K. Sridhar Moorthy

Theoretical Modeling in Marketing

Over the last 10 years or so, theoretical modeling has rapidly become an important style of research in marketing. To many people, however, this style is still a mystery. This article is an attempt at explaining theoretical modeling. The author argues that even though theoretical modeling is quantitative, it is closer to behavioral marketing in purpose and methodology than to quantitative decision support modeling. Whereas behavioral marketing involves empirical experiments, theoretical modeling involves logical experiments. Using this framework, the author addresses such issues as the internal and external validity of theoretical models, the purpose of theoretical modeling, and the testing of model-based theories. The agency theory explanation of salesforce compensation is used as a case study.

AN essentially new style of research has sprung up in marketing recently: mathematical theoretical modeling.[1] Scarcely an issue of *Marketing Science* passes without an article in this style. Some examples are the articles by McGuire and Staelin (1983), Moorthy (1984), Basu et al. (1985), Mahajan and Muller (1986), Hess and Gerstner (1987), Hauser (1988), Wilson and Norton (1989), and Rao (1990). Lately, theoretical modeling seems to have invaded the *Journal of Marketing Research* as well (e.g., Hauser and Wernerfelt 1989; Lal 1990; Wilson, Weiss, and John 1990).

To the nonparticipant, the popularity and growth of theoretical modeling may seem like an oddity, a passing fad. The method seems to violate all the norms of good research. The articles are (generally) all theory, no data. The assumptions are unrealistic. Managerial implications are difficult to find. To make matters worse, the reader must wade through countless lemmas, propositions, theorems, proofs. It is legitimate to ask: What is all this in aid of? How does the methodology work? Why is it useful to marketing? How can we apply these models? How can we test these models? How does quantitative theorizing differ from the verbal theorizing in the "behavioral" literature and the quantitative models in the decision support system literature?

This article is an attempt at answering these questions in an informal way. It is not meant to be a philosophical discussion of research methodology, but rather a user's guide to one style of research. (For a more formal treatment, see Cook and Campbell 1979; Hunt 1991; Suppe 1977.) The principal aim is to relate theoretical modeling to the other research paradigms in marketing, so that the method becomes accessible to

[1]In the title of the article and elsewhere, "mathematical" is dropped and the term "theoretical modeling" is used. A theoretical model need not be mathematical (cf. the verbal and graphic models in "behavioral marketing": Bettman 1979; Puto 1987; Sujan 1985; Wright 1975) and a mathematical model need not have a theoretical purpose. For example, most mathematical models in marketing are really *measurement* models—models set up to estimate demand functions (Hanssens, Parsons, and Shultz 1990).

K. Sridhar Moorthy is Associate Professor of Marketing, William E. Simon Graduate School of Business Administration, University of Rochester. The article was written while the author was visiting the Anderson Graduate School of Management at UCLA in May 1990. Preliminary ideas were presented at the American Marketing Association's 1989 Doctoral Consortium at Harvard University, the 1989 Summer Marketing Educators' Conference in Chicago, the Wharton School, and the University of Toronto. The author thanks the participants at these meetings for their comments and encouragement. He is especially grateful to Joao Assuncao, Rajiv Lal, Gary Lilien, Rick Staelin, and the anonymous *JM* reviewers for their detailed comments on a previous version of the article.

Journal of Marketing
Vol. 57 (April 1993), 92–106

a broad array of marketing academicians and practitioners. At the same time, the perspective provided here should be useful to theoretical modelers as they appraise their own and others' works.

The main theme of the article is that mathematical theories are built by a process of *logical experimentation*, much like the empirical experimentation in behavioral marketing research (e.g., Bettman 1979; Puto 1987; Sujan 1985; Wright 1975). Theoretical models are the "treatments" in this logical experiment. They are thus quite different from the quantitative models in the decision support and measurement literature in marketing (e.g., Bultez and Naert 1988; Little 1970, 1975, 1979; Lodish 1971; Silk and Urban 1978). Even though both types of models use mathematics, in purpose and methodology they are far apart. By the same token, even though theoretical modeling is different from behavioral marketing research in its mathematical level, in purpose and methodology they are very similar.

In the next section, this view of theoretical modeling is developed further. Then a lengthy case study of theoretical modeling is presented. The example used is the "agency theory" explanation of salesforce compensation practices, which has received a lot of attention lately (Basu et al. 1985). This example is used throughout the article. Among the topics discussed are the internal and external validity of theoretical modeling experiments, the purpose of theoretical modeling, the relative importance of internal versus external validity given the purpose of theoretical modeling, the realism of modeling assumptions, and the testing of model-based theories. Finally, theoretical modeling is related to the other two major research styles in marketing: behavioral marketing and decision support modeling.

Overview of Theoretical Modeling

Theoretical modeling begins with the need to understand some marketing phenomenon. For example, we may want to understand why stores have sales, or why some manufacturers are vertically integrated into distribution and others are not. The researcher then constructs an *environment*—which he or she calls a *model*—in which the actions to be explained take place. A model is specified by a series of *assumptions*. Some assumptions are purely mathematical; their purpose is to make the analysis tractable. Other assumptions are *substantive*, with verifiable empirical content. They can describe such things as who the actors are, how many of them there are, what they care about, the exogenous conditions under which they make decisions, what their decisions are about, and so on. (In marketing models, the actors usually are manufacturing firms, channel intermediaries, or consumers.) Only the substantive assumptions participate in the explanation being offered.

For example, in Hauser's (1988) model of product and price competition, two (or three) manufacturers are deciding on the configuration of their products. Each is assumed to offer only one product. Only two attributes have to be set for each product, and the two attributes are related by the production technology. Consumers' preferences are additive and linear in the two attributes. Every feasible product configuration has the same constant production cost for each firm. The firms choose their products first, simultaneously. Then, after committing to a product, each firm simultaneously chooses its price.

These assumptions, clearly, *do not describe* real-world markets. At best they define an artificial world with some connections to the real world. Thus the concept of a model in theoretical modeling is different from the concept of a model in decision support systems and behavioral marketing research. In decision support modeling, a model is a "mathematical *description* of how something *works*" (Little 1979, italics added); in theoretical modeling, a model is simply a *setting* in which a question is investigated, a "laboratory."[2] Hence, whereas decision support models—because they are descriptions of how things work—emphasize realism, theoretical models—because they are laboratories—are necessarily unrealistic. A theoretical model is also different from a behavioral model. The latter is a verbal or graphic *description* of the researcher's *theory*. For example, Puto (1987) describes his "proposed conceptual model of the buying decision framing process" graphically. Sometimes behavioral researchers refer to their measurement model (e.g., a regression equation) as *the* model, even though a model describing their theory also exists.

Once a theoretical model has been built, the researcher analyzes its *logical* implications for the phenomenon being explained. Then another model, *substantively different* from the first, is built—very likely by another researcher—and its implications are analyzed. The process continues with a third and a fourth model, if necessary, until all ramifications of the explanation being proposed have been examined. By comparing the implications of one model with those of another, and *tracing* the differences to the model design, we hope to understand the cause-effect relationships governing the phenomenon in question. This is as though a logical experiment were being run, with the various models as the treatments and the phenomenon being explained as the "dependent variables." The key difference from empirical experiments is that

[2]Subsequently a distinction is made between a supermodel and a model and the laboratory interpretation is reserved for the supermodel. For the present, this distinction is not necessary.

in empirical experiments the subjects produce the effects, whereas here the researcher produces the effects by logical argument. Theoretical modeling also resembles meta-analysis in some ways (Assmus, Farley, and Lehmann 1984). Both involve "post-model" analyses, but meta-analysis is used to discover the patterns in *empirical* results across a number of situations, whereas theoretical modeling's purpose is to construct cause-effect *explanations* of marketing phenomena. Figure 1 summarizes this view of theoretical modeling.

Theoretical Modeling in Action: A Case Study

Firms compensate their salesforces in a variety of ways, for example, salaries, commissions, quotas, sales contests, and free vacation trips. A natural question to ask is: Why is there such a variety of compensation schemes and what function does each compensation component serve in a compensation package? Suppose we consider just salaries and commissions. Basu et al. (1985) have used agency theory to explain these features of salesforce compensation plans.

Agency theory originates from economics, where it was developed to address situations in which a "principal" must use an "agent" to carry out certain actions. The principal cannot observe the agent's actions costlessly, so the question arises: What kind of contract should the principal offer the agent so that the agent is motivated to act in the principal's interest?[3] Notable contributors to the theory include Wilson (1969), Spence and Zeckhauser (1971), Ross (1973), Mirrlees (1976), Harris and Raviv (1979), Holmstrom (1979, 1982), Shavell (1979), Grossman and Hart (1983), Nalebuff and Stiglitz (1983), and Holmstrom and Milgrom (1987, 1990).

FIGURE 1
Overview of Theoretical Modeling

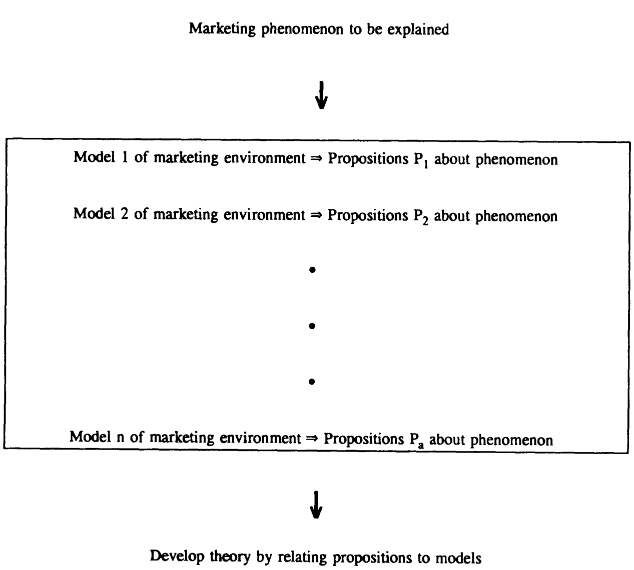

Stated verbally, the agency theory explanation of salaries and commissions is as follows. Salespeople, like most human beings, are risk averse. They prefer a stable, known income to a fluctuating, uncertain income, even if the latter is the same on average as the former. Salaries, by definition, lend stability and predictability to a compensation package, so they are used to reduce the income risk borne by salespeople. If all of the salesperson's income came as salary, however, he or she would have no incentive to work hard *given* that the firm cannot observe how hard each salesperson works. Commissions are used to motivate salespeople to work hard in situations where their effort cannot be observed. Thus, the use of salaries and commissions in compensation packages represents a tradeoff between reducing the income risk borne by salespeople and providing them the incentives to work hard.

Let us see how theoretical modeling helps us gain this understanding. The first step is to construct a "supermodel" specifying the overall environment in which the explanation will be constructed. Subsequently, we specify submodels of this supermodel and derive the logical implications of these submodels. This procedure is analogous to a behavioral researcher first specifying the overall boundaries of his or her experiment—which variables will be manipulated, what the context will be, how many (and which) levels of the variables will be used—and then actually running the experiment.

Supermodel

The following assumptions describe our supermodel.

- *Assumption 1*: A sales manager, representing the firm, is designing a compensation package for salespeople working independently.
- *Assumption 2*: The compensation package consists of a salary and/or commissions on the revenues generated by the salesperson. The sales manager designs the package and commits to it. The salesperson then accepts or rejects the compensation package offered. If he or she rejects, he or she will work somewhere else and get expected utility U_0.
- *Assumption 3*: Each salesperson's utility from income I and selling effort W is given by $U(I,W) = V(I) - W$.

[3]Strictly speaking, agency theory addresses situations in which the agent's actions cannot be *verified* by the principal, that is, the principal cannot prove in a court of law whether or not the agent carried out the desired actions. Observability of the agent's actions is necessary for verifiability, but not sufficient. We will, however, continue to use "observable" in place of "verifiable."

V is an increasing, twice-continuously differentiable, concave function of I. Salespeople decide how hard to work by maximizing their expected utility.

- *Assumption 4*: The manager designs the compensation package to maximize the firm's expected net profits, anticipating the salesforce's reaction. The net profits are given by $\pi - nI$, where π denotes the gross profits of the firm and n denotes the number of salespeople.
- *Assumption 5*: The gross profits of the firm are a function of W, the work put in by each salesperson, and ϵ, a random variable representing the uncertainty in the revenues generated. ϵ is independently and identically distributed across salespeople. Neither the manager nor the salesperson observes the resolution of this uncertainty. Both can, however, observe the revenues obtained. As the salesperson works harder, he or she shifts the distribution of π such that higher gross profit outcomes are more likely.
- *Assumption 6*: Assumptions 1 through 5 are known to the sales manager and the salespeople and both know this.

These assumptions have substantive and mathematical components. The distinction between the two is that the former are verifiable empirically (in principle), whereas the latter are not.[4] It is the substantive assumptions that define the marketing environment:

- Each salesperson's output is independent of any other salesperson's output (think about a situation in which each salesperson is selling a unique product).
- Only salaries and commissions are available as compensation elements.
- The manager proposes the compensation package on a take-it-or-leave-it basis (i.e., there is no room for negotiating compensation after the revenues have been realized).
- The manager is risk neutral; he or she is indifferent between getting *x* dollars for sure or getting a gamble with the same expected value.
- Salespeople, however, can be risk neutral (as just defined) or risk averse (i.e., prefer x dollars for sure in preference to the gamble)—both possibilities are admitted by the (weak) concavity assumption.
- Salespeople dislike putting in effort, and their dislike is independent of the amount of money they make.
- Salespeople cannot completely control the revenues they produce.

The mathematical content of assumptions 1 through 6 resides in assumptions 3 and 4. V is assumed to be twice-continuously differentiable, which means that the salesperson's utility is a sufficiently smooth function of income I. (It has no kinks or discontinuities, its slope has no kinks or discontinuities, and the slope of its slope has no discontinuities.) This assumption enables the researcher to use calculus as the primary analytical tool. The "maximization assumptions"—manager maximizing expected profits, salesperson maximizing expected utility—have mathematical and substantive content; they are difficult to verify empirically, but we can find situations in which the "stakes" are high enough for optimizing behavior to be a reasonable assumption.

Running the Experiment

Let us now construct a series of submodels (hereafter, simply "models") from this supermodel by specializing assumptions 1 through 6,[5] and state their logical implications for the optimal salesperson compensation contract. (The derivations of these implications are in the articles cited previously.) The models are the *treatments* in the experiment defined by the supermodel; various aspects of the optimal compensation scheme (e.g., salary, commissions, expected income, the firm's expected profits) are the *dependent variables*. See Figure 1.

Model 1 (salespeople are risk-neutral and their effort is observable). This is the simplest model to analyze. Given the salesperson's risk neutrality, it is immaterial whether the compensation package is all salary, all commissions, or any combination, as long as all options yield the same expected income to the salesperson. Furthermore, given that the salesforce's work is observable, the manager will design the compensation package such that if a salesperson does not work as hard as the manager would like, that person will be penalized severely. So the salesforce *will* work as hard as the manager would like and each member of the salesforce will get an expected income yielding utility U_0, the utility they would have gotten from the alternative job. The firm's expected profits will be as high as they can be.

Model 2 (salespeople are risk averse and their effort is observable). Borch (1962) has shown that under these circumstances an all-salary plan (with penalties as in model 1) is optimal whereas an all-commissions plan is not. The reason is that with all commissions the salesperson's income will fluctuate, so for any effort level his or her expected utility will be lower than it would be if he or she were given the same expected income in salary. Therefore, the manager who wants the salesforce to put out effort W and get an expected utility U_0 will have to pay them more compensation on average with commissions than with salary. With the optimal all-salary plan, however, the model 1 results are replicated.

[4]The standards of verifiability are the generally accepted standards of evidence in the profession. Formal empirical tests may not be necessary to meet this standard. For example, most marketing researchers would willingly accept, even without a formal empirical test, an assumption such as: *ceteris paribus*, a salesperson's utility increases with monetary income.

[5]For each model, the defining special assumptions are in parentheses.

Model 3 (salespeople are risk neutral and their effort is not observable). Now, a pure commissions compensation scheme is optimal for the firm and the salesperson *will* work as hard as he or she did under model 1 (Harris and Raviv 1979). The commission scheme, however, cannot be the same as in model 1. Now the commission rate will be such as to give the salesperson *all* of the firm's gross profits from the product he or she sells. The firm will make its money by asking the salesperson to pay a lump-sum amount equal to the firm's net profits under model 1. Essentially the manager is selling the product to the salesperson for a lump-sum price. Both firm and salesperson will be as well off as they were in model 1.

Model 4 (salespeople are risk averse and their effort is not observable). This is the most complicated case. To analyze it, Holmstrom (1979) makes two additional substantive assumptions: the distribution of gross profits (1) satisfies the monotone likelihood ratio property and (2) is "convex" (Grossman and Hart 1983). The distribution of sales satisfies the monotone likelihood ratio property if an observation of high sales is more likely to reflect high effort on the part of the salesperson than low effort. Convexity means (loosely) that the probability of observing high gross profits is higher with "average" effort than with a 50-50 combination of high and low effort. (The gamma distribution used by Basu et al. 1985 has these properties.) Holmstrom then shows that any additional signal of salesperson effort will increase the firm's expected profits if and only if it adds information. Basu et al. (1985) show that the optimal compensation package must involve both salaries and commissions. Furthermore, the salesperson will not work as hard as he or she did under model 2.

In addition, the following "comparative-statics" results obtain with a gamma distribution[6] for π and a specific power function for the utility function of the salesperson ($V(I) = I^{\delta}/\delta$ with $\delta = 1/2$):

1. The greater the responsiveness of gross profit variance to the salesperson's effort, the less the salesforce works, the less its expected income, the less the firm's expected profits, and the greater the proportion of salary to expected income.[7]
2. As the salesperson's work effectiveness increases, the greater the firm's expected profits and the harder the salesforce works.
3. As the expected utility from the alternative job increases, the less the salesforce works, the more its expected income, the less the firm's expected profits, and the greater the proportion of salary to expected income.

[6]Most of these comparative-statics results are replicated for a binomial distribution as well (Basu et al. 1985).

[7]The variance of gross profits for the gamma distribution is given by $g^2(t)/q$, where $g(t)$ is some increasing function of the salesperson's effort, t, and $q > 0$ is a parameter of the gamma distribution. Thus, the smaller the q, the more responsive is the variance of gross profits to the salesperson's effort.

Interpreting the Results

The four submodels can be seen as a 2 × 2 full-factorial experimental design with two factors and two levels of each factor (Figure 2). Comparing the implications of model 1 versus model 2 and model 3 versus model 4, we see that the salesperson's risk preference—whether he or she is risk neutral or not—has a "main effect" on the optimal compensation plan. With risk neutrality, salaries are not needed; with risk aversion, salaries are needed. Similarly, comparing model 1 with model 3 and model 2 with model 4, we see that the observability of the salesperson's effort has a main effect on the optimal compensation plan. If the salesperson's work is observable, commissions are not needed; otherwise they are. There are also interaction effects. For example, for the dependent variable "how hard the salesperson works," there is an interaction effect between risk aversion and observability: lack of observability results in less work if the salesperson is risk averse, but with risk neutrality the observability has no effect on how hard the salesperson works.

What about the comparative-statics results from model 4? They, too, seem to indicate something about how certain independent variables affect certain dependent variables. All of them are *ceteris paribus* results—for example, when considering the effect of uncertainty, we fix the salesperson's utility from the alternative job and his or her work effectiveness—and all of them hold regardless of the values at which we fix the other independent variables. So, even though all of this analysis is being conducted within model 4, it is as though several models with model 4's de-

FIGURE 2
Experimental Design for Theoretical Modeling of Salesforce Compensation[a]

Observability of salesperson effort	Salesperson's attitude toward risk: RISK-NEUTRAL	RISK-AVERSE
OBSERVABLE	MODEL 1 All salary, all commissions, or any mixture; penalty for shirking; salesperson puts out desired effort	MODEL 2 All salary; penalty for shirking; salesperson puts our desired effort
UNOBSERVABLE	MODEL 3 All commissions; salesperson puts out desired effort	MODEL 4 Specific mixture of salary and commissions; salesperson shirks

[a]Entries in the cells give the nature of the optimal compensation scheme and the salesperson's effort level under various treatments.

fining characteristics (salesperson's work not observed, salesperson risk averse) are being analyzed for their main effects with respect to certain independent variables. For example, comparative-statics result 2 says that the salesperson's work effectiveness has a main effect on the firm's profits and how hard the salesperson works, and we could have discovered this result—albeit approximately and much more laboriously—by analyzing a "large" number of model 4's, each with a different level of salesperson work effectiveness. Hence, comparative-statics analysis is essentially an *efficient* way to run an experiment when the "causes" being manipulated are continuous variables.[8] Model 4 is the supermodel now and the (sub)models are the ones defined by various combinations of levels of the independent variables on which the comparative statics is run.

Supermodels and Models

As is apparent from the example just considered, a supermodel is a *framework* for interpreting the implications of models. Without a supermodel it would be difficult to compare one model with another. Nevertheless, the choice of a supermodel is not easy. One issue is the tradeoff between generality and clarity. For example, in physics, the general field theory that is still being developed can be thought of as a supermodel comprising the following models: electromagnetic theory, quantum mechanics, and general theory of relativity. Each of these models, however, can be thought of as a supermodel in its own right. For example, the general theory of relativity is a supermodel for the special theory of relativity. (Similarly, model 4 is a supermodel for its comparative-statics results.) These successive attempts at generalization have as their goal the explanation of *more* phenomena within a common framework; the general theory has greater external validity (discussed subsequently) and is easier to test. However, for understanding the specific effects captured in the less general theory, *that* theory is better.

Similar tradeoffs are faced by behavioral experimenters. They must decide how many effects to "throw in" to a given experiment. For example, Rao and Monroe (1988) examined the relationships among product familiarity, objective quality, price, and perceived quality in a product class in which there is a *strong* market correlation between price and objective quality (women's blazers). They used three levels of familiarity as a covariate, four price levels, and two objective quality levels. In other words, they performed a fairly complex experiment with at least three effects. Nevertheless, by constructing the experiment as they did, they were unable to determine whether the relationships they found would apply in a product class in which there is a *weak* correlation between market price and objective quality. The point is, however, that Rao and Monroe had to limit their experiment somewhere. Their experiment is already much more complex than previous studies of the relationship between price and perceived quality (Olson 1977).

In the salesforce compensation context, our supermodel assumes that a salesperson's productivity is independent of other salespeople's productivity. Though in some situations this assumption is empirically true, in many others it is false. It is often false because one or the other of the following conditions holds: (1) salespeople work as a team in selling to an account or (2) even though different salespeople work independently, their productivity is affected by the same underlying environmental factors (e.g., state of the economy). Such dependence among salespeople is the key to explaining why salespeople in a team are all compensated alike and why sales contests are used (as we learned from Holmstrom 1982 and Nalebuff and Stiglitz 1983). By restricting our supermodel as we did, we were unable to explain *these* salesforce compensation phenomena. However, the restrictions helped us isolate and understand the effects of observability of salesperson's actions and risk aversion on the choice between salary and commissions.

The other complication in specifying a supermodel is that the supermodel is constantly changing. It evolves as our understanding evolves. Each successive study is based on a "big picture" (as it exists then), but it also contributes to the big picture. In the physics context, the development of the special theory of relativity made possible the general theory of relativity, and, in turn, the general field theory. In the Rao and Monroe (1988) study, the inclusion of familiarity as a covariate is testimony to the evolution of our understanding of consumer behavior; familiarity does not appear as a construct in the studies reviewed by Olson (1977). In the salesforce compensation context, recent research by Holmstrom and Milgrom (1990) suggests that some salespeople may be compensated by salary alone—even though their effort cannot be observed—because output is multidimensional and some of the dimensions cannot be measured (e.g., missionary work). If compensation were based solely on the observable output dimensions, salespeople may misallocate their effort with respect to the unobservable

[8]"Essentially" because comparative statics usually requires strong differentiability and convexity assumptions (Milgrom and Shannon 1991). For example, if we assume that $V(I) = -e^{\gamma I}$ with γ as the risk aversion parameter, a comparative-statics analysis of model 4 with respect to γ does not give us the expected result for risk neutrality. The optimal compensation plan turns out to be $(1/\gamma)\log[\gamma(A + B\pi)]$, where A and B are constants independent of π, and this expression is not well defined for γ equal to zero. Constructing model 3 becomes imperative if we want to see what happens with risk neutrality.

output dimensions. In effect, this research identifies another "cause" for salaries, but because this learning is new, we do not see an awareness of this issue in previous supermodels.[9] Supermodels from now on, however, must explicitly assume unidimensional or multidimensional output, depending on which cause is the focus of study.

In summary, a supermodel defines a manageable experiment, with the word "manageable" left deliberately vague. It builds in a set of potential "causes" for the phenomenon in question and the submodel analyses then help identify the implications of *those* causes. Which potential causes to include and which to leave out depends on our knowledge of potential causes—knowledge that changes as our understanding improves—and the tradeoff between generality and clarity. How this tradeoff is resolved depends on one's purpose. If understanding of specific effects is the goal, a less general supermodel focusing on *those* effects is indicated; if external validity and empirical testing are the relevant goals, a more general supermodel is appropriate.

Internal Validity of Theoretical Models

The internal validity of the salesforce compensation experiment just described is very high. Because the models were chosen carefully to form a factorial design of the two forces underlying the explanation and the conclusions were derived logically (as shown by the fact that the results have stood the test of time), there is essentially no question about the cause-effect relationships established. The qualifier "essentially" is used because there is a potential problem in model 4 with the use of environmental assumptions that are more restrictive than those in models 1, 2, and 3. For some results we assumed that the distribution of π satisfied the monotone likelihood ratio property and was convex, and for the comparative-statics results we assumed in addition that this distribution was gamma and the salesperson's utility function was a power function. One might legitimately ask: Are our results due to the model differences that we emphasized, or are they due to these additional, more restrictive assumptions that we "sneaked in" to the analysis? Because this question is really about the *generality* of the results, however, it is better to address it as an external validity issue (discussed subsequently) rather than as an internal validity issue (Cook and Campbell 1979). To see this, observe that we can get rid of any questions about internal validity by the following trick: any special assumptions made under model 4 could have been made up front, under the supermodel, as assumption 5.1! That would have given us a less general supermodel, but internal validity would be secure.

Realism of Theoretical Models

The need for internal validity in theoretical modeling necessarily implies that theoretical models will be unrealistic to some extent. The reason is twofold. First, it is easier to infer cause-effect relationships when other distracting forces—other "causes"—that could affect the phenomenon in question are not present. This fact creates a demand for spareness in the modeling, and hence unrealism. For example, in the agency theory research cited, Basu et al. (1985) omit certain real-world considerations (such as multidimensional output or the firm not knowing how productive the salesperson is) because they want to focus on two issues: how the risk-aversion characteristics of salespeople and the observability of their effort affect the compensation scheme. Inclusion of the other features would have reduced the internal validity of the research by making it more difficult to judge "what causes what." This problem is like the internal validity problems a behavioral researcher would face if asked to incorporate "field conditions" in his or her (theory-testing) laboratory experiment.

Theoretical models must be unrealistic also because variation in models is a necessary aspect of deducing cause-effect relationships (cf. Figure 1). Hence, if model A is realistic for a given situation, then model B—forced to differ from A in order to establish causality—cannot be. (Similarly, if different researchers analyze different models, at least some researchers must be analyzing unrealistic models.) Model B could be realistic for a different situation, but it need not be. For example, of the four models analyzed here, model 1 is the most unrealistic—it is difficult to find situations in which the salesperson's effort is observable and he or she is risk neutral. Model 4 is the most realistic because salespeople tend to be risk averse and their effort is generally not observable. The point is that we could not just analyze model 4; we also had to analyze the comparatively unrealistic models 1, 2, and 3.

This situation is similar to that in behavioral experiments in which not all groups of subjects can be assigned the most realistic levels of the factors. For example, in a behavioral experiment on the effects of repetition in advertising, some groups may be exposed to an advertisement zero times, others may be exposed once, still others two times, and so on. Clearly, not all levels of repetition are realistic in the same setting.

This built-in artificiality of theoretical models contrasts sharply with the quest for realism in decision

[9]In fact, in the first version of this article, there was no discussion of this issue.

support models. Because decision support models are meant to serve as operational models, they tend to be *inclusive* in their choice of variables and the variables are set at their most realistic levels (Little 1975). Theoretical models, however, tend to exclude variables that are not part of the explanation being proposed.

External Validity of Theoretical Models

The appropriate interpretation of external validity in theoretical experiments is whether the cause-effect relationships obtained in one setting—the supermodel setting under which the submodels are defined—will generalize to other settings (other supermodels). This question must be assessed effect by effect—dependent variable by dependent variable—because one cause-effect relationship may have high external validity but another may not. Which of our salesforce compensation results hold even if assumptions 1 through 6 do not hold? As Cook and Campbell (1979) have pointed out, this question boils down to whether there are *interaction* effects between the specific setting chosen and the cause-effect relationship obtained. Think of a mythical larger experiment in which assumptions 1 through 6 are also a factor, and then ask whether our results will show up as main or interaction effects in this experiment. If the former, we have externally valid (also called "robust") results; if the latter, then our results are at least somewhat externally invalid.

In the present context, we can say categorically that some of our results will not be robust to changes in assumptions 1 through 6. For example, it is possible to choose a distribution for π that does not satisfy the monotone likelihood ratio property and prove in model 4 that a salesperson's total compensation *should decrease* with sales for some range of sales (Grossman and Hart 1983). In contrast, the result that a salesperson's expected income should rise with the expected utility of his or her alternative job has high external validity (but, of course, this is hardly a distinctive contribution of agency theory).

Requiring that our results be robust to *any* change in assumptions 1 through 6, however, is an overly harsh requirement. Ultimately, the researcher must think about what kinds of robustness are good to have and what kinds one can live without. One can live without lack of robustness over *unrealistic* variations in the assumptions. By this argument, the lack of robustness of the salesforce compensation results when the monotone likelihood ratio property is not satisfied would not be too troublesome. The monotone likelihood ratio assumption *is* a realistic assumption: higher sales ought to signal higher salesperson effort rather than lower salesperson effort.

Moreover, we may not even want to consider all realistic variations of the supermodel if we are willing to restrict the applicability of the theory. The substantive aspects of assumptions 1 through 6 serve only to delimit the *observable scope* of the theory—the subset of real-world situations being examined—and one could *define* the target area of applicability of the theory to be this subset of the real world (Cook and Campbell 1979, p. 71). For example, we could say: "We are explaining the use of salaries and commissions for situations in which (1) a salesperson's output is measurable and depends only on *his* or *her* work (and not other salespeople's work) and random factors unique to him or her, (2) the firm commits to a compensation package, etc." As long as we can find salesforce compensation situations that match these restrictions, the theory has applicability and its external validity can be assessed for these situations.

The advantage of restricting the scope of our theory to the empirically correct assumptions of the supermodel is that it enables us to focus on results that are sensitive to the unverifiable assumptions of the supermodel. Assumptions whose sole purpose is to make the analysis mathematically tractable are generally without empirical content. Therefore they are also unverifiable. For example, in assumption 4, the requirement that the utility function V be twice-continuously differentiable makes the analysis tractable. Subsequently we assumed that the function V(I) was actually $2I^{1/2}$, thus making it easier to do the comparative-statics analysis. These assumptions are difficult to verify because our measurement techniques are too coarse to pick up such details of utility functions.

How do we check whether our theory is robust to mathematical assumptions? Unfortunately, there is no easy way. If the simplifying mathematical assumptions were really made to make the analysis tractable,[10] *proving* that the results would not change without them is difficult. (The proof would consist of analyzing a model without the simplifying assumptions, and that analysis would be intractable.) The only "solution" is to try to replicate the results, logically or by simulation, with several versions of these assumptions. (Replication has a similar role in assessing the external validity of behavioral experiments.) The more replication attempts a result survives, the more robust it becomes. In the case of our salesforce compensation theory, the Holmstrom finding that informative signals are valuable is robust, but the comparative-statics results on the shape of the compensation function are not (see Basu and Gurumurthy 1989; Hart and Holmstrom 1987).

Some supermodel assumptions that make the anal-

[10]This is where the mathematical ability of the modeler comes in. The stronger the modeler's mathematical ability, the fewer tractability assumptions he or she needs to make.

ysis mathematically easier also have empirical content. If this content is empirically verifiable, we treat such assumptions exactly the same way as any verifiable assumption. Either we limit the applicability of our theory to situations in which they are true or else evaluate the robustness of our results to realistic variations of them. If the empirical content of the assumption is not verifiable, it is essentially a mathematical assumption and we need to assess the robustness of our results over variations of it. For example, in our salesforce compensation theory, though the gamma distribution for the distribution of sales enables us to get closed-form comparative-statics results, it also implies that the *variance* of sales increases with the salesperson's effort. One could argue that reality is just the opposite. In this example, then, the mathematical assumption has empirical content, but the empirical content may not be true. Hence the results that *depend* on the gamma distribution may have no real-world applicability.

To summarize, the external validity of a model-based theory has two aspects: (1) applicability (can we find any real-world situations that fit the verifiable assumptions of the supermodel?) and (2) robustness of the theory with respect to its unverifiable assumptions.

Usefulness of Theoretical Modeling

Though questioning the external validity of the salesforce compensation theory just discussed, one could nevertheless argue that it has served its main purpose. *It has given us one explanation of the observed phenomenon*. The interaction effects uncovered while discovering the lack of external validity indicate simply that there are other explanations of the observed phenomenon;[11] they do not preclude the explanation proposed. The *admissibility* of the proposed explanation has been upheld on internal validity grounds. The main purpose of theoretical modeling is *pedagogy*—teaching us how the real world works. That purpose is always served by internally valid theoretical experiments. (This is probably the basis for the folklore that theories are rarely rejected by data, only by other theories.) Theoretical modeling is a way to think clearly, and that is always valuable.

The pedagogical use of theoretical modeling—and the preeminence of internal validity considerations in that use—is analogous to the theory-testing purpose of empirical experiments. When theory testing, as opposed to application, is the purpose of the empirical experiment, internal validity considerations dominate external validity considerations (Calder, Phillips, and Tybout 1981, 1982; Cook and Campbell 1979, p. 83; Lynch 1982).

Does theoretical modeling have any practical use for managers? Yes, as long as the theory's observable scope assumptions cover the manager's situation and one focuses on the robust results. (This may require a more general supermodel than the one that is optimal for learning.) The usefulness comes in two ways: (1) as direct *qualitative guidance* for managerial policy and (2) as the *basis for a decision support system* that will "fine-tune" the theory to the manager's particular decision-making environment and generate quantitative prescriptions.

How does theoretical modeling provide qualitative guidance for managerial policy? Theoretical modeling helps managers learn about the forces that determine the "bottom-line" effect of managerial decision variables. Such knowledge is crucial in deciding how to set those decision variables in a given situation and how to change them if the decision-making environment changes. For example, the salesforce compensation models just discussed teach managers why it is generally optimal to have salaries and commissions in their salesforce compensation plans. The "why" part of this learning is the crucial learning. After all, the use of salaries and commissions is standard industry practice.[12] (Otherwise we would not have developed a theory to explain them.) Managers could develop salary-commission plans for their firms by simply copying other firms' policies. Knowing the theory, however, they can do better. They know that the optimal compensation policy depends on the risk aversion of the salesperson and whether he or she can be monitored directly, and that it also depends on the uncertainty in deducing the salesperson's effort from sales when direct monitoring is too costly. So in deciding how their compensation policy should differ from average industry practice, *they know what differences to look for*. Moreover, once they know the nature of these differences, they can determine how their plans should differ from average industry practice. Similarly, they can determine how they should change their company's compensation policy if the selling environment for their firm changes in the future.

Qualitative guidance, however, is only one part of the benefit from theoretical models. Managers can ask that decision support systems be built to apply the theory more precisely to their situation (Little 1970, 1979). The starting point is to decide *which* theories to apply. This decision must be based on an assessment of the match between the scope of various theories and the manager's particular situation. Different theories em-

[11] The interaction effect uncovered is itself new learning about the phenomenon and that is another benefit of theoretical modeling.

[12] In John and Weitz's (1989) study, 76% of the sample used salary-plus-commissions plans.

phasize different forces. The manager must decide which forces are key in his or her *particular* situation and select the appropriate theories. This requires judgment. Multiple theories may be necessary. For example, in the salesperson compensation context, there are at least three theories that differ in scope: (1) the theory just outlined, (2) Nalebuff and Stiglitz's (1983) theory explaining the use of sales contests in environments where several salespeople share some common selling factors (e.g., the same territory or the same product in different territories), and (3) the theory explaining the use of *menus* of compensation plans in environments where the manager does not know the salesperson's ability (Lal and Staelin 1986; Rao 1990). Which of these theories to use depends on answers to such questions as: Is the salesperson selling a unique product in a unique territory? How much is known about the salesperson's productivity on this job and other related jobs?

Once the appropriate theories have been picked, the decision support system must use the teachings of those theories in setting up its "measurement module"—what variables to measure and how to measure them. Then, algorithms must be devised for carrying out the optimizations of interest to the manager. See Dobson and Kalish (1988) for a concrete illustration of this process.

Empirical Testing of Model-Based Theories

The pedagogical purpose of theoretical modeling is served even if the theory has not been tested. The problem with leaving a theory untested, however, is that then the theory has limited empirical content and therefore cannot be used to provide guidance to the manager. Moreover, if there are multiple theories, they are also likely to have multiple prescriptions for optimal policy, so, again, what should we ask the manager to do?

Theories are tested by their predictions, not by the realism of their supermodel assumptions *per se* (Friedman 1953). If a theory's supermodel assumptions are unrealistic, the theory is not applicable, and hence also untestable. A more general theory with weaker supermodel assumptions is called for. Empirical studies that test only the realism of assumptions are evaluating only the applicability of the theory (see, e.g., John and Weitz 1989). What is missing is the critical next step: testing whether the theory makes correct predictions in its area of applicability.

A prediction is any result of the theory. For example, one prediction of the theory just considered is: "As output becomes a poorer indicator of the salesperson's effort, the salary component of compensation increases, *ceteris paribus*." Predictions may be observationally testable or experimentally testable. An *observationally testable* prediction of a theory is a prediction $X \rightarrow Y$ that holds under a set of verifiable assumptions A and *any* unverifiable assumptions B. The researcher *finds* a setting to conform to A and then *observes* whether X leads to Y. In contrast, if the theory were such that under A and B, $X \rightarrow Y$, but under A and B′, $X \rightarrow Y'$, he or she would not have an observationally testable proposition. Even if the researcher knew that the testing situation satisfied A, he or she would not know whether B or B′ was true, and the theory makes different predictions in each case. However, even this situation may be testable experimentally. The researcher chooses his or her setting to conform to A and *creates* B or B′, and then tests whether X implies Y or Y′. (For examples of this construction see Smith 1982, 1986.) It is precisely because of this extra "man-made" manipulation in an experiment that "demand effects" sometimes result. That is why observational tests, if possible, are preferred to experimental tests.

The procedure just outlined is changed only slightly if one is testing among alternative theories. One starts by making sure the competing theories are in fact competing theories—that is, they are meant to explain the same phenomenon—and that they have some common scope—that is, environments can be found (or created) in which all of the theories can operate. Then one develops lists of *distinctive* predictions of the competing theories and checks which of them hold. The qualification "distinctive" is necessary because often theories share some predictions. For example, one could hardly claim that the prediction "the expected income of the salesperson rises with the utility he or she expects from his or her other opportunities" is a distinctive prediction of agency theory. Any self-respecting theory of salesforce compensation would make the same prediction.

It is in developing observationally testable predictions that lack of external validity really hurts. It makes it difficult to *find* real-world situations that fit the observable scope of the theory and in which the theory makes predictions that are robust to its unverifiable assumptions. For example, we may have a hard time finding a salesforce compensation situation in which a salesperson's output is independent of other salespeople's outputs. What do we do in such situations? There are two possibilities. First, we could follow the Friedman dictum of ignoring the supermodel assumptions and simply see whether the $X \rightarrow Y$ predictions hold. If the predictions hold under a wide variety of situations, we consider the theory corroborated even though it is possible that our theory is false and we just happened to find situations B′ in which *another* theory predicts $X \rightarrow Y$. Second, we could develop a more general, more realistic, supermodel and develop predictions within it. Such a supermodel typically will carry *several* forces simultaneously, making it more difficult to *understand* specific effects, but it may provide a way to assess how the cause-effect relationships of interest are affected by other forces, making the theory testable.

In the salesforce compensation theory just discussed, the observationally testable predictions arise from Holmstrom's (1979) finding that the optimal compensation contract must use all informative signals of the salesperson's effort, and only those (see Antle and Smith 1986; Eisenhardt 1985; Rosen 1990).

Theoretical Modeling Versus Behavioral Theories

Thus far we have drawn a close analogy between the thought experiments that underlie theoretical modeling and the empirical experiments that test behavioral theories in marketing. What about the theory-building process in behavioral marketing? How does *that* compare with model-based theorizing?

One obvious difference is that behavioral theories are largely *verbal* whereas theoretical modeling is mathematical.[13] This difference produces two effects. First, because the language of verbal reasoning is necessarily less precise than mathematics, the verbal theorist has a greater chance of going wrong in his or her reasoning. This does not mean that wrong conclusions are drawn in every instance of verbal reasoning or that correct conclusions are drawn in every instance of theoretical modeling, only that the probability of mistakes is higher with verbal reasoning. Verbal arguments are also more difficult to check than mathematical arguments. Different researchers looking at the same verbal theory may disagree on what the theory is saying because they interpret the terms differently. Again, the chances of this kind of confusion are less in theoretical modeling because the assumptions, definitions, and arguments are all stated mathematically.

The other, subtler, difference between behavioral theories and model-based theories is in their use of the researcher's intuition in the theory-building process. In both cases, the researcher's intuition plays a role. A behavioral theory is essentially an amalgamation of previous empirical findings, other theories, and the researcher's intuition. Similarly, theoretical modelers draw on their empirical knowledge, other theories, and intuition in formulating a model and "looking" for certain results. The difference, however, is that whereas the possibility exists that the theoretical modeler will be surprised in the theory-building process,[14] that possibility is less likely in verbal theorizing. In the process of proving the results that he or she conjectured on the basis of intuition, the theoretical modeler may discover something he or she did not expect. This discovery will enhance his or her intuition. Though behavioral researchers may also discover something they were not looking for *when they do their empirical experiment*, such discoveries are less likely at the theory-development stage. The difference comes from the fact that mathematical reasoning is a much more *searching* process than verbal reasoning. The rules of mathematical argument *require* that *all* feasible paths be explored—and they may include paths that are easy to overlook when one is thinking verbally.

The greater precision and the attendant promise of "deeper" theories afforded by theoretical modeling also impose costs on such theorizing. Some models are simply not tractable for mathematical reasoning. The researcher is unable to prove his or her results. To achieve tractability, the researcher may simplify his or her model, but then the generality of results becomes a question. Without generality, the theory is less applicable and less testable.

Another cost of theoretical modeling is that some phenomena do not lend themselves well to mathematical modeling. Prime examples are the "framing effects" uncovered by psychologists (Slovic, Fischoff, and Lichtenstein 1977; Tversky and Kahneman 1979) and "procedural irrationality" (Simon 1978). The former refers to the phenomenon of people solving problems differently depending on how the problems are set up. Theoretical modelers have difficulty representing such "set-up differences" mathematically. Procedural irrationality refers to the fact that people may not optimize. Much of the logical tightness in theoretical modeling comes from the assumption of optimization. In some cases, apparently irrational behavior is rationalizable. For example, the finding that consumers do not gather much information before making brand choice decisions (Wilkie and Dickson 1985) can be explained as a rational, utility-maximizing response to the costs of information acquisition (Stigler 1961; Tirole 1989). In other cases, the assumption of optimization is much more difficult to rationalize. For example, Akerlof (1991) notes how the assumption of dynamic optimization fails when people procrastinate because the current costs of act-

[13]Verbal theories are by no means unique to behavioral marketing. Darwin's theory of evolution is an early example of a verbal theory. In economics, the most well-known verbal theory is Williamson's (1975) transaction costs theory.

[14]This point pertains to model builders themselves being surprised. Of course, the likelihood of readers of the theory being surprised is higher. If even the reader can anticipate the results of the theory by just looking at the model, the modeling effort has not been very useful. It has lent precision to—and verified the correctness of—an argument that was intuitive to begin with. Such modeling exercises have limited pedagogical value for the reader.

ing are more salient than the current benefits. In these cases, procedural rationality should be treated as a supermodel assumption (as we did in assumptions 3 and 4), and we should view the theory as applicable only for situations in which procedural rationality can be expected to hold. This is really why in laboratory tests of economic theories, researchers are advised to provide their human subjects with sufficient incentives to optimize (Smith 1986). In observational tests, the "disciplinary forces of real-world markets" serve the same purpose.

One could still make the case that theoretical modeling with procedural rationality assumptions is a useful thought experiment. The argument would be that, to appreciate the effects of irrationality, we need to know the effects of rationality. However, this also means that we need to develop alternative theories that do *not* demand such extreme rationality and to examine empirically how their predictions compare with the rationality-based theories. For an example of such a comparison, see Shiller (1990).

Theoretical Models Versus Decision Support Models

Unlike theoretical modeling and behavioral research in marketing, decision support models are designed to help managers make decisions in their operating environment. If the first two can be thought of as developing the science of marketing, the latter is engineering. Little (1979, p. 11) describes a marketing decision support system as "a coordinated collection of data, systems, tools, and techniques with supporting software and hardware by which an organization gathers and interprets relevant information from business and environment and turns it into a basis for marketing action." Much of the quantitative model building in marketing is decision support modeling. In some instances of such modeling, all of its components are developed—measurement model setup, estimation, profit computation, and optimal policy determination—but more typically only the measurement model is estimated (e.g., Guadagni and Little 1983). Other examples of decision support systems include Lodish's (1971) CALLPLAN model for scheduling a salesforce, Little's (1975) BRANDAID system for brand management, Silk and Urban's (1978) ASSESSOR system for assessing the sales potential of new, frequently purchased consumer goods, Dobson and Kalish's (1988) model for designing a product line, and Bultez and Naert's (1988) SHARP model for allocating retail shelf space among various products.

How does decision support modeling compare with theoretical modeling? The fundamental difference is in the objectives. The practical objectives of decision support modeling translate to a preference for realistic representations of the manager's decision situation. In contrast to theoretical modeling, there is no need to create unrealistic models because cause-effect inference is not the goal. The goal of a decision support model is to capture mathematically the essentials of the manager's decision-making situation, so that the model can then be manipulated to derive prescriptions for managerial action.[15] Therefore, unlike a theoretical modeler, who is trying to create a "spare" environment by *excluding* variables, a decision support modeler is trying to capture as much of reality as possible by *including* variables. For example, Little and Lodish's (1969) MEDIAC model has variables representing the effectiveness of different media in reaching various target segments, the sizes and sales potentials of various market segments, seasonality effects, and so on. Lodish's (1971) extension of MEDIAC adds to this list competitors' media schedules and associated parameters. Aaker's (1975) ADMOD system is a further extension, *simultaneously* addressing budget, copy, and media allocation decisions.

Is there any connection between theoretical modeling and decision support modeling beyond their use of mathematics? Just as mechanical engineering builds on physics, decision support systems build on the conceptual framework and cause-effect relationships provided by empirically tested theories. For example, advertising decision systems such as MEDIAC incorporate the advertising carryover effects emphasized in the theoretical work of Nerlove and Arrow (1962) and empirically tested by Lambin (1976). Similarly, Dobson and Kalish (1988) develop a decision support system for product line design, using the self-selection framework in Moorthy (1984).

This process of adapting theories for decision support will face its severest test in the new theories being developed to account for *strategic* behavior. Such behavior involves the interaction of two or more self-interested parties, quite unlike the "single-person" decision problems characteristic of most marketing decision support systems. The single-person problem is amenable to general "operations-research"-based methods that transcend specific situations. Such an approach can hardly work for strategic situations. As indicated previously, even for the relatively simple salesforce compensation problem, only situation-specific answers are possible. Moreover, the dimensionality of the "situations space" is very large. There is no alternative to studying the particular situation carefully—on a case-by-case basis—before offering recommendations.

[15]Both decision support systems and theoretical models may solve for the manager's optimal strategy. However, whereas the theoretical modeler sees the optimal strategy as the phenomenon being explained by his or her model, the decision support modeler sees it as the solution to the manager's problem.

TABLE 1
Theoretical Modeling of Competitive Product Strategy[a]

Author(s)	Type of Product	Consumer Choice Deterministic or Stochastic?	Price Competition?	Cost Differences Among Products?	Result: Product Differentiation?
Hotelling (1929)	Heterogeneous ideal points	Deterministic	No	No	No
d'Aspremont, Gabszewicz, and Thisse (1979)	Heterogeneous ideal points	Deterministic	Yes	No	Yes
Hauser (1988)[b]	Heterogeneous ideal points	Deterministic	Yes	No	Yes
de Palma et al. (1985)	Heterogeneous ideal points	Stochastic	Yes	No	No
Shaked and Sutton (1982)	Homogeneous ideal points	Deterministic	Yes	No	Yes
Moorthy (1988)	Homogeneous ideal points	Deterministic	Yes	Yes	Yes

[a]This is a selected listing of models centering on a few key dimensions of competitive product strategy. All of these models share the following supermodel assumptions: (1) product competition is on a single attribute, (2) consumer segments are uniformly distributed over the relevant space, and (3) the number of competing firms is two (though some results are more general).
[b]Even though the Hotelling, d'Aspremont et al., and Hauser models have the same type of product and consumer model in the classification system used here, they differ in some details. In particular, consumers' utility functions are different in the three models, and Hauser, furthermore, begins with two product attributes which then collapse into one because of exogenous restrictions.

Conclusion

Real-world marketing situations are incredibly complex. Many forces operate and what we observe as managerial actions is the aggregate effect of all these forces. Theoretical modeling is a way to learn the specific effect of each force. As described here, it works by a process of experimentation. The analyst constructs a series of models, each capturing a different subset of the real world, and determines by logical argument what the managerial actions would be in each of these artificial worlds. Then, by relating the managerial implications to the model design, he or she infers how various forces affect managerial actions.

The running example used here is the salesforce compensation theory developed by Basu et al. (1985). Other examples could have been chosen. Table 1 shows the development of the competitive theory of product differentiation over a period of nearly 60 years. Five models are described on four dimensions: type of product, whether consumer behavior is deterministic or stochastic, whether price competition is allowed or not, and whether there are cost differences among the various products. By studying the experimental design underlying these models, one can infer the following cause-effect relationships:

- Two forces determine the competitive product strategy of a firm.
- One force is the desire of *each* firm to choose a product that best reconciles consumer preferences and costs. This force brings the firms' products together.
- The other force is price competition, which pushes them apart.

To illustrate, in the study by de Palma et al (1985), only the first force is effective because price competition is weakened by the stochastic nature of consumer behavior; hence their no-differentiation result. Finally, observe that the type of product manipulation, and the different utility functions used in the various models, do not have any causal effects; they only increase the external validity of the theory (but see Economides 1986).

Theoretical modeling is both an art and a science. The scientific part is the use of logical arguments and the affinity to experimental design. The artistic part is the choice of the model itself. The modeler must strike a careful balance between realism and the need to isolate the interesting forces. On the one hand, there is the danger of including so many effects in the model that cause-effect relationships are impossible to infer. On the other hand, a theoretical model can be so spare that the results are obvious. Perhaps worst of all, a model can be bad because it focuses on "uninteresting" forces. (That, however, is a matter of taste.)

REFERENCES

Aaker, David (1975), "ADMOD: An Advertising Decision Model," *Journal of Marketing*, 12 (February), 37–45.

Akerlof, George (1991), "Procrastination and Obedience," *American Economic Review*, Papers and Proceedings, 81 (May), 1–19.

Antle, Richard and Abbie Smith (1986), "An Empirical Investigation of the Relative Performance Evaluation of Corporate Executives," *Journal of Accounting Research*, 24 (Spring), 1–39.

Assmus, G., J. Farley, and D. Lehmann (1984), "How Advertising Affects Sales: Meta-Analysis of Econometric Results," *Journal of Marketing Research*, 21 (February), 65–74.

Basu, Amiya and Kalyanaraman Gurumurthy (1989), "A Study of a Class of Simple Salesforce Compensation Plans," BEBR Working Paper No. 89-1565, University of Illinois at Urbana-Champaign.

————, Rajiv Lal, V. Srinivasan, and Richard Staelin (1985), "Salesforce Compensation Plans: An Agency Theoretic Perspective," *Marketing Science*, 4 (Fall), 267–91.

Bettman, James (1979), *An Information Processing Theory of Consumer Choice*. Reading, MA: Addison-Wesley Publishing Company.

Borch, Karl (1962), "Equilibrium in a Reinsurance Market," *Econometrica*, 30 (July), 424–44.

Bultez, Alain and Philippe Naert (1988), "SH.A.R.P.: Shelf Allocation for Retailers' Profit," *Marketing Science*, 7 (Summer), 211–31.

Calder, Bobby, Lynn Phillips, and Alice Tybout (1981), "Designing Research for Application," *Journal of Consumer Research*, 8 (September), 197–207.

————, ————, and ———— (1982), "The Concept of External Validity," *Journal of Consumer Research*, 9 (December), 240–4.

Cook, Thomas and Donald Campbell (1979), *Quasi-Experimentation*. Boston: Houghton-Mifflin Company.

d'Aspremont, C., J. Jaskold Gabszewicz, and J.-F. Thisse (1979), "On Hotelling's 'Stability in Competition'," *Econometrica*, 47 (September), 1145–9.

de Palma, A., V. Ginsburgh, Y. Y. Papageorgiou, and J.-F. Thisse (1985), "The Principle of Minimum Differentiation Holds Under Sufficient Heterogeneity," *Econometrica*, 53 (July), 767–81.

Dobson, Gregory and Shlomo Kalish (1988), "Positioning and Pricing a Product Line," *Marketing Science*, 5 (Spring), 107–25.

Economides, N. (1986), "Minimal and Maximal Product Differentiation in Hotelling's Duopoly," *Economics Letters*, 21, 67–71.

Eisenhardt, Kathleen M. (1985), "Control: Organizational and Economic Approaches," *Management Science*, 31 (February), 134–49.

Friedman, Milton (1953), "The Methodology of Positive Economics," in *Essays in Positive Economics* by M. Friedman. Chicago: The University of Chicago Press.

Grossman, Sanford and Oliver Hart (1983), "An Analysis of the Principal-Agent Problem," *Econometrica*, 51 (January), 7–45.

Guadagni, Peter M. and John D. C. Little (1983), "A Logit Model of Brand Choice Calibrated on Scanner Data," *Marketing Science*, 2 (Summer), 203–38.

Hanssens, Dominique, Leonard Parsons, and Randall Shultz (1990), *Market Response Models: Econometric and Time Series Analysis*. Boston, MA: Kluwer Academic Publishers.

Harris, Milton and Artur Raviv (1979), "Optimal Incentive Contracts With Imperfect Information," *Journal of Economic Theory*, 20 (April), 231–59.

Hart, Oliver and Bengt Holmstrom (1987), "The Theory of Contracts," in *Advances in Economic Theory, Fifth World Congress*, T. Bewley, ed. New York: Cambridge University Press.

Hauser, John (1988), "Competitive Price and Positioning Strategies," *Marketing Science*, 7 (Winter), 76–91.

———— and Birger Wernerfelt (1989), "The Competitive Implications of Relevant-Set/Response Analysis," *Journal of Marketing Research*, 26 (November), 391–405.

Hess, James and Eitan Gerstner (1987), "Loss Leader Pricing and Rain Check Policy," *Marketing Science*, 6 (Fall), 358–74.

Holmstrom, Bengt (1979), "Moral Hazard and Observability," *Bell Journal of Economics*, 10 (Spring), 74–91.

———— (1982), "Moral Hazard in Teams," *Bell Journal of Economics*, 13 (Autumn), 324–40.

———— and Paul Milgrom (1987), "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55 (March), 303–28.

———— and ———— (1990), "Multi-Task Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," Technical Report No. 6, Stanford Institute for Theoretical Economics, Stanford, CA.

Hotelling, Harold (1929), "Stability in Competition," *Economic Journal*, 39 (March), 41–57.

Hunt, Shelby (1991), *Modern Marketing Theory*. Cincinnati, OH: South-Western Publishing Co.

John, George and Barton Weitz (1989), "Salesforce Compensation: An Empirical Investigation of Factors Related to Use of Salary Versus Incentive Compensation," *Journal of Marketing Research*, 26 (February), 1–14.

Lal, Rajiv (1990), "Manufacturer Trade Deals and Retail Price Promotions," *Journal of Marketing Research*, 27 (November), 428–44.

———— and Richard Staelin (1986), "Salesforce Compensation Plans in Environments With Asymmetric Information," *Marketing Science*, 5 (Summer), 179–98.

Lambin, Jean-Jacques (1976), *Advertising, Competition, and Market Conduct in Oligopoly Over Time*. Amsterdam: North-Holland Publishing Company.

Little, John D. C. (1970), "Models and Managers: The Concept of a Decision Calculus," *Management Science*, 16 (April), B466–85.

———— (1975), "BRANDAID: A Marketing Mix Model," *Operations Research*, 23 (July-August), 628–73.

———— (1979), "Decision Support Systems for Marketing Managers," *Journal of Marketing*, 43 (Summer), 9–27.

———— and Leonard Lodish (1969), "A Media Planning Calculus," *Operations Research*, 17 (January-February), 1–34.

Lodish, Leonard (1971), "CALLPLAN: An Interactive Salesman's Call Planning System," *Management Science*, 18, Part II (December), 25–40.

Lynch, John, Jr. (1982), "On the External Validity of Experiments in Consumer Research," *Journal of Consumer Research*, 9 (December), 225–39.

Mahajan, Vijay and Eitan Muller (1986), "Advertising Pulsing Policies for Generating Awareness for New Products," *Marketing Science*, 5 (Spring), 89–106.

McGuire, Timothy and Richard Staelin (1983), "An Industry Equilibrium Analysis of Downstream Vertical Integration," *Marketing Science*, 2 (Spring), 161–92.

Milgrom, Paul and Christina Shannon (1991), "Monotone Comparative Statics," Technical Report No. 11, Stanford Institute for Theoretical Economics, Stanford, CA.

Mirrlees, Jim (1976), "The Optimal Structure of Authority and Incentives Within an Organization," *Bell Journal of Economics*, 7 (Spring), 105–31.

Moorthy, K. Sridhar (1984), "Market Segmentation, Self-Selection, and Product Line Design," *Marketing Science*, 3 (Fall), 288–305.

——— (1988), "Product and Price Competition in a Duopoly," *Marketing Science*, 7 (Spring), 141–68.

Nalebuff, Barry and Joseph Stiglitz (1983), "Prizes and Incentives: Toward a General Theory of Compensation and Competition," *Bell Journal of Economics*, 14 (Spring), 21–43.

Nerlove, Marc and Kenneth Arrow (1962), "Optimal Advertising Policy Under Dynamic Conditions," *Econometrica*, 29 (May), 129–42.

Olson, Jerry (1977), "Price as an Informational Cue: Effects on Product Evaluations," in *Consumer and Industrial Buying Behavior*, Arch G. Woodside et al., eds. New York: North-Holland Publishing Company.

Puto, Christopher (1987), "The Framing of Buying Decisions," *Journal of Consumer Research*, 14 (December), 301–15.

Rao, Ram (1990), "Compensating Heterogeneous Salesforces: Some Explicit Results," *Marketing Science*, 9 (Fall), 319–41.

Rosen, Sherwin (1990), "Contracts and the Market for Executives," Working Paper No. 90-12, Department of Economics, University of Chicago.

Ross, Stephen (1973), "The Economic Theory of Agency: The Principal's Problem," *American Economic Review*, 63, 134–9.

Shaked, Avner and John Sutton (1982), "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, 49, 3–13.

Shavell, Steve (1979), "Risk Sharing and Incentives in the Principal and Agent Relationship," *Bell Journal of Economics*, 10 (Spring), 55–73.

Shiller, Robert (1990), "Market Volatility and Investor Behavior," *American Economic Review*, Papers and Proceedings, 80 (May), 58–62.

Silk, Alvin and Glen Urban (1978), "Pre-Test-Market Evaluation of New Packaged Goods: A Model and Measurement Methodology," *Journal of Marketing Research*, 15 (May), 171–91.

Simon, H. A. (1978), "On How To Decide What To Do," *Bell Journal of Economics*, 9 (Autumn), 494–507.

Slovic, Paul, B. Fischhoff, and Sarah Lichtenstein (1977), "Behavioral Decision Theory," *Annual Review of Psychology*, 28, 1–39.

Smith, Vernon (1982), "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72 (December), 923–55.

——— (1986), "Experimental Methods in the Political Economy of Exchange," *Science*, 234 (October), 167–73.

Spence, Michael and Richard Zeckhauser (1971), "Insurance, Information, and Individual Action," *American Economic Review*, Papers and Proceedings, 61, 380–7.

Stigler, George (1961), "The Economics of Information," *Journal of Political Economy*, 69 (June), 213–15.

Sujan, Mita (1985), "Consumer Knowledge: Effects on Evaluation Strategies Mediating Consumer Judgments," *Journal of Consumer Research*, 12 (June), 31–46.

Suppe, Frederick (1977), *The Structure of Scientific Theories*. Urbana: University of Illinois Press.

Tirole, Jean (1989), *The Theory of Industrial Organization*. Cambridge, MA: The MIT Press.

Tversky, Amos and Daniel Kahneman (1979), "Prospect Theory: An Analysis of Decisions Under Risk," *Econometrica*, 47 (March), 263–91.

Wilkie, William L. and Peter Dickson (1985), "Shopping for Appliances: Consumers' Strategies and Patterns of Information Search," reprinted in *Perspectives in Consumer Behavior*, Harold Kassarjian and Thomas Robertson, eds. Englewood Cliffs, NJ: Prentice-Hall, Inc.

Williamson, Oliver (1975), *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: The Free Press.

Wilson, Lynn and John Norton (1989), "Optimal Entry Timing for a Product Line Extension," *Marketing Science*, 8 (Winter), 1–17.

———, Allen Weiss, and George John (1990), "Unbundling of Industrial Systems," *Journal of Marketing Research*, 27 (May), 123–38.

Wilson, Robert (1969), "The Structure of Incentives for Decentralization Under Uncertainty," *La Decision*, 171.

Wright, Peter (1975), "Consumer Choice Strategies: Simplifying vs Optimizing," *Journal of Marketing Research*, 11 (February), 60–7.

Reprint No. JM572106