assumption is based on his perceptions of the market as delivered by his information system.

Finally, the manager *installs another round of information collection*. He continues successful measurement techniques and tries out new ones. He runs consumer surveys, obtains store audits, and taps into the company's internal record keeping stream.

Clearly, we are describing a system for controlling marketing activities. It is an adaptive system since new data is collected from the environment and used to modify actions. We shall find it helpful to develop a more formal view of the system. This is sketched diagrammatically in Figure 1.

Here we have called the assemblage of marketing information a *data bank*. The conception of the market developed from the data is called a *model* of the market. On the basis of the model *decision making* takes place. This involves setting values for marketing variables and choosing courses of action from various sets of alternatives. The next step is the *design of measurements* which, along with the various chosen courses of action, determine what happens in the *market*. The results are *sales*, *profit* and a variety of marketing information, which are inserted into the data bank. The new data is used to update the model and continue the feedback process.

2. LITERATURE

Obviously every company has some procedure for determining its marketing actions and the procedure must include many of the above elements. The need for adaptive systems has frequently been recognized in a qualitative way. For example, Robinson and Luck [1] have an eight step "Adaptive

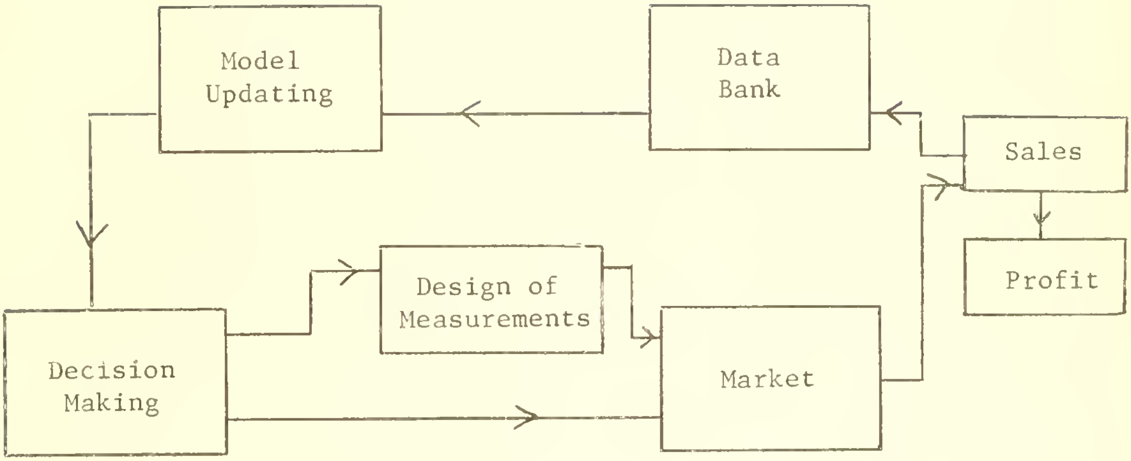
Model Updating
Data Bank
Sales
Profit
Design of Measurements
Decision Making
Market

Figure 1. An Adaptive Control System for Marketing Activities

Planning and Control Sequence." In such qualitative structures, however, the relationships between data inputs and decision outputs are not at all formally specified. Our interest is in quantitative structures for which we can study possible inputs and possible relations to determine their effect on overall company performance. We are particularly interested in systems that are designed to handle changing marketing conditions.

In the quantitative realm a good deal of work has been done in the area of adaptive models for forecasting. See for example Harrison [2] and the bibliography in his paper. Forecasting involves data collection and updating, but in most situations it does not explicitly involve decision making, i.e. the criterion for goodness of a forecast is usually the closeness between predicted and actual values without reference to the decisions that are to be made on the basis of the forecast. Consequently we do not think of the usual company forecasting process as a formal adaptive control process. However, from a mathematical point of view, a criterion is a criterion, whether it be a dollar loss or predicted minus actual, and so prediction can be viewed as a problem in optimal control. As a result the literature on forecasting and on adaptive control are intermixed.

The new product area has seen the development of a number of models which can be called adaptive. This is not surprising since a company obviously must learn as it goes along during the development and introduction of a new product. The decision tree analysis of Green [3], which deals with the size and timing of production facilities for a new product, is one example. Urban [4] and Charnes et al. [5] have developed new product decision models which have adaptive aspects. In each case a series of

decision points is encountered by the firm. At each point one of three choices is made: (1) to go ahead with the product, (2) to stop it, or (3) to collect further information. The decision is made on the basis of the accumulated information to date. A review of these models is given by Montgomery and Urban [6].

Little [7] has developed an adaptive control model for controlling promotional spending in the case of an ongoing product. Compared to the new product models mentioned above, his model is more restricted in scope but much more explicit in operating details. It will be discussed more fully below. Fitzroy [11] has worked on different but similar models as have Hurst [12] and de Andrade [13].

A large body of mathematical work has grown up in the field of optimal control. In cases where the underlying processes have unknown elements that must be learned as the process unfolds, optimal control problems may be called adaptive Much of the work to date has been stimulated by engineering problems but the potential applications are far wider. Typical of work that may eventually have relevance in marketing is that of Sworder [8] on optimal adaptive control systems and that of Aoki [9] on optimization of stochastic systems. Little's model falls in the general class of optimal control models

3. EXAMPLE: ADAPTIVE CONTROL OF PROMOTIONAL SPENDING

To my knowledge there is nowhere today in operation a comprehensive formal system of the type shown in Figure 1. However, at least one company I know of does something quite similar to the model about to be described. They use the process to set their annual advertising budget. The company

does not go through all of the mathematical manipulations of the model but the sequence of operations is essentially the same.

The model is briefly as follows. Company sales are a function of promotional spending, but the relationship changes with time. Information about sales response is collected by performing a field experiment. The experimental results are used to update a sales response model. On the basis of the model promotion rate is then set to maximize profit in the next time period. The cycle is then repeated.

In order to design the control system we must specify sales and profit models plus a change mechanism for the sales response model, find a way to set the optimal promotion rate, and determine how to design the experiment. In the discussion below we describe each of these steps verbally. For a mathematical description, see [7].

3.1 Sales and Profit Models. We suppose that sales response to promotion at a specific point in time has the general shape of the diminishing returns curve of Figure 2. (Specifically, we assume that sales response in the neighborhood of the national average promotion rate is reasonably well represented by a quadratic function of promotion rate.)

Profit is given simply by a gross margin times the sales rate less the promotion cost and less any fixed costs.

3.2 Change of sales response with time. If the constants in the sales response function were known, it would be a simple matter to find the promotion rate that maximizes profit (for example, by plotting profit against promotion and picking off the highest point.) However, the constants are presumably hard to measure and so we ordinarily expect to end up with a less than ideal promotion rate and therefore incur a loss relative to perfect information.

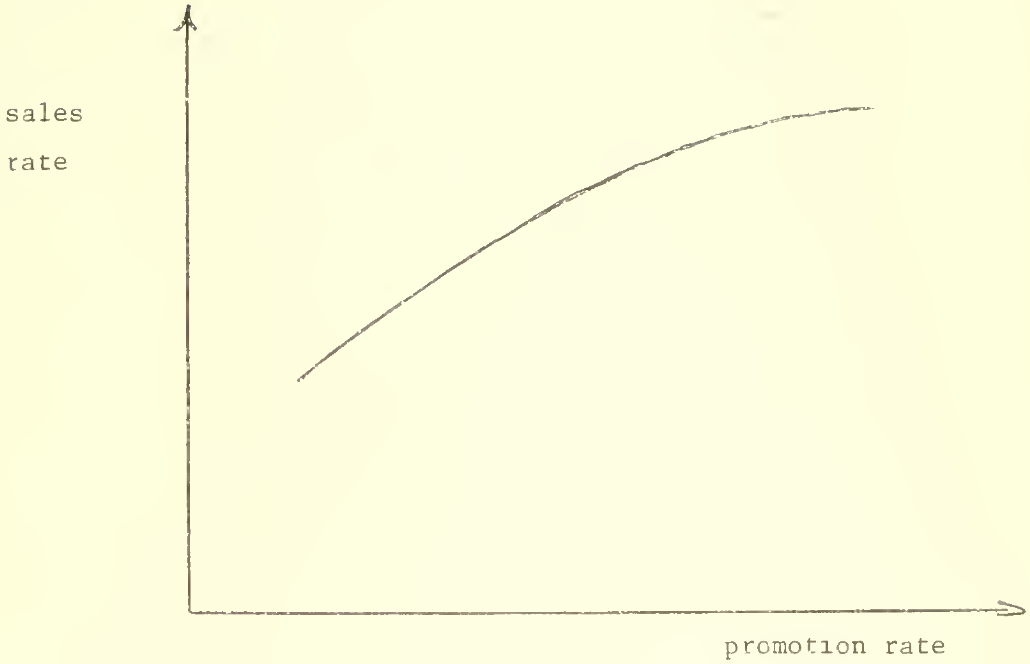

Figure 2 Sketch of possible sales response to promotion

If sales response did not change with time, the company should put a big effort into measuring it right away, because the extra profit from increased accuracy would last indefinitely. However, it is difficult to believe that in practice sales response stays the same. For example, competitive activity, product changes, changes in the quality of promotion, and shifts in economic conditions lead us to expect shifts in response. As a result, an expensive effort to learn the sales response curve of the moment cannot be justified. On the other hand, if the constants of the curve change fairly slowly with time, some effort is worthwhile and, in fact, very important.

A specific change mechanism must be assumed if we are to analyze the system. The assumption we shall make is that one of the constants in the sales response function has a random disturbance added (or subtracted) from it in each time period. (To be precise, we assume that the coefficient of the linear term changes.) The net effect is that even if the company holds its promotion rate constant, sales will change with time. The changes might appear as in Figure 3. Here time is treated in discrete periods, say of six months or a year.

3.3 Measurement by field experiment. For a measuring device, we shall think in terms of direct sales experiments. That is, different groups of customers are given different experimental treatments and the effect on sales is observed. This has now been done many times in practice. For a discussion of some of the techniques involved, see, for example, Hoofnagle [10]. Other, less direct devices for measuring sales response are, of course, possible and in some situations might be preferable.

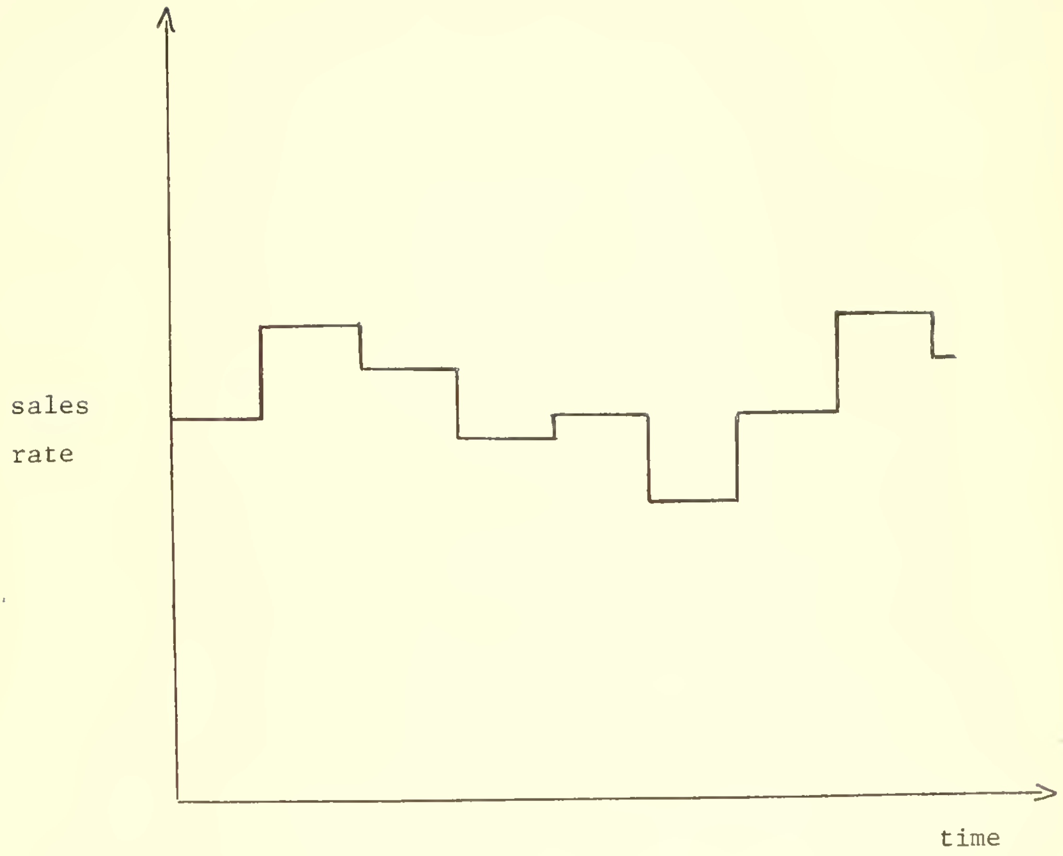

Figure 3. Sketch of how sales rate might change with time under constant promotion rate because of unanticipated market disturbances.

The experimental setup is sketched in Figure 4. A group of market areas, say n of them, is operated at a promotion rate higher than the national average and another group of n is operated lower. The resulting differences in average sales rates between these two groups is used to estimate the unknown constant of the sales response curve.

The precision of the measurement depends on the sample size, n, on the spread between the high and low promotion rates, and on the inherent variability in sales from market area to market area. Sample size and spread are design parameters of the experiment. Market to market variance is more or less a fact of life, except that we can frequently measure and remove some of the extraneous factors that influence sales. This will increase precision. Quite a lot can be done along these lines using the theory of experimental design, analysis of covariance, and well thought out data collection.

The high and low promotion rates are defined relative to the national rate, which normally can be expected to change from period to period. Thus, the whole experiment, in a sense, moves around.

3.4 Updating and resetting. We need to use the experimental results to update the model and then use the model to pick the best national promotion rate. As is turns out we can combine these two tasks into one step rather simply in the present case.

The experimental results by themselves imply an apparent optimum promotion rate. It can be shown that the overall optimal promotion rate for the coming time period is a weighted sum of the apparent optimum rate from the experiment and the promotion rate used in the period just completed.

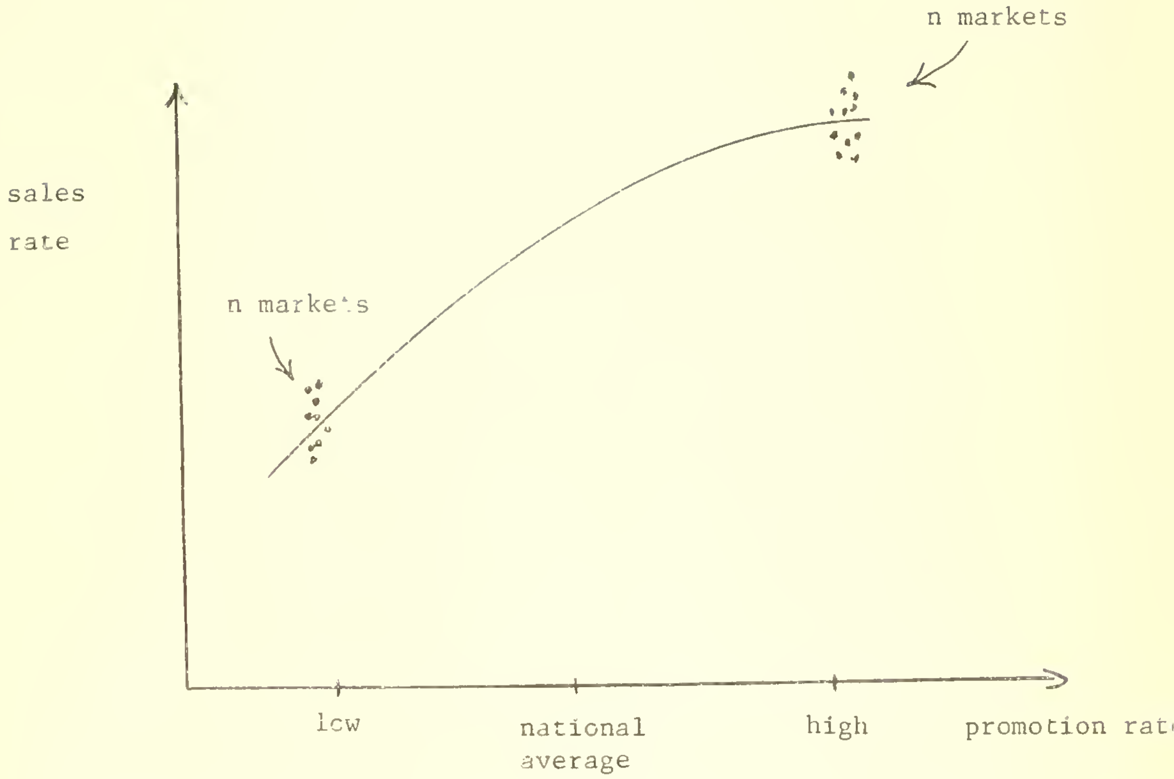

Figure 4. Field experiment to estimate effect of promotion on sales

This type of weighting of old and new data is known as exponential smoothing. It appears in many fields and finding it here is an interesting and perhaps even comforting surprise. The weights themselves can be determined by mathematical analysis (using Baysian statistics.) The weights have the following intuitively satisfying property: If the experiment is precise, most of the weight will automatically go to the current experimental data and little to the previous promotion rate. On the other hand, if the precision of the experiment is low, most of the weight goes to the previous promotion rate, which, in turn, represents a summary of considerable past experience.

Although an appropriate choice of weights is necessary for reasonable operation, the general form of the rule provides an adaptive control system that might be expected to work fairly well for a variety of underlying sales response mechanisms. One might expect that, if the weights were chosen with one mechanism in mind, they would work well with other mechanisms not too different. Our later results indicate just this and are in conformity with the experience with exponential smoothing in other contexts, where it has been found to be quite robust.

3.5 Designing the experiment. The next job is to specify the experimental design details. These are the sample size and the spread between high and low promotion rates. They will be picked to minimize the sum of two losses; the loss incurred because we do not know the sales response perfectly and the loss incurred by trying to learn sales response better. The losses will be calculated relative to the profits that could be obtained under perfect information.

It is clear enough that there will be, on the average, a loss when sales response is not perfectly known, since perfection is our reference point. The reason for the loss from performing the experiment is less obvious. It arises because the experimental markets are deliberately being held away from the rate thought to be best. Therefore, on the average, they incur an extra loss relative to that of the national rate. (In addition, there may be an out-of-pocket cost for running the experiments. We have not included this but could without difficulty.)

For the present model a fairly simple mathematical expression can be found for the sum of the losses and a straight-forward calculation will determine the optimal sample size and spread. Once the design of the optimal experiment is known, the weights for the updating rule can be calculated and the system is ready to go.

3.6 Numerical example. To try out the system and see how it might behave in practice, we have constructed a numerical example. In doing this we have tried to pick realistic values for the constants involved. We have not, however, tried to represent any specific product.

The constants have been chosen to give sales of about 25 million dollars/year. Promotion rate is about 1.5 million dollars/year or roughly 6% of sales. The performance of the system will be sensitive to the variability of sales from market to market. For a numerical statement of this, we have set the coefficient of variation of sales to 7%, a value that has been achieved in practical field experiments. We permit an experimental deviation in promotion rate up to ± 30% of the long term average. Much larger deviations have actually been used by companies in single experiments. This value seems workable for continuous use.

The values of the constants in the sales response function were chosen after considering a variety of actual experiments with which we have had contact. The values lie within the rather wide range that would be consistent with these experiments. The random disturbance to sales response was picked so that from time period to time period there would be about a 3% coefficient of variation in sales in the absence of other changes.

When the appropriate numbers are substituted into the formulas the optimal experiment comes out to involve 21 market areas at a high rate that is 30% above the national average and 21 at a low rate 30% below. These numbers then determine the weighting factors for new and old information. In the example, the weights turn out to be roughly equal.

To see how the system behaves we have simulated the underlying market over time and operated according to the adaptive control system. The results are shown in Figure 5. The top curve shows the random disturbances representing the environmental changes to sales response that the company was unable to anticipate. The company's experiments try to measure these changes. The second curve shows how the company sets promotion rate as a result of the experiments. The third curve gives the resulting sales. The bottom curve shows the loss relative to operation with perfect information. This is expressed, for convenience, as a % of the long run average promotion rate.

The response of the adaptive system is seen to be quite good. Note that the reaction of the promotion rate lags the disturbance by one time period since the information gained during one period is not available until the next. The losses are generally small, although occasional

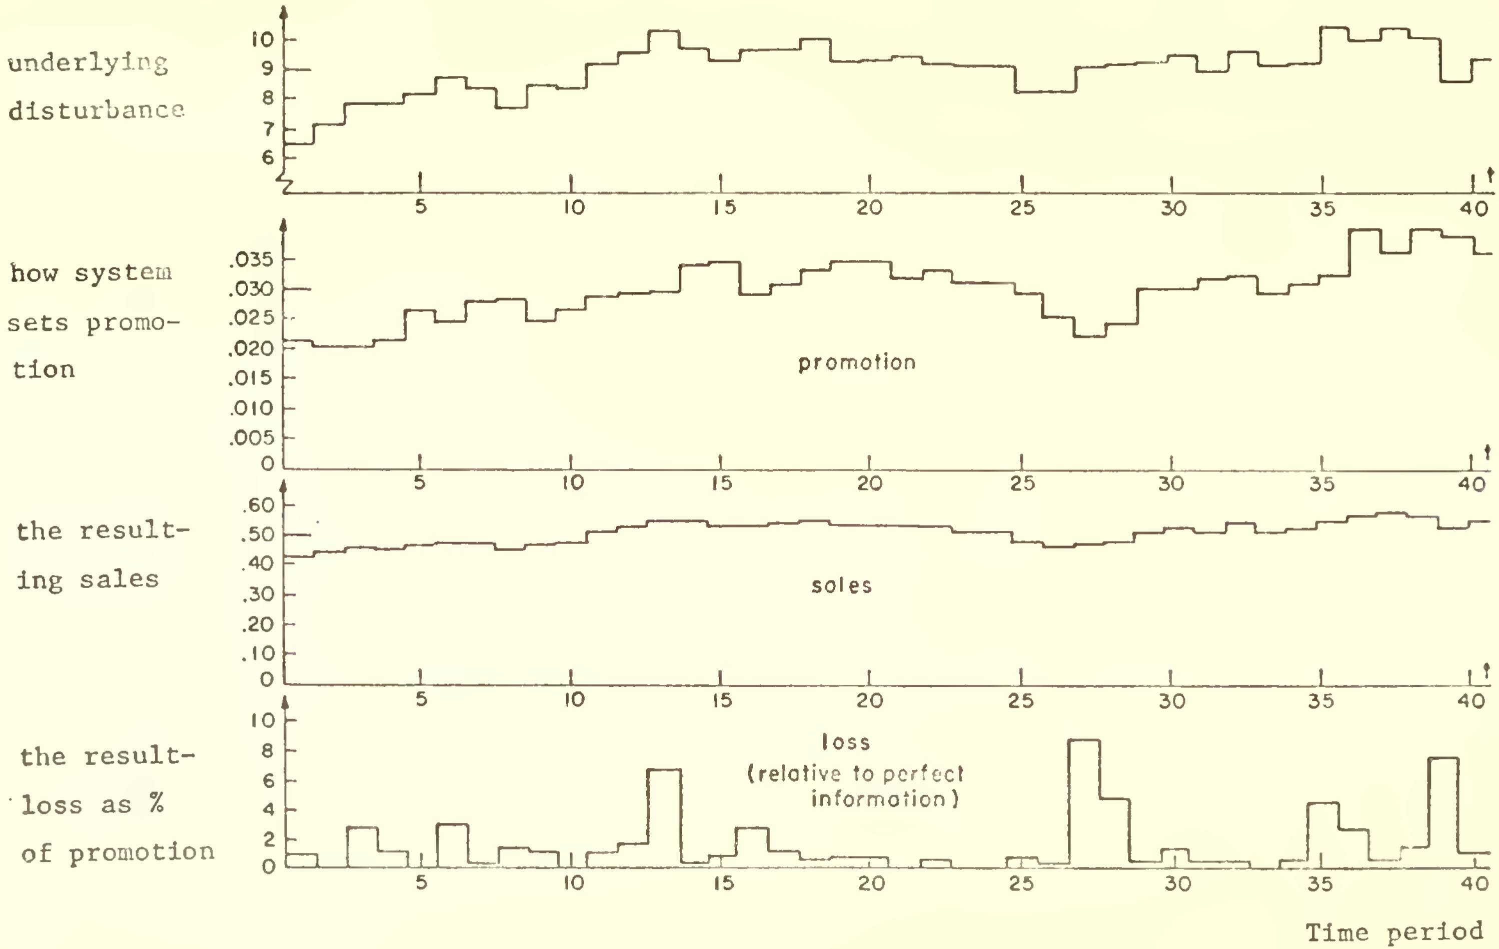

Figure 5. Simulation of Adaptive System for Setting Promotion Rate.

peaks occur when response is disturbed substantially and the promotion rate has not yet caught up. The long run average loss in this example is 1.61% of the average promotion rate or 24,200 dollars/year on a promotion budget of 1,500,000 dollars/year. This seems rather small considering that our standard is perfect information. Of the 1.61% loss, 1.23% comes from poor knowledge of sales response and .38% from the deliberate non-optimal operation involved in experimentation.

3.7 Comparison with other policies With the help of the numerical example we can make some comparison of the adaptive system with other types of operation. One possibility is to set promotion rate to a constant value. The question is: what value? Having perfect knowledge we can determine the best constant value. However, it is not clear how the company could do this (unless the president is clairvoyant.) Consequently we shall try not only the best constant rate but also values 50% higher and lower.

The results are shown in Table 1. The best constant rate is worse than the adaptive system by an amount equal to 2% of the promotional budget. However the loss rate in either case is really quite small. On the other hand, if the constant rate differs substantially from the "clairvoyant" value, the loss can be substantial; it equals 27% of the promotional budget in the case of a 50% deviation.

Suppose that circumstances prevent an experiment of the proposed size. Instead of 21 markets in each group suppose only 10 are possible. The experimental precision decreases and the weighting factors for new and old information must be adjusted, but the system is still adaptive. As shown in Table 1, the deterioration in performance is small.

Table 1

ADAPTIVE SYSTEM COMPARED TO OTHER POLICIES

Policy	Expected loss (relative to perfect information) as % of promotional budget
Adaptive System	1.6%
Best constant promotion rate	3.7%
Constant promotion rate 50% high (or low)	28.7%
Adaptive but experiment half size	1.7%

3.8 Sensitivity analysis. The adaptive system performed well in the example but the system is based on various assumptions about the underlying operation of the market. Many of the assumptions can be questioned. How will the system perform if some of them are incorrect? To investigate this, we use the experiment and operating rules as derived but change the underlying marketing models in various ways. This is the same as supposing that the company's models are, in certain ways, erroneous.

We consider four possible changes in the underlying world: (1) Perhaps sales response is really constant or essentially so. (2) Perhaps it is changing but there is no correlation from time period to time period. (3) Perhaps it has less curvature (diminishing returns) than we have hypothesized. (4) Perhaps it has greater curvature.

In the first two situations the learning aspects of the adaptive system would appear to be worthless: there are no trends to perceive. In the second two cases we test the question of how serious it is to make a poor estimate of the curvature (i.e. the coefficient of the quadratic term in the sales response function) in setting up the system.

The results are shown in Table 2. They are more or less all the same. For these cases, even when the adaptive system is based on wrong assumptions its operating results are quite good. Sometimes it is slightly better, sometimes slightly worse than the best constant promotion rate. In contrast, however, a constant rate which deviates considerably from the best gives rather substantial losses.

The explanation for the good performance, at least in this case, is that it is better to keep trying to learn the sales response, even if it is not changing and even if it has a different shape from what is expected, than it is to make a bad guess about the promotion rate directly.

Table 2

SENSITIVITY ANALYSIS: PERFORMANCE OF SELECTED POLICIES WITH VARIOUS UNDERLYING MODELS

Policy	Expected loss rate (relative to perfect information) as % of promotional budget				
	Original Model	Constant Sales Response	No Correlation	Less Curvature	More Curvature
Adaptive System	1.6	0.7	1.7	7.4	1.8
Best Constant Rate	3.7	0	0.7	14.6	1.8
Constant Rate: 50% high (or low)	28.7	25.0	25.7	20.6	51.8

4. DISCUSSION

In viewing the model of the previous section it seems perhaps surprising that such simple operating rules can cope successfully with various different environments. Other examples, however, are not hard to find. The home thermostat controls temperature under widely varying conditions of heat loss and does it without solving complicated heat flow equations on a digital computer. In our case there are several reasons why the adaptive system operates well. In the first place, the exponential smoothing procedure enhances the value of the measurements by combining past and present data. In the second place, sales response is presumed (reasonably enough) to vary smoothly with spending. As a result, the profit maximum is also smooth and small deviations from best operation cost very little. Substantial losses are produced only by large deviations and these tend to be avoided by the adaptive system. Thirdly, various empirical studies suggest that, once some nominal promotion rate is achieved, spending rate per se is not a big factor in sales. This is not to say that increases or decreases may not be profitable but rather to say that they are not likely to cause the jumps in sales that are caused, for example, by product changes or, sometimes by changes in promotional treatment.

How can the adaptive system ideas be put to use? They can be introduced with varying degrees of formality. Perhaps the key general principles are: (1) some kind of objective measurement is necessary and (2) such measurements should be continuously made so as to accumulate information and permit adaptation to changing conditions. As mentioned earlier, there are companies today who monitor the spending level in their

major promotional activity by keeping an experiment in the field at all times. The present model offers guidance on the best way of doing this. A particularly key place where guidance is needed is on the basic design questions of how many markets to use and what size of experimental deviation to pick. The model approaches this by considering the costs and benefits involved.

5. FUTURE DEVELOPMENTS

The basic ideas discussed here are being developed in several directions, both with respect to new theory and new applications. On the theoretical side, the model as reported has only one control variable (promotional spending) but extensions are being made to take account of many control variables. Another limitation of the current model is that a variable like promotional spending changes continuously in small steps whereas many marketing decisions are discrete (this package or that, this campaign theme or another, etc.) Developments are under way to accomodate discrete choices.

Considerably more complex models are under development, encompassing, for example, consumer awareness, brand switching, and retail distribution. Furthermore it would be helpful to have such models available for on-line manipulation and updating at the remote terminal of a time-shared computer. Urban [14] discusses the conceptual framework of such a model.

Applications are growing, particularly in the new product and advertising areas. Furthermore, since marketing-like problems are not confined to business, applications are currently being considered in such diverse areas as political campaigning and fund-raising for non-profit institutions.

CT 20 '70

REFERENCES

1. P. J. Robinson and D. J. Luck, Promotional Decision Making: Practice and Theory, McGraw Hill, New York 1964

2. P. J. Harrison, "Exponential Smoothing and Short-Term Sales Forecasting," Management Science, 13, 821-842 (July 1967)

3. W. Alderson and P. E. Green, Planning and Problem Solving in Marketing, Irwin, Homewood Illinois, 1964 (Chapter 8)

4. G. L. Urban, "SPRINTER: A Tool for New Product Decision-Making," Industrial Management Review, 8, 43-55 (Spring 1967)

5. A. Charnes, W. W. Cooper, J. K. DeVoe, and D. B. Learner, "DEMON: Decision Mapping Via Optimum GO-NO Networks - A Model for Marketing New Products," Management Science, 12, 865-887 (July 1966)

6. D. B. Montgomery and G. L. Urban, Management Science Models in Marketing, Prentice Hall, Englewood Cliffs, New Jersey (forthcoming)

7. J. D. C. Little, "A Model of Adaptive Control of Promotional Spending," Operations Research, 14, 1075-1097 (November 1966)

8. D. Sworder, Optimal Adaptive Control Systems, Academic Press, New York, 1966

9. M. Aoki, Optimization of Stochastic Systems, Academic Press, New York, 1967

10. W. S. Hoofnagle, "Experimental Designs in Measuring the Effectiveness of Promotion," Journal of Marketing Research, 2, 154-162 (May 1965)

11. P. T. FitzRoy, "An Adaptive Model for Promotional Decision Making," MSI Special Report, Philadelphia 1967

12. E. G. Hurst, Jr., "A Class of Models for Adaptive Experimentation and Control," Ph D. Thesis, M.I.T., Cambridge, Mass. June 1967

13. A. B. deAndrade, "An Adaptive Control Model for Promotional Spending," S.M. Thesis, M.I.T., Cambridge, Mass. February 1968

14. G. L. Urban, "Market Response Models for the Analysis of New Products," Working Paper 336-68, Sloan School of Management, M.I.T., Cambridge, Mass. June 1968

~~OCT 20 '70~~

Date Due

~~FEB 25 '76~~

NOV 05 '78

~~79~~

SEP 11 '79

JUL 10 1980

NOV 25 [illegible]

NOV 29 '84

JA[illegible]

FEB 19 1990

Lib-26-67

MIT LIBRARIES DUPL
3 9080 003 874 358
342-68

MIT LIBRARIES DUPL
3 9080 003 905 384
343-68

MIT LIBRARIES DUPL
3 9080 003 874 390
344-67

MIT LIBRARIES DUPL
3 9080 003 905 343
345-68

MIT LIBRARIES DUPL
3 9080 003 874 317
346-68

MIT LIBRARIES DUPL
3 9080 003 873 988
347-68

MIT LIBRARIES DUPL
3 9080 003 904 957
348-68

MIT LIBRARIES DUPL
3 9080 003 905 327
349-68