

Copy 1

Dewey

A Model of Adaptive Control of Promotional Spending

122-65

John D. C. Little

Working Paper

May 7, 1965

HD28

.M414

no. 122-65

RECEIVED
NOV 9 1965

Abstract

A simple model of adaptive control of promotional spending is analyzed. In the model, company sales (and therefore profits) are functions of promotional spending rate. Sales response to the promotion changes with time as a result of changes in a parameter of the sales response function. Information about sales response is collected in each time period by performing an experiment. On the basis of present and past information, the promotion rate is set to maximize expected profit in the next period. The experiment is chosen to minimize the combined costs of imperfect information and experimentation.

A numerical example is studied analytically and by simulation the adaptive system appears to work well. In a sensitivity analysis, the system based on one underlying model of the market is found to perform well when the underlying model is changed considerably.

A Model of Adaptive Control of Promotional Spending

John D. C. Little^{*}

1. Introduction

A company must assemble marketing information, use it to modify its conception of the market, use the revised conception to make marketing decisions, and then arrange for the collection of new information. In short, a company needs a control system for its marketing variables.

Obviously every company has some procedure for determining its marketing actions, but usually the relationships between data inputs and decision outputs are not at all formally specified. Our interest is in studying possible inputs and possible relationships to determine their effect on overall company performance. Presumably, by careful systems design, companies can achieve better marketing performance than they do now.

Formal systems design in the sense we mean will require considerable development. Marketing variables are many and so are the possible sources and forms of information. However, we can at least start the job by investigating a simple marketing system that involves some of the important ideas.

The system consists of a model of the marketing process to be controlled, a means of using the model to set values for the marketing

* Sloan School of Management, MIT, Cambridge, Massachusetts, 02139. Work reported was supported in part by Project MAC, an MIT research program sponsored by the Advanced Research Project Agency, Department of Defense, under the Office of Naval Research Contract No. Nonr-4102(01). Reproduction in whole or in part is permitted for any purpose of the United States Government.

variables, and a measurement device for keeping the model up to date. The model of the process is briefly as follows: company sales (and therefore profits) are functions of a single variable, a rate of spending money on promotion. The sales response to promotion changes with time. The change is the result of a single changing parameter in the sales response function.

The control cycle operates as shown in Figure 1: Based on the sales model, a calculation is made that sets the promotion rate to maximize expected profit in the next time period. An experiment is then designed to monitor the effectiveness of the promotion. The results are then implemented, the market responds, and some sales rate is produced. The sales data thereby generated represents new information, which is then combined with old information to update the sales model. The cycle is then repeated.

In designing the control system, we must specify how to combine new and old information and we must determine what size of experiment to perform. We plan to combine new and old information so as to maximize the expected profit from the decisions that use the information. The size of the experiment will be chosen to minimize the sum of the losses arising from imperfect information about sales response and from the cost of performing the experiment.

Thus, the problem is set up according to the principles of statistical decision theory. The problem has a complication not found in elementary applications, however, because sales response is changing with time, and, as a result, the value of a piece of information deteriorates with its age.

For a measuring instrument, we shall work with direct sales experiments. That is, we shall take different groups of customers,




Figure 1. Cycle of adaptive control of promotion.

say, people in different geographical areas, give them different experimental treatments, and observe the effect on sales. Other, less direct devices for monitoring sales response are of course possible.

The overall system is certainly an idealization of real world operations, but it is perhaps not too far from practicality. At least one company we know goes around the feedback loop of Figure 1 in an informal way now.

2. Sales and Profit Models

The profit model expresses company profit rate in terms of its sales rate, promotion rate, fixed cost rate, and incremental profit on sales, the latter taken to be constant. To make the quantities as concrete as possible we shall give them specific units. In particular, money flows will be expressed in dollars per household per year (dol./hh.yr.). Let

s = sales rate. (dol./hh.yr.)

x = promotion rate. "

p = profit rate. "

c = fixed cost rate. "

m = gross margin, the incremental profit as a fraction of sales. (dimensionless)

The model of company profit is

$$(2.1) \quad p = ms - x - c .$$

Notice that promotion enters here as a fixed cost. Thus we are not planning to consider variable cost promotional activities (e.g. price-off deals), although there is no conceptual difficulty in so doing.

The sales model is set up as a simple quadratic in promotion rate. We suppose that, for a given fixed time period, sales response has the general shape shown in Figure 2 and that the curve can be approximated, at least near the current operating point, by a quadratic function of x :

$$(2.2) \quad s = \alpha + \beta x - \gamma x^2 .$$

The parameters α , β , and γ are constants for the fixed time period (they may be thought of as giving the average sales rate over the period) but some of them may vary from period to period.

The value of x , say x^* , that maximizes profit is easily found to be

$$(2.3) \quad x^* = (m\beta - 1)/2m\gamma .$$

If the company uses x instead of x^* , the loss rate, ℓ , relative to maximum profit is

$$\ell(x) = p(x^*) - p(x) ,$$

which, using (2.1), (2.2), and (2.3), becomes

$$(2.4) \quad \ell(x) = m\gamma(x - x^*)^2 .$$

3. Changes in Sales Response

If the sales response parameters α , β , and γ were known, we would set promotion rate to x^* and obtain the loss $\ell(x^*) = 0$. However, the parameters are presumably fairly difficult to measure and we ordinarily expect to come up with some non-optimal x and therefore to incur a relative loss.

If the parameters were constant over time, we would put a big effort into measuring them right away, because the extra profit from increased accuracy would extend far into the future. However, it is difficult to believe that in practice the parameters stay constant.

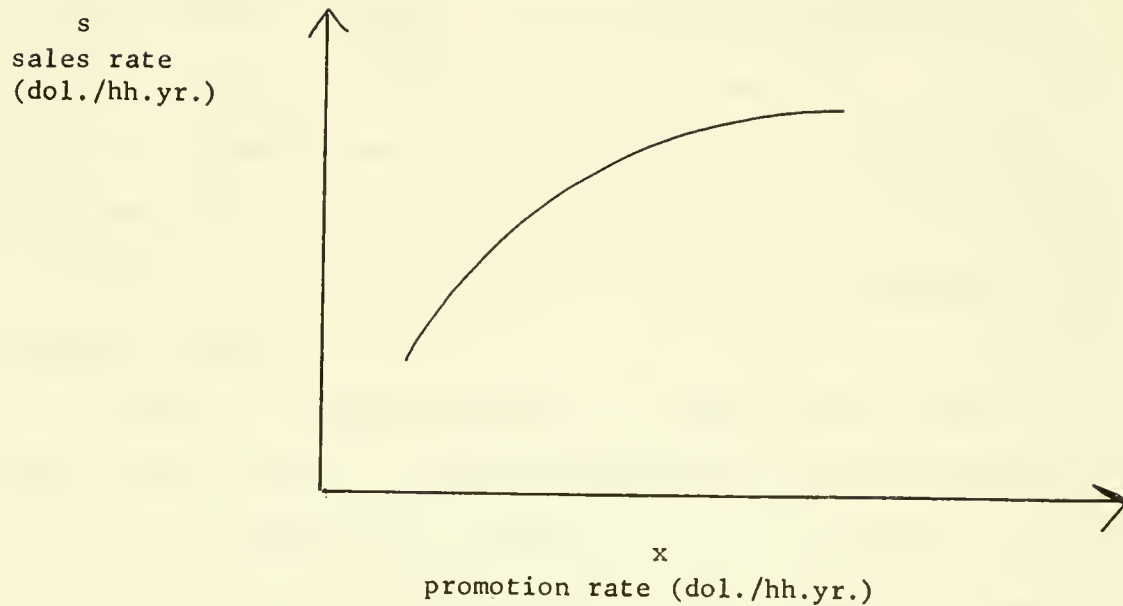


Figure 2. Shape of sales response to promotion in a fixed time period.

For example, competitive activity, product changes, changes in the quality of the promotion, and shifts in economic conditions lead us to expect shifts in response. Consequently, an expensive effort to learn the parameters immediately cannot be justified. On the other hand, the parameters may change fairly slowly with time, in which case some effort is worthwhile. In each time period new information is collected, combined with the old and used to set operations in the immediate future.

To build a fairly simple model of changing sales response, we shall suppose that α and β change with time in a specified way involving some randomness, but that γ does not. As a matter of notation, we shall use the tilde (\sim) when we wish to emphasize that some quantity is being viewed as a random variable. Furthermore, α , β , s and x usually depend on t , and when this needs emphasis we shall write $\alpha(t)$, $\beta(t)$, etc.

At a fixed time period, t , we assume that national sales rate for the product is

$$(3.1) \quad \tilde{s} = \tilde{\alpha} + \tilde{\beta}x - \gamma x^2 \quad (\text{dol./hh.yr.})$$

and that

$\tilde{\alpha}(t)$ has relatively high variance from time period to time period

$\tilde{\beta}(t)$ is dependent on $\tilde{\beta}(t-1)$,

$\tilde{\alpha}$ and $\tilde{\beta}$ are independent,

γ is a known constant.

As implied by (2.3) and as will be seen, the current information about α does not directly affect the optimal x . However, information about α may make it possible to learn more about β in an experiment. The assumption of high variance for α , (high conditional

variance given the previous α 's if successive α 's are considered dependent) may be fairly realistic and, in any case, simplifies the statistical analysis by removing α as a contributor to information about β .

The assumption that γ is known and constant seems quite unrealistic, but we shall argue later that this may not be too serious to the successful operation of the system.

The response parameter β will be considered to be generated by a random walk. One possibility is $\tilde{\beta}(t) = \tilde{\beta}(t-1) + \tilde{\epsilon}_{\beta}(t)$, where

$$\tilde{\epsilon}_{\beta}(t) = \text{a random variable with mean} = 0 \text{ and variance} = \sigma_{\beta}^2.$$

We shall take $\tilde{\epsilon}_{\beta}(t)$ to be normally distributed and independent of previous values of β and ϵ_{β} . The difficulty with the above random walk is that β is likely to wander unrealistically far from its starting value. Therefore, we shall hypothesize a long run average value and a tendency for β to return to that value. Specifically, let

$$\beta^0 = \text{the long run average value of } \tilde{\beta}(t).$$

$$k = \text{a constant, } 0 \leq k \leq 1.$$

We take as our model of changing β :

$$(3.2) \quad \tilde{\beta}(t) = k \tilde{\beta}(t-1) + (1-k)\beta^0 + \tilde{\epsilon}_{\beta}(t)$$

As $k \rightarrow 1$, $\beta(t)$ is increasingly dependent on $\beta(t-1)$ and also wanders more and more freely from β^0 . As $k \rightarrow 0$, β becomes independent from time period to time period. Figure 3 shows a possible sequence of $\beta(t)$ versus t . To make $\beta(t)$ more operationally meaningful, we observe that if the promotion rate is held constant, the sales rate consists of a constant plus a term directly proportional to $\beta(t)$.




Figure 3. Sketch of a possible variation of the sales response parameter, β , with time period, t .

4. An Experiment to Measure Current Sales Response

In the previous two sections we have described the environment in which the company is operating. Next we describe the control system. Our starting place in the cycle of Figure 1 will be the point of designing an experiment to measure sales response.

Measuring sales response essentially means measuring β , since α does not enter decisions and γ is taken to be a known constant. Information about β will be collected by operating different groups of market areas at different promotion rates.

Although the national average sales rate is given by (3.1) we suppose that individual markets differ in sales rate because of local random variations. For some fixed t suppose that the national parameters take on the specific values $\tilde{\alpha} = \alpha$ and $\tilde{\beta} = \beta$. We assume sales in a market are then given by

$$(4.1) \quad \tilde{s} = s(x) + \tilde{\epsilon}$$

where \tilde{s} = sales rate in the market. (dol./hh.yr.)

$s(x) = \alpha + \beta x - \gamma x^2$ = national sales rate. (dol./hh.yr.)

$\tilde{\epsilon}$ = a random variable for the market. (dol./hh.yr.)

We assume that $\tilde{\epsilon}$ is normally distributed, is independent from market to market, and has mean = 0 and variance = σ^2 .

The experiment is sketched in Figure 4. Suppose that at t we have picked a promotion rate, $x_0(t)$. This will be used everywhere except that in n markets we shall use a deliberately low value, x_1 , and in another n a deliberately high value, x_2 . We take

$$(4.2) \quad x_1 = x_0(t) - \Delta/2$$

$$x_2 = x_0(t) + \Delta/2$$

where Δ is a design constant yet to be selected.

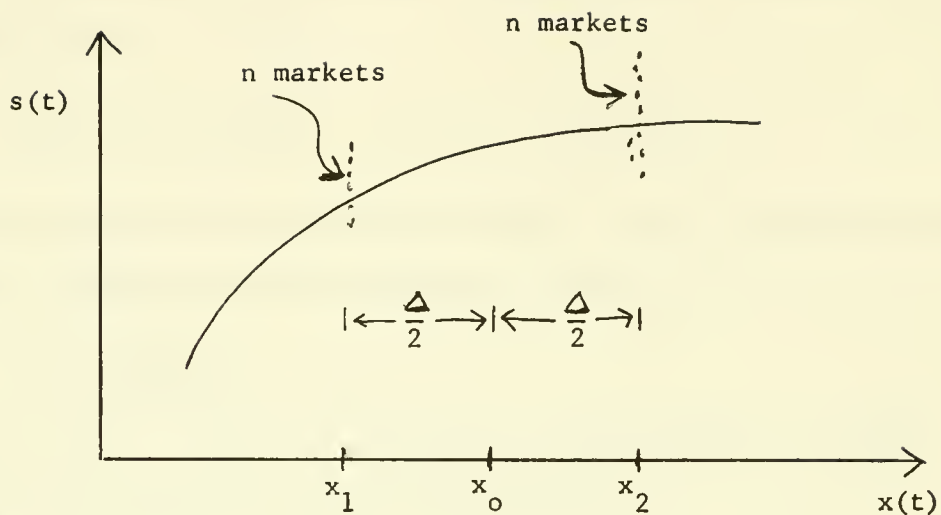


Figure 4. The sales experiment. A group of n markets is given a promotion rate $\frac{\Delta}{2}$ greater than the national average, x_0 , and another group of n is given a rate $\frac{\Delta}{2}$ less.

Let \bar{s}_1 and \bar{s}_2 be the observed mean sales rates in the groups of markets x_1 and x_2 respectively. An estimate of $\beta(t)$ can be computed from the experimental data:

$$(4.3) \quad \hat{\beta}(t) = (1/\Delta)(\bar{s}_2 + \delta x_2^2 - \bar{s}_1 - \delta x_1^2) .$$

This will be called the "experimental mean". It is a random variable because \bar{s}_1 and \bar{s}_2 are random variables. Let *

$$v = V[\hat{\beta}(t)] .$$

From (4.1), (4.2), and (4.3) we find that, given $\tilde{\beta}(t) = \beta(t)$, $\hat{\beta}(t)$ is normally distributed with mean and variance:

$$(4.4) \quad E[\hat{\beta}(t)] = \beta(t) \\ v = 2\sigma^2/n\Delta^2$$

Notice that v does not depend on t .

The experimental result (4.3) does not represent all our information about β . Even before doing the experiment we had the information developed in previous experiments. The information will be summarized in a prior distribution for $\tilde{\beta}(t)$. This distribution will be taken to be normal with

$$E[\tilde{\beta}(t)] = \text{mean of prior distribution of } \tilde{\beta}(t), \\ v' = \text{variance of prior distribution of } \tilde{\beta}(t).$$

At the beginning of period t , when the promotional rate, $x(t)$, is to be set, we have only the prior distribution. At the end of t , the experimental results are at hand and we can construct a posterior distribution. The additional information about $\beta(t)$, however, is of no use in t even though it will be helpful in $t+1$.

* We use $V[\]$ and $E[\]$ to denote the variance and expectation operators, respectively.

5. Decision Rules for Updating the Model and Setting the Promotion Rate

Returning to the control cycle, suppose that the experiment for period t has been implemented, the market has responded, and we have in hand the sales results. We are now ready to update the model and go on to pick a promotion rate for period $t+1$.

Decision rules for updating the model and setting promotion can be determined by formal decision theory arguments, but we prefer to defer these and begin intuitively. One reason is that the optimality of formally derived rules is entirely dependent on the specific assumptions of the model, some of which are rather restrictive. The form of our decision rules, however, appears reasonable (although not necessarily optimal) for a wide class of situations.

The decision rules are, first, to update the mean of our prior distribution by an exponential smoothing process and, second, to set promotion rate by a formula analogous to (2.3). Specifically, we choose a number, a , such that $0 \leq a \leq 1$ and update the mean by

$$(5.1) \quad E^0 \left[\tilde{\beta}(t+1) \right] = a E^0 \left[\tilde{\beta}(t) \right] + (1-a) \hat{\beta}(t) .$$

Then we choose as our promotion rate for $t+1$:

$$(5.2) \quad x_0(t+1) = \left(m E^0 \left[\tilde{\beta}(t+1) \right] - 1 \right) / 2m\delta .$$

Suppose the process starts at $t=1$. It is necessary to pick a starting value $E^0 \left[\tilde{\beta}(0) \right]$, but, thereafter, promotion rate is set mechanically by the rule. Since we are dealing with an exponential smoothing process, the effect of the starting value on later operations decays exponentially with t .

Notice that the rule as stated makes no assumptions about the underlying mechanisms generating $\alpha(t)$ and $\beta(t)$ and, if we think

of δ as an arbitrary positive constant, the rule is not tied down to any specific sales response process. The rule can be applied to any situation in which the experiment of Section 4 is performed, provided that somebody is willing to pick \underline{a} and $\underline{\delta}$ plus $E' \left[\tilde{\beta}(0) \right]$ and the experimental design parameters Δ and n .

The behavior of the decision rules will be clearer if we express the resulting promotion rate somewhat differently. Let us define a quantity, $\hat{x}_0(t)$, to be called the "experimental \hat{x}_0 " by

$$(5.3) \quad \hat{x}_0(t) = \left[m \hat{\beta}(t) - 1 \right] / 2m\delta.$$

This is the promotion rate that would be best if β actually equalled the experimentally determined $\hat{\beta}$. Using (5.3), (5.2) can be rewritten:

$$(5.4) \quad x_0(t+1) = a x_0(t) + (1-a) \hat{x}_0(t)$$

Thus the decision rules amount to using a weighted combination of last period's promotion rate and the experimental \hat{x}_0 . If we have a tight, accurate experiment, we should use a small \underline{a} and so rely mostly on the current experiment. If the accuracy of the experiment is low, a large \underline{a} is appropriate. Then this year's promotion rate depends mostly on last year's, which, in turn, represents a summary of considerable past experience.

Still another way of writing the decision rule brings out more sharply the role of the sales data. Using (4.3) and (5.3) we get

$$(5.5) \quad x_0(t+1) = x_0(t) + \left[(1-a)/2m\delta\Delta \right] \left[m(\bar{s}_2 - \bar{s}_1) - \Delta \right].$$

The quantity $\left[m(\bar{s}_2 - \bar{s}_1) - \Delta \right]$ is simply the experimentally estimated net profit for increasing the promotion rate from x_1 to x_2 , i.e. by

an amount Δ . If the net profit is positive, $x_0(t+1)$ is made larger than $x_0(t)$; if it is negative, $x_0(t+1)$ is made smaller. The amount of the adjustment is controlled by the constant $(1-a)/2m\gamma\Delta$. If the constant is large, the promotion rate will be sensitive to the most recent experiment; if the constant is small, insensitive.

Although an appropriate choice of constants is necessary for reasonable operation, the general form of our 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 constants were chosen with one mechanism in mind, they might work satisfactorily with other mechanisms not too different. Whether or not this is so in a specific case can be explored by simulation and sensitivity analysis.

For the sales response model being assumed we now wish to motivate the decision rules more carefully and go on to pick optimal values for \underline{a} , n , and Δ .

6. Choosing the Smoothing Constant

The smoothing constant, \underline{a} , was used in (5.1) to combine new and old information about β . We shall now show that \underline{a} can be expressed in terms of the experimental design constants and the constants of the $\beta(t)$ process.

Consider first the problem of finding the posterior distribution of $\beta(t)$ given its prior distribution and the experimental results at t . Let

$$\begin{aligned}
 E' \left[\tilde{\beta}(t) \right] &= \text{mean of prior} \\
 \hat{\beta}(t) &= \text{experimental mean} \\
 E'' \left[\tilde{\beta}(t) \right] &= \text{mean of posterior}
 \end{aligned}$$

v' = variance of prior

v = experimental variance

v'' = variance of posterior

Since the prior of $\beta(t)$ is normal and the distribution of $\hat{\beta}(t)$ given $\beta(t)$ is normal, the posterior of $\beta(t)$ is normal and has mean and variance (see [1] p. 294-5):

$$(6.1) \quad E'' [\tilde{\beta}(t)] = [v/(v+v')] E' [\tilde{\beta}(t)] + [v'/(v+v')] \hat{\beta}(t)$$

$$(6.2) \quad v'' = v v' / (v+v')$$

The process generating $\beta(t+1)$ has been specified in (3.2):

$$\tilde{\beta}(t+1) = k \tilde{\beta}(t) + (1-k)\beta^0 + \tilde{\epsilon}_{\beta}(t+1) .$$

As of the beginning of $t+1$, we know the posterior distribution for $\beta(t)$. We also know the distribution of $\tilde{\epsilon}_{\beta}(t+1)$. Since the two random variables are independent, the prior distribution of $\beta(t+1)$ has mean and variance:

$$(6.3) \quad E' [\tilde{\beta}(t+1)] = k E'' [\tilde{\beta}(t)] + (1-k)\beta^0$$

$$(6.4) \quad v' [\tilde{\beta}(t+1)] = k^2 v' [\tilde{\beta}(t)] + \sigma_{\beta}^2 .$$

Furthermore, the prior of β at $t+1$ is normal and so the normality of β is preserved as time passes.

Substituting (6.1) into (6.3), we obtain

$$(6.5) \quad E' [\tilde{\beta}(t+1)] = k \left\{ [v/(v+v')] E' [\tilde{\beta}(t)] + [v'/(v+v')] \hat{\beta}(t) \right\} + (1-k) \beta^0 .$$

If k is near one, a good approximation may be obtained by setting $k=1$ in the above expression. This has the advantage of eliminating k and β^0 as parameters in the decision rules and so we shall use it.

However, there is no fundamental difficulty in carrying along k and β^0 , except that we must specify their values.

Under the $k=1$ approximation, (6.5) is in the same form as (5.1) and so we see that

$$(6.6) \quad a = v/(v+v') .$$

Thus, for this case, we have justified the decision rule (5.1) and found a way to compute \underline{a} . It remains to express v and v' in terms of known parameters.

We already know v in terms of the experimental design parameters by (4.4). With respect to v' , we first observe that v' and v'' will not change with t once steady state operation is achieved. This is because σ_β^2 and v do not change with t . In steady state, (6.4) and (6.2) become

$$\begin{aligned} v' &= k^2 v'' + \sigma_\beta^2 \\ v'' &= v v''/(v+v') . \end{aligned}$$

Using the $k=1$ approximation and solving for v' , we obtain

$$(6.7) \quad v' = \frac{1}{2} \sigma_\beta^2 \left\{ 1 + \left[1 + (4v/\sigma_\beta^2) \right]^{\frac{1}{2}} \right\} ,$$

where $v = 2\sigma^2/n\Delta^2$. Substitution of v and v' into (6.6) gives the smoothing constant, \underline{a} , in terms of σ , n , Δ , and σ_2 . This completes the job of finding \underline{a} .

7. Setting the Promotion Rate

Next we wish to justify the decision rule (5.2) for setting promotion rate. Profit rate is a random variable because sales rate is:

$$(7.1) \quad \tilde{p}(t) = m \tilde{s}(t) - x(t) - c$$

$$(7.2) \quad \tilde{s}(t) = \tilde{\alpha}(t) + \tilde{\beta}(t) x(t) - \gamma x^2(t)$$

We cannot maximize true profit but choose instead to maximize expected profit. At the start of t , the company's view of expected profit in t is (simplifying the notation by suppressing t):

$$(7.3) \quad E' [\tilde{p}] = m \left\{ E' [\tilde{\alpha}] + E' [\tilde{\beta}] x - \gamma x^2 \right\} - x - c .$$

The company can maximize this by setting $x(t)$ to be

$$(7.4) \quad x_0(t) = (m E' [\tilde{\beta}(t)] - 1) / 2m\gamma .$$

This is (5.2).

At this point we have justified the decision rules of Section 5 for our specific model, at least to the extent of $k=1$ approximation. It remains to pick the parameters n and Δ of the experiment.

8. Designing the Experiment

The experimental design parameters will be picked to minimize the sum of two losses: the loss incurred because we do not know β exactly and the loss incurred trying to learn β better. The losses will be calculated relative to the profits obtainable under perfect information.

With perfect information we would choose the promotion rate:

$$(8.1) \quad \tilde{x}^*(t) = [m \tilde{\beta}(t) - 1] / 2m\gamma .$$

Instead we choose

$$(8.2) \quad x_0(t) = [m E' [\tilde{\beta}(t)] - 1] / 2m\gamma .$$

Notice that

$$(8.3) \quad E' [\tilde{x}^*(t)] = x_0(t) .$$

The loss rate compared to perfect information is seen from (2.4) to be

$$\tilde{Q}(t) = m\gamma [x_0(t) - \tilde{x}^*(t)]^2 .$$

Therefore,

$$E' [\tilde{Q}(t)] = m\gamma V' [\tilde{x}^*(t)] = (m/4\gamma) V' [\tilde{\beta}(t)]$$

or

$$(8.4) \quad E' \left[\tilde{\ell} \right] = (m/4\delta) v' ,$$

independent of t . This is the expected loss rate relative to perfect information for those markets where we follow the decision rule to use $x_0(t)$.

In the $2n$ experimental markets, the expected loss rate is higher because the promotion rate is deliberately set to be different from the best available value, $x_0(t)$. Suppressing t for the moment, the experimental promotion rates are

$$x_1 = x_0 - 1/2 \Delta$$

$$x_2 = x_0 + 1/2 \Delta .$$

Consider a market at x_1 . Let its loss rate relative to perfect information be $\tilde{\ell}_1$.

$$\tilde{\ell}_1 = m\delta \left[x_1 - \tilde{x}^* \right]^2 = m\delta \left[(x_0 - \tilde{x}^*)^2 - \Delta(x_0 - \tilde{x}^*) + (\Delta^2/4) \right] .$$

Therefore, using (8.3) and (8.4),

$$E' \left[\tilde{\ell}_1 \right] = E \left[\tilde{\ell} \right] + m\delta \Delta^2/4 .$$

Letting $\tilde{\ell}_{ex} = \tilde{\ell}_1 - \tilde{\ell}$ be the extra cost rate of the experimental deviation, we see that

$$(8.5) \quad E' \left[\tilde{\ell}_{ex} \right] = m\delta \Delta^2/4 . \text{ This same expression holds for a market at } x_2 .$$

The total expected loss rate can now be computed. The loss rates above apply to individual markets and have the dimensions dol./hh.yr.

Let

N = total number of markets in the country,

$2n$ = the number of experimental markets,

P = the average number of households in a market,

T = total expected loss rate (dol./yr.)

We shall assume, for simplicity, that all markets have the same size, but this is not essential. Then,

$$T = N P m v^i / 4\gamma + (1/2) P m \gamma n \Delta^2 \quad (\text{dol/yr})$$

In later discussion of numerical results, it will be convenient to express loss rates in percentage terms. A problem is that most quantities that are candidates for forming the denominators of such percentages are fluctuating with time. However, for several quantities we can construct reference values about which actual values fluctuate. Let

$$\beta^0 = \text{long run average of } \beta(t)$$

$$\alpha^0 = \text{long run average of } \alpha(t)$$

$$(8.6) \quad x^0 = (m \beta^0 - 1) / 2m\gamma$$

$$(8.7) \quad s^0 = \alpha^0 + \beta^0 x^0 - \gamma (x^0)^2$$

Here x^0 is the optimal promotion rate if $\beta = \beta^0$, and s^0 is the corresponding rate when $\alpha = \alpha^0$. These quantities make convenient reference points.

Now, let

L = expected loss rate relative to perfect information (and no experiment) as a fraction of the long run average promotion rate, x^0 .

$$(8.8) \quad L = \frac{\left[(N P m v^i / 4\gamma) + (1/2) P m \gamma n \Delta^2 \right] / N P x^0}{m v^i / 4\gamma x^0 + m \gamma n \Delta^2 / 2 N x^0}$$

The experimental design parameters n and Δ will be picked to minimize L . First, we observe from (8.8) and (4.4) that n and Δ always appear in L in the combination $n \Delta^2$. We shall therefore find $n \Delta^2$ to minimize L . Then we can more or less trade off n against Δ any way we wish as long as $n \Delta^2$ is kept to its minimizing value. Practically, there are limitations (e.g. x_1 must be non-negative and $2n$ not greater

than N , to cite two extremes) but the flexibility implied is interesting and valuable.

Rather than minimize L with respect to n_{Δ}^2 directly, we minimize with respect to the dimensionless quantity $z = 8\sigma^2/\sigma_{\beta}^2 n_{\Delta}^2$, from which n_{Δ}^2 can be immediately calculated, by means of

$$(8.9) \quad n_{\Delta}^2 = 8\sigma^2/\sigma_{\beta}^2 z .$$

Substituting this into (4.4) to get v , v into (6.7) to get v' , and then v' and n_{Δ}^2 into (8.8), we obtain L in terms of z :

$$(8.10) \quad L = (m\sigma_{\beta}^2/8 \gamma x^0) \left[1 + (1+z)^{\frac{1}{2}} \right] + 8m\gamma \sigma^2/N\sigma_{\beta}^2 x^0 z$$

Setting $dL/dz = 0$, we obtain an equation that optimal z must satisfy:

$$(8.11) \quad z/(1+z)^4 = 8\gamma \sigma/\sigma_{\beta}^2 \sqrt{N} .$$

The equation can be solved for z graphically, or simply by trial and error. Therefore, given the system constants, z can be determined from (8.11) and the optimal experimental design constant, n_{Δ}^2 , can be found from (8.9).

9. Numerical Example

The behavior of the system will next be illustrated by a numerical example . Given a set of values for the constants, we design the optimal experiment, determine decision rules for setting promotion rate, and simulate system operation. Average losses can be calculated from the simulation or directly from expected loss formulas. In this section we suppose that the underlying model on which we have based the system design is correct. We compare optimal operation with various other policies. In a later section we consider examples of what happens when the underlying model is not what we supposed.

Constants. In constructing the example we have tried to pick realistic values for the constants. We have not, however, tried to

represent any specific product. Sales and other figures will be given in absolute units, but the example is constructed from assumptions about percentages. Therefore, to the extent that the results are realistic, they are intended to be realistic under changes of scale that multiply sales and promotion by the same factor.

The constants have been chosen to give sales of about 25 million dollars/year ($s^0 = .50$ dol./hh.yr., based on 50 million households). Promotion rate is about 1.5 million dollars/year ($x^0 = .03$ dol./hh.yr.) or roughly 6% of sales. Gross profit margin has been taken as 1/3 of the selling price.

The performance of the system is sensitive to the accuracy of the experiment, which in turn depends on the variance of sales among markets. We have chosen $\sigma/s^0 = 7\%$, a value that has been achieved by some companies in field experiments.* We permit an experimental deviation in promotion rate of $\pm 25\%$ of the long term average, x^0 . Then $\Delta = .5x^0$. Much larger deviations have been used in practice for single experiments. The proposed deviation seems workable for continuous use.

The hardest question in picking constants is how to set reasonable values for the sales response parameters. There are a few sales experiments that have given an indication of diminishing returns for the case of advertising. As might be expected, the results show considerable variation. We have chosen β^0 and γ so that they lie within the rather wide range of values that are consistent with these experiments. The constant σ_β determines the period to period variance in sales response. The choice here has been quite subjective. However

* Field experiments are usually analyzed by regressions that take into account as many relevant variables as possible. The σ here refers to the residual standard deviation.

it can be given a concrete interpretation as follows: If promotion rate is held at its reference value, x^0 , there will be sales fluctuations (arising solely from fluctuations in the sales response parameter β) that are normally distributed and have a standard deviation of 3% of s^0 . The constant k has been taken as .9. If k had been set at 1.0, sales response in one period would be made up of last period's response plus a random fluctuation of the type just discussed. A value of .9 gives something fairly close to this but provides a tendency for sales to move toward the reference value, s^0 . Specifically, in the absence of new fluctuations, the difference between actual sales and the reference value would be reduced by 10% in each time period.

To summarize and add detail:

$$\alpha^0 = .32 = \text{sales rate in absence of any promotion. (dol./hh.yr.)}$$

$$\beta^0 = 9 = \text{long run average of sales response parameter } \beta. \\ \text{(dimensionless)}$$

$$\delta = 100 = \text{curvature parameter of sales response function.} \\ \text{(dol./hh.yr)}^{-1}$$

$$m = 1/3 = \text{gross profit as a fraction of selling price.} \\ \text{(dimensionless)}$$

These lead to reference values, calculated from (8.6) and (8.7):

$$x^0 = .03 = \text{reference promotion rate. (dol./hh.yr.)}$$

$$s^0 = .50 = \text{reference sales rate. (dol./hh.yr.)}$$

$$x^0/s^0 = 6\% .$$

The model of sales response fluctuations has parameters:

$$\sigma_\beta = .5 = \text{period to period standard deviation of } \beta(t). \\ \text{(dimensionless)}$$

$$k = .9 = \text{persistence constant for } \beta(t). \text{ (dimensionless)}$$

The data related to experimental design are:

$\sigma = .07 s^0 = .035$ = standard deviation of sales rate for an individual market. (dol./hh.yr.)

$\Delta = .50 x^0 = .015$ = range of experimental deviation in promotion rate. (dol./hh.yr.)

$N = 1000$ = number of individual markets of "average" size required to make up national sales.

Experiment: To design the experiment, the appropriate constants are substituted into (8.11), which can be solved to give $z = 5.70$, then by (8.9), $n\Delta^2 = .00687$. Since we have arbitrarily fixed Δ at .015, n becomes 30.5 or, rounding,

$n = 30$ markets.

This completes the experimental design. In each time period 30 markets will be run at a promotion rate $.25x^0$ above the national rate, $x^0(t)$, and another 30 markets will be run the same amount below.

The resulting experimental standard error of β is calculated from (4.4) to be .602. This does not make a particularly accurate experiment, considering that we would like to operate on the sales response curve at the point with slope $1/m = 3.0$. Under the above standard error, we could fairly easily get a measured slope of 2.9 while we were actually operating at 3.5. However, two factors are present that make the situation less serious than it might appear. First of all, the current experiment does not represent all the available information about sales response. Before setting the promotion rate, current information is combined with past information by the decision rule. Secondly, profit maximizations of the present type have the property that, as long as the control variables can be kept fairly close to their optimal values, losses will be small.

Decision Rule. In order to set up the decision rule, we need the smoothing constant \underline{a} . First v' is found from (6.7) to be .451 . Then (6.6) gives

$$a = .446$$

and $(1-a) = .554$. This means we shall be giving fairly equal weight to new and old information. By combining (5.3) and (5.4), we can write the decision rule for $x_o(t)$ as

$$\dot{x}_o(t) = a x_o(t-1) + [(1-a)/2\delta] [\hat{\beta}(t-1) - (1/m)]$$

or, substituting numbers,

$$(9.1) \quad x_o(t) = .446 x_o(t-1) + .00277 [\hat{\beta}(t-1) - 3.0] .$$

We are now set to operate. Given a starting value, $x_o(0)$, and a sequence of experimental results, $\hat{\beta}(t)$, we can generate the promotion rate $x_o(t)$.

Simulation: To see how the system behaves, we have simulated the underlying market over time and have operated with the above rule. The steps in the simulation are briefly as follows: The values for k and β^o are put into (3.2) to give:

$$(9.2) \quad \tilde{\beta}(t) = .9 \tilde{\beta}(t-1) + .9 + \tilde{\epsilon}_{\beta}(t) .$$

Then random numbers, $\epsilon_{\beta}(t)$, having mean zero and variance $\sigma_{\beta}^2 = 1/4$, are substituted into (9.2) to generate a sequence of values for $\beta(t)$. For the $\alpha(t)$ process we take $\alpha(t) = \alpha^o = .32$ dol./hh.yr. for all t . A varying of $\alpha(t)$ might be more realistic but, since $\alpha(t)$ does not enter into any decisions, we have simply made it a constant.

Given $\beta(t)$, the experimental results at t are simulated by

$$\hat{\beta}(t) = \beta(t) + \epsilon_{ex}(t) .$$

Here $\epsilon_{ex}(t)$ is a random normal number with mean zero and variance $v = .363$. (We could simulate each of the $2n$ test markets separately to generate $\hat{\beta}(t)$, but it is equivalent and much easier to simulate the final experimental result directly.)

The company uses $\beta(t)$ to generate $x_o(t)$. The final sales outcome can be calculated from this, $\beta(t)$, α , and γ :

$$s(t) = .32 + \beta(t) x_o(t) - 100 [x_o(t)]^2 \text{ dol./hh.yr.}$$

The sales, in turn, can be used to calculate profit. Our principal criterion for judging the system, however, is the loss relative to the profit that could be made under perfect information. The best promotion rate under perfect information can be developed as a side calculation using (8.1).

$$x^*(t) = [\beta(t) - 3] / 200 \text{ dol./hh.yr.}$$

The loss rate relative to this is, from (2.4):

$$l(t) = 33.3 [x_o(t) - x^*(t)] \text{ dol./hh.yr.}$$

The further loss resulting from the experimental deviations has not been simulated, but its expected value can be calculated from (8.5).

The simulation results are shown in Figure 5. 40 time periods are shown. (The series was started with $x(0) = x^0$ and $\beta(0) = \beta^0$ and run for 10 periods before the present data were taken.) Plotted are $\beta(t)$, which is driving the system, $x^0(t)$, by which the company responds, and the resulting $s(t)$ and $l(t)$. The latter is expressed as a % of x^0 .

We see that the response of the adaptive system is quite good. Notice that responses in $x(t)$ lag changes in $\beta(t)$ by a time period since the information gained during one period is not available until the next. The losses are generally small, although occasional peaks




Figure 5. Simulation of adaptive control

occur where $\beta(t)$ has changed substantially and $x(t)$ has not yet caught up. The 40 period average loss is 1.55% of x^0 .

Expected losses. The expected losses relative to an operation based on perfect information can be calculated from various formulas. We shall use the exact loss formulas given in the Appendix instead of those of Section 8, which are based on the $k=1$ approximation. The differences, however, are slight.

| <u>Source</u> | <u>Expected Loss Rate</u> (compared to perfect information) | | |
|-------------------------|--|-------------------|--------------------------------|
| | <u>(dol./hh.yr.)</u> | <u>(dol./hr.)</u> | <u>(% of x^0)</u> |
| $x_0(t)$ not perfect | .000376 | 18,500 | 1.23 |
| experimental deviations | .001875 | 5,700 | .38 |
| Total | | 24,200 | 1.61% |

The 40 period average loss rate of 1.55% in the simulation is reasonably close to the calculated 1.23%. These losses are encouragingly small, especially since the standard is perfect knowledge of the response curve and since our experiment is not terribly precise.

Comparison with other policies: The results using the optimal adaptive system may be compared with other types of operations. One possibility is to set $x(t)$ to a constant value, say x^0 , for all t . The rate x^0 has the property of being optimal when $\beta(t)$ takes on its long run average value. It is not clear, of course, how the company would figure out x^0 . Consequently the values $.5x^0$, x^0 , and $1.5x^0$ will be tried. Expected losses compared to perfect information can be computed from (2.4) and (A.10). For a constant promotion rate there is no cost of experimental deviations.

| | <u>Expected Loss Rate</u> (compared to perfect information) | |
|---------------------------|--|--------------------------------|
| | <u>(dol./yr.)</u> | <u>(% of x^0)</u> |
| adaptive system: $x_0(t)$ | 24,200 | 1.61% |
| constant rate: x^0 | 54,700 | 3.65 |
| constant rate: $.5x^0$ | 430,000 | 28.65 |
| constant rate: $1.5x^0$ | 430,000 | 28.65 |

Thus the increased loss ranges from 2% of x^0 for the lucky guess to 27% of x^0 for the substantial deviation.

Suppose that circumstances prevent an experiment of the proposed optimal size. Instead of 30 markets in each group, suppose we have to get along with 15. Then the experimental error would increase from .602 to .726. The value of \underline{a} would decrease from .561 to .446, implying more reliance on past information. The system, however, is still adaptive.

| | <u>Expected Loss Rate</u> (compared to perfect information) | |
|---------------------------|--|--------------------------------|
| | <u>(dol./yr.)</u> | <u>(% of x^0)</u> |
| optimal experiment (n=30) | 24,200 | 1.61% |
| smaller experiment (n=15) | 25,500 | 1.70 |

The deterioration in performance is small.

10. Sensitivity Analysis

The optimal adaptive system performed well in the example of the previous section, but the optimality of the system is based on various assumptions about the underlying operation of the market. Many of these assumptions are questionable. How will the system perform if some of them are incorrect? To investigate this question,

we use the experiment and decision rule as derived but change the underlying market model in various ways.

Constant Sales Response. The original model assumed that sales response fluctuates over time. Suppose that, instead, sales response is constant and has $\beta(t) = \beta^0$. Then the optimal promotion rate is $x(t) = x^0$. If we use the adaptive system of the previous section some relative loss will be incurred. The loss can be calculated using the formulas in the appendix by setting $\sigma_\beta = 0$.

| | <u>Expected Loss Rate</u> (compared to perfect information) | |
|-------------------------|--|--------------------------------|
| | <u>(dol./yr.)</u> | <u>(% of x^0)</u> |
| adaptive system | 11,000 | .74 |
| constant rate: x^0 | 0 | 0 |
| constant rate: $.5x^0$ | 375,000 | 25. |
| constant rate: $1.5x^0$ | 375,000 | 25. |

We see that the average loss rate using the adaptive system is only .74% of x^0 . This is quite small; in fact, it is less than the 1.61% that occurred when sales response fluctuated. (The reason is that the stable sales response is inherently easier to optimize.) Of course, if the company president is clairvoyant, he will pick $x(t) = x^0$ and have no loss at all. On the other hand, if he misses and sets $x(t) = .5x^0$ or $1.5x^0$, a substantial loss is incurred.

Changes in Sales Response Have No Persistence. The earlier model assumed that sales response changed but that the starting point for the change was the previous sales response. Thus a goal of the adaptive system was to follow sales response as it drifted about. Suppose instead that sales is subject to fluctuation but has no persistence.

In other words, $\beta(t)$ equals β^0 plus an independent random variable for each t . This situation may be obtained in our model by setting $k = 0$. Under these circumstances the optimal policy is $x(t) = x^0$.

| | <u>Expected Loss Rate</u> (compared to perfect information) | |
|-------------------------|--|--------------------------------|
| | <u>(dol./yr.)</u> | <u>(% of x^0)</u> |
| adaptive system | 25,800 | 1.72 |
| constant rate: x^0 | 10,400 | .70 |
| constant rate: $.5x^0$ | 386,000 | 25.7 |
| constant rate: $1.5x^0$ | 386,000 | 25.7 |

The adaptive system is not as good as the constant rate x^0 but the difference is small. On the other hand if a constant rate substantially different from x^0 is picked, large losses are incurred.

Different curvature. Perhaps the most serious assumption in the model is that γ , the curvature parameter in sales response, is a known constant. In practice, γ will usually be unknown or poorly known. We suggested earlier that perhaps our ignorance of γ was not too serious. To test this possibility we make some fairly drastic changes in γ but set promotion rate by the adaptive system derived from the old value.

First, γ is reduced from 100 to 25, a factor of 4. Sales response is now much more linear.

| | <u>Expected Loss Rate</u> (compared to perfect information) | |
|-------------------------|--|--------------------------------|
| | <u>(dol./yr.)</u> | <u>(% of x^0)</u> |
| adaptive system | 75,000 | 5.01 |
| constant rate: x^0 | 219,000 | 14.6 |
| constant rate: $.5x^0$ | 312,000 | 20.6 |
| constant rate: $1.5x^0$ | 312,000 | 20.6 |

The losses compared to perfect information are found to be higher than they were for $\gamma = 100$ for both the adaptive system and the constant rate, x^0 . The more linear response means that a change in β can be made the basis for a substantial and profitable change in promotion rate (if response is known perfectly). Without perfect information, both ways of operating fumble more than formerly. However, the adaptive system does much better than the best constant value.

Next we increase γ by a factor of 4 to 400. Sales response is now much more nonlinear.

| | <u>Expected Loss Rate</u> (compared to perfect information) | |
|-------------------------|--|--------------------------------|
| | <u>(dol./yr.)</u> | <u>(% of x^0)</u> |
| adaptive system | 27,100 | 1.81 |
| constant rate: x^0 | 13,600 | .91 |
| constant rate: $.9x^0$ | 73,600 | 4.91 |
| constant rate: $1.1x^0$ | 73,600 | 4.91 |

Here a clairvoyantly picked constant rate is somewhat better than the adaptive system. However sales and profit rates are very sensitive to promotion. Constant rates of $.5x^0$ and $1.5x^0$ lead to unrealistic positions on the sales response curve and consequent large losses. Therefore, we have substituted the milder rates of $.9x^0$ and $1.1x^0$. These values, although quite close to the best constant rate are seen to perform worse than the adaptive system.

Behavior with a Step Function. Insight into the behavior of the adaptive system can be gained by seeing how it reacts to a step change in β . Figure 6a shows a $\beta(t)$ that goes along as $\beta^0 = 9$ until $t = 1$ and then jumps to 15. Under perfect information optimal




Figure 6. Behavior of average $x(t)$ when $\beta(t)$ jumps.

operation would be $x^0 = .03$, jumping to $.06$ at $t = 1$. Figure 6b shows the expected value for $x(t)$ under the adaptive system of the previous section. (Because of randomness in the experimental results, actual operation would fluctuate about the expected value.) We see that average $x(t)$ lags one time period before starting to respond and then responds fairly quickly.

In summary, we have devised sensitivity tests to see how the adaptive system performs when the assumptions on which it is based are violated in various ways. These tests have been applied to the numerical example of the previous section. The performance of the adaptive system is found to be good.

11. Discussion

Perhaps it is surprising that such simple operating rules can cope successfully with a changing environment. 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 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 system overcomes some of the inaccuracy of its measuring device by working from an accumulation of present and past information. In the second place, sales response is presumed (reasonably enough) to vary smoothly with spending. As a result, the profit maximum is also smooth. This means that small deviations from best operation cost very little: underspending gives fewer sales but saves almost an equal amount in out-of-pocket expense; overspending is almost counterbalanced by increased sales. 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. (The rapid diminishing returns of the log curve used by Benjamin and Maitland [2] can be interpreted this way.) 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.*

The use of the adaptive system requires a tolerance of explicit uncertainty on the part of a company's management. The system will specify some definite number for promotion rate, but the available information on response may be sufficiently ambiguous that other, rather different numbers look almost as good. Such a situation may be disconcerting. Some people prefer the pseudo-certainty of a plan that is defended as exactly right (even though something quite different was done the previous year under roughly the same circumstances.) The adaptive system must be regarded as a set of operating rules that produce a good average return rather than as a device for producing the perfect number each time. In fact, uncertainty plays an essential role in the system since inaccuracy is deliberately accepted in the optimal design.

11.1 Practical Problems. We shall try to anticipate and discuss some practical questions that might come up using a system like the one presented here. First of all, there is a class of questions relating to setting up the system: Is the model as it stands sufficiently

* There is an extremely interesting question implicit here: If some aspect of promotional treatment, for example, advertising copy, is exceptionally good, should a company spend more or less on promotion? The answer does not seem obvious.

good or should certain complications be added? Several of the parameters of the model will almost certain be unknown; how should they be picked? The decision rules have been derived under the assumption of steady state operation; should anything different be done in starting up?

One approach to handling the uncertainties behind these questions would be to develop a formal Bayesian analysis of them. However, the approach suggested by the work here is to gather together whatever relevant information can be found, build a specific model, and investigate its behavior by simulation and sensitivity analysis. For starting up, the past promotion rate of the company could be used as the past promotion rate required by the model. It might be desirable to design extra accuracy into the first few experiments. As a practical matter, however, experimental accuracy will probably increase rather than decrease with time because of increasing experience.

Over a number of time periods, information can be built up on the validity of the model and better estimates can be made of the constants. Particularly important is information about δ , since δ expresses the degree of diminishing returns. Notice that our proposed operation collects information about δ even though we have not acknowledged this fact in the analysis. The information is the result of using three spending levels ($x_0 - (\Delta/2)$, x_0 , $(x_0 + (\Delta/2))$). These permit an estimate of δ in each time period. Individual estimates are likely to be quite unreliable, however. This is one reason we have chosen not to build the measurement of δ directly into the analysis.

Another practical problem is that the decision rules may occasionally call for a really substantial change in spending rate. In

some cases the proposed rate may fall outside the range included in the experiment. Substantial changes are often disturbing to an organization and changes beyond the experimental conditions are on less solid empirical ground than those within. One way to handle the situation is by exercise of the managerial override that obviously exists on the whole system. A better way, however, and one that can be pre-planned and pre-studied is to clamp the amount of change permitted, say, by limiting it to $\pm 15\%$ or $\pm 20\%$. This will tend to slow down system response, but perhaps not excessively so.

Different market areas, it may be argued, will have different sales responses to promotion, whereas the measurements discussed here produce an average response. It seems certainly true that markets will differ, at least to some degree. However, whether the differences are appreciable will depend on the situation. Where they are, it would be desirable to take advantage of them. Possibly, one can develop an adaptive system that applies to individual markets, but the measurement problems appear to be quite difficult. A feasible approach would be to use whatever empirical and judgmental information is available to develop individual market adjustments to apply to the average curve. In this case each market would have its own $x_o(t)$ and experimental deviations would be made relative to this.

A related problem concerns the nature of the test markets. Assuming, as we generally have, that the basic experimental unit is a market area, the test markets themselves will usually be medium-sized markets. The very large and very small tend to be ruled out by various operating considerations. Yet, we wish to set promotional

rate throughout the country. This is basically a problem of individual market adjustments and can be handled as discussed above.

The experimental design discussion of Section 8 omits a cost that may sometimes be appreciable, namely, the out-of-pocket cost of running the experiment. If the cost is substantial, it can easily be included in the analysis for determining the optimal experiment. If there is a cost proportional to the number of test markets, the quantity $n\Delta^2$ will no longer be an invariant constant. Instead, the calculation will show that total cost is least when the number of markets, n , is as small as possible and the experimental deviation, Δ , in the promotion rate is as large as possible. We do not want Δ to become so large that sales are substantially reduced in the low markets. Although this consideration could be included in the formal analysis, a simpler procedure is to place an arbitrary upper limit on Δ . Then n can be calculated in a straightforward way.

The effect of promotion may be delayed. Empirical data (see references 3, 4, 5, and 6) suggest that the principal response is frequently fairly rapid, say, within one to three months. If sales are measured by factory shipments, a pipeline delay will also be encountered. A nominal value might be 1-1/2 months. In the discussion here we have obviously assumed that delays are small enough that a measurement of the promotional effort can be made within the experimental time period. A period of a year seems appropriate for many situations. A year also eliminates certain problems of seasonality and fits in with the budgeting process.

Any control system should be examined for stability. Conceivably, an inappropriate choice of constants relative to the underlying process could result in too small a value of the smoothing

constant, a, and so an inappropriately large dependence on the most recent data. This could perhaps lead to an oscillation of under-spending alternating with overspending. However, this does not seem likely for the type of decision rule being used. In any case, suspected sources of oscillation can be investigated in advance by simulation.

11.2 Competition. Businessmen frequently are concerned about a possible self-defeating aspect in promotional competition. If the effect of promotion is primarily on market share and not on total demand and, if promotional efficiency is comparable from one company to another in the industry, then promotional increases may appear profitable in experiments. Yet, when the increases are applied nationally, they may be countered by competing companies in such a way that nobody's sales change much but everybody's spending is increased. Models of this sort of process have been built by a number of writers. Mills'⁸ paper contains an example.

Several remarks can be made. Let us consider the extreme case where promotion affects only market share and not total industry sales. If the companies are operating at competitive equilibrium, i.e., the companies are already individually operating (within their measurement capabilities) in the neighborhood of their independent maximum profit points, then the adaptive system will continue this type of operation in an efficient manner. Whenever a company starts over- or under-spending, the adaptive system will tend to return the spending rate to the maximum profit position.

If, on the other hand, all companies are spending less than they would at competitive equilibrium, i.e., an accurate experiment would

indicate that increases would be profitable for each company if the others stayed constant, then spending rates tend to be unstable. A company not wishing to disturb the situation may choose to hold promotional spending constant or to adopt a figure close to the industry average spending as a per cent of sales. Either way, there is no need for an adaptive system to set promotional spending for it is already set. (There may be some other interesting applications, however, in the allocation of funds between promotional alternatives.)

Another policy for this situation, but one that is more likely to be misunderstood by the competition, is to set stiffer return requirements on promotional spending than those implied by the conventional maximum profit calculation. This could be done by using a smaller value for the gross margin than actually was the case.

It is a rare company that knows whether it is in competitive equilibrium with respect to promotional spending and whether its spending appreciably affects industry sales. Some information about these questions can be obtained in experiments of the type we have been considering, although, as we have discussed, the information is likely to deteriorate unless kept up to date.

11.3 Extensions. We mention three particularly desirable extensions of the model. First it would be helpful to have an adaptive system to allocate a fixed budget between several promotional alternatives. Second, it is unrealistic and undesirable to rely solely on past experiments to estimate sales response. Certain other pertinent information is usually available; knowledge of product changes, forecasts of economic conditions, etc. A step could be introduced into the feedback loop to bring this information into the

prior distributions used to set promotional rate. Finally, it would be worthwhile to study the possibility of monitoring individual market response by time series analysis as a basis for individual market adjustments.

12. Conclusions

The adaptive system discussed here is directed toward a major continuing problem: the setting of spending rate for promotion. The concept of the adaptive system seems basically correct: a company should learn from experience in an organized way. The model studied is a simple one but it may be useful as it stands and it certainly is capable of extension. The operating rules that have been developed are simple, feasible, and seem intuitively reasonable. In examples using realistic numbers, system performance has been good, despite a relatively inaccurate measurement process. Of particular interest is the insensitivity of performance to substantial changes in the underlying model.

Appendix

The mathematical development so far has been directed toward finding optimal operation, particularly under the $k=1$ approximation. However, for sensitivity analyses it is helpful to have exact expressions for the expected loss rate under various kinds of non-optimal operation. Accordingly we here solve the following problem:

Given (1) the basic environmental models

$$(A.1) \quad s = \alpha + \beta(t)x - \gamma x^2$$

$$(A.2) \quad \beta(t) = k \beta(t-1) + (1-k)\beta^0 + \epsilon_{\beta}(t)$$

(2) the decision rules

$$(A.3) \quad x_0(t+1) = a x_0(t) + (1-a) \hat{x}_0(t)$$

$$(A.4) \quad x_0(t) = \left[m \hat{\beta}(t) - 1 \right] / 2m\gamma$$

(3) the stochastic process describing the experiment

$$(A.5) \quad \hat{\beta}(t) = \beta(t) + \epsilon_{ex}(t)$$

(4) the loss rate formulas

$$(A.6a) \quad \lambda(t) = m\gamma \left[x_0(t) - x^*(t) \right]^2$$

$$(A.6b) \quad \lambda_i(t) = m\gamma \left[x_i(t) - x^*(t) \right]^2, \quad i=1,2$$

Find the steady state expected loss rate L .

We shall do this for an arbitrary (not optimal) smoothing constant, a , and an arbitrary (not optimal) experiment of accuracy $V(\epsilon_{ex}) = v = 2\sigma^2/n\Delta^2$. The calculation will proceed by solving difference equations to express $\beta(t)$, $x^*(t)$, and $x_0(t)$ as infinite sums of independent random variables. The desired expected loss can be calculated from these.

Successive substitutions in (A.2) give

$$(A.7) \quad \beta(t) = \beta^0 + \sum_{j=0}^{\infty} k^j \epsilon_{\beta}(t-j),$$

so that in steady state

$$(A.8) \quad E [\beta(t)] = \beta^0$$

$$(A.9) \quad V [\beta(t)] = \sigma_{\beta}^2 / (1-k^2).$$

Putting (A.7) into

$$x^*(t) = [m \beta(t) - 1] / 2\delta m$$

gives

$$(A.10) \quad x^*(t) = x^0 + \frac{1}{2\delta} \sum_{j=0}^{\infty} k^j \epsilon_{\beta}(t-j)$$

Therefore in steady state

$$(A.11) \quad E [x^*(t)] = x^0$$

$$(A.12) \quad V [x^*(t)] = (1/4\delta^2) \sigma_{\beta}^2 / (1-k^2)$$

The calculation of $x_0(t)$ is a little longer. Successive substitution in (A.3) gives

$$(A.13) \quad x_0(t+1) = (1-a) \sum_{j=0}^{\infty} a^j x_0(t-j) .$$

From (A.5) and (A.7) we have

$$\hat{\beta}(t) = \beta^0 + \sum_{j=0}^{\infty} k^j \epsilon_{\beta}(t-j) + \epsilon_{ex}(t),$$

so that, from (A.4)

$$\hat{x}_0(t) = x^0 + (1/2\delta) \sum_{j=0}^{\infty} k^j \epsilon_{\beta}(t-j) + [\epsilon_{ex}(t)] / 2\delta .$$

Putting this in (A.13), we finally obtain

$$(A.14) \quad x_o(t+1) = x^o + \left[(1-a)/2\delta(k-a) \right] \sum_{s=0}^{\infty} (k^{s+1} - a^{s+1}) \epsilon_{\beta}(t-s) \\ + \left[(1-a)/2\delta \right] \sum_{j=0}^{\infty} a^j \epsilon_{ex}(t-j) .$$

In steady state, therefore,

$$(A.15) \quad E [x_o(t+1)] = x^o$$

$$(A.16) \quad V [x_o(t+1)] = \left[(1-a)/4\delta^2(1+a) \right] \left[\sigma_{\beta}^2(1+ak)/(1-ak)(1-k^2) + \sigma_{ex}^2 \right]$$

Some manipulation using (A.14) and (A.10) gives the following result to be used shortly:

$$(A.17) \quad E \left[\{x_o(t+1) - x^o\} \{x^*(t+1) - x^o\} \right] = \\ \left[\sigma_{\beta}^2/4\delta^2 \right] \left[k(1-a)/(1-k^2)(1-ak) \right] .$$

Turning now to the loss rates, we observe that (A.6a) can be written:

$$\ell(t) = m\delta \left\{ \left[x_o(t) - x^o \right]^2 - 2 \left[x_o(t) - x^o \right] \left[x^*(t) - x^o \right] \right. \\ \left. + \left[x^*(t) - x^o \right]^2 \right\} ,$$

so that

$$E [\ell(t)] = m\delta \left\{ V [x_o(t)] + V [x^*(t)] \right. \\ \left. - 2 E \left[\{x_o(t) - x^o\} \{x^*(t) - x^o\} \right] \right\} .$$

Using now (A.16), (A.12) and (A.17), we have

$$(A.18) \quad E [\ell(t)] = \left[m/4\delta \right] \left[(1-a)\sigma_{ex}^2/(1+a) + \right. \\ \left. 2\sigma_{\beta}^2/(1+k)(1+a)(1-ak) \right]$$

In the case of the experimental markets the analysis of Section 8 holds so that

$$(A.19) \quad E [l_1(t)] = E [l_2(t)] = E [l(t)] + m\delta\Delta^2/4 .$$

Weighting (A.18) and (A.19) by the number of markets and expressed loss rate as a fraction of x^0 , we obtain the desired quantity

$$(A.20) \quad L = [m/4\delta x^0] \left\{ [(1-a)/(1+a)] [2\sigma^2/n\Delta^2] + 2\sigma_\beta^2 / (1+k)(1+a)(1-ak) \right\} + m\delta n\Delta^2 / 2Nx^0$$

Finally it is of some interest to find $E[s(t)]$:

$$(A.21) \quad E [s(t)] = s^0 + [(1-a)/(1+a)4\delta] [\sigma_\beta^2 (ak+2k-1) / (1-ak)(1-k^2) - \sigma_{ex}^2]$$

Notice that only under special circumstances will $E[s(t)] = s^0$.

References

1. H. Raiffa and R. Schlaiffer, Applied Statistical Decision Theory, Boston: Graduate School of Business Administration, Harvard University, 1961.
2. B. Benjamin and J. Maitland, "Operational Research and Advertising: Some Experiments in the Use of Analogies," Operational Research Quarterly, 9, 207-217, (September 1958).
3. M. L. Vidale and H. B. Wolfe, "An Operations Research Study of Sales Response to Advertising," Operations Research, 5, 370-81, (June 1957). (Reprinted in Ref. 7).
4. B. Benjamin, W. P. Jolly, and J. Maitland, "Operational Research and Advertising: Theories of Response," Operational Research Quarterly, 11, 205-218, (December 1960).
5. A. A. Brown, F. T. Hulswit, and J. D. Kettelle, "A Study of Sales Operations," Operations Research, 4, 296-308 (June 1956). (Reprinted in Ref. 7).
6. P. L. Henderson, J. F. Hind, and S. E. Brown, "Sales Effect of Two Campaign Themes," Journal of Advertising Research, 1, 6, 2-11 (December 1961).
7. Bass et al., editors, Mathematical Models and Methods in Marketing, Richard D. Irwin, Inc., Homewood, Illinois, 1961.
8. H. D. Mills, "A Study in Promotional Competition," 271-288 in Ref. 7.

