

Marketing Decision Support and Intelligent Systems: Precisely Worthwhile or Vaguely Worthless?

ERIC M. EISENSTEIN and LEONARD LODISH

Our goal in this chapter is to review the marketing decision support system (MDSS) literature so as to provide maximal guidance to researchers and practitioners on how best to improve marketing decision-making using decision support systems. In order to achieve this goal, we lay out a taxonomy of decision support systems, create an integrative framework showing the drivers that maximally aid successful implementation, and propose future research that will help to resolve the inconclusive results in the literature. Throughout the chapter, we also attempt to reunite the divided decision support system literature by examining the assumptions underlying different research traditions in a broad, integrative context.

The chapter is structured as follows. We first place decision support system research in context by characterizing the major assumptions used in the two research traditions that form the backbone of the DSS field. We also make the case that decision support can be helpful even to expert decision-makers. Second, we define the basic constructs that are antecedent to decision support system design. Third, in light of these basic constructs, we construct a taxonomy of the types of decision support systems that we will be discussing. Fourth, we propose a metric for evaluating DSS success and we present an integrative framework that reflects the factors that appear most directly to affect the success of marketing decision support systems. We also relate existing systems to the taxonomy and framework. Fifth, we review the factors in our framework to identify the most important and to provide guidance about best practices. Sixth, we answer the question: do MDSS

work? Seventh, we examine the research on knowledge-based systems in marketing, assess whether they work, and discuss the relationship between these systems and other types of decision support systems. Eighth, we briefly review recent developments in machine learning and data mining. Finally, we provide general comments and our thoughts on future directions for research. Readers who are interested in another recent review should see Wierenga et al. (1999).

DECISION SUPPORT SYSTEMS – OVERVIEW

There are two conceptually distinct traditions of decision support system research, which have origins in the disciplines of computer science and decision theory. Computer science gave rise to knowledge-based system design (also called expert systems, or artificial intelligence). Decision theory developed across multiple fields, including psychology, operations research, engineering, and a variety of business disciplines. Because of the wide differences in training of researchers, decision-theory research has been broader in scope than computer science. However, the most typical decision-theoretic approaches have concentrated on providing statistical or optimization-based models to improve decision-making. In this section, we examine the different assumptions that researchers trained in these traditions have brought to the design of decision support systems. We also justify

an assumption common to both decision support traditions, that decision support is necessary even for experienced experts.

Who needs Decision Support

Success in business is based on making better decisions than one's competitors. The volume and complexity of information that is available as an input to decisions has been increasing, particularly in marketing and other fields where computerized data collection is the norm (Huber, 1983; King, 1993; Little, 1970). The overall greater complexity of decision environments has brought about increased reliance on specialists and experts in many disciplines (Wright & Bolger, 1992), including marketing, and this increased reliance on experts motivated researchers to examine the decision-making process more rigorously.

These research efforts have been motivated by two (often opposing) goals. The first is to understand how people *actually* make decisions – a descriptive goal. The second is to understand how people should *ideally* make decisions – a normative goal. A smaller group of researchers have been motivated by a desire to improve decisions – a prescriptive goal. Prescriptive research takes as given that there is a discrepancy between normative and descriptive decision-making that leaves room for improvement.

It is intuitively clear why a novice needs a decision support system. It is less obvious why an experienced, senior decision-maker would need one. Both the computer science and the decision theory traditions assume that decisions by seasoned decision-makers can be improved (a prescriptive stand). Furthermore, both traditions agree that some experienced decision-makers (experts) are better than others. Where the groups disagree is on the normative theory, the 'gold standard,' that represents best performance in a field.

Assumptions in the Tradition of Computer Science

The computer-science tradition starts with two plausible assumptions. First, it is assumed that experienced, senior decision-makers (such as experienced marketing managers) are able to achieve excellent decision-making performance and that these decision-makers outperform novices (such as recent graduates) in their field of expertise. Second, it is assumed that the most experienced and expert individuals in a field represent the normative standard against which both other practitioners and computerized systems should be judged (Hayes-Roth et al., 1983; Turban & Aronson, 1998). These assumptions lead to the conclusion that to improve decisions, we should capture the knowledge and process of reasoning of the best experts, codifying

it in computer code. Systems that attempt to accomplish this task are termed knowledge-based (or expert) systems. Once developed, these systems can be used by less seasoned decision-makers or novices to improve their decision-making, or as replacements for the original expert(s). The assumption of outstanding expert judgment is rarely challenged in the computer-science tradition.

Assumptions in the Decision Theory Tradition

By contrast, researchers in decision theory have concluded that even the top experts in a field rarely represent the appropriate gold standard. Decision theorists base this conclusion on the substantial evidence of a 'process-performance paradox' (Camerer & Johnson, 1991). The paradox is that although seasoned experts display considerable advantages in memory, cue use, richness of problem structure, and other cognitive aspects of expert decision-making, they frequently demonstrate little or no performance advantage in decision quality. In some studies, experts are found to perform no better than novices (Oskamp, 1962). In other studies, it is found that experts are out-performed by simple linear models (Dawes & Corrigan, 1974; Dawes et al., 1989; Einhorn & Hogarth, 1975; Meehl, 1954). Experts can fall prey to the same array of cognitive biases that affect novices, resulting in sub-optimal performance and unreliability (Carroll & Payne, 1976; Christensen-Szalanski & Bushyhead, 1981; Einhorn, 1974; Northcroft & Neale, 1987; Oskamp, 1962). More controversial research implies that experts should be completely replaced by statistical models, where possible, since the models have been shown to be more accurate under certain circumstances (Dawes & Corrigan, 1974; Dawes et al., 1989).

The decision-theory literature is not monolithic with respect to the opinion of experts. The expert-disparaging results are challenged by a smaller number of studies that indicate that expert judgment can have utility. Experts appear to be especially good when they can take advantage of rare but highly predictive information in the environment – so-called 'broken-leg cues' (Blattberg & Hoch, 1990; Johnson, 1988; Meehl, 1954; Phelps & Shanteau, 1978). Shanteau (1987, 1988, 1992a) shows that in domains such as weather forecasting, auditing, chess, physics, etc., experts can be quite good. These domains are characterized by greater opportunities for feedback, greater problem structure, and less noise. Shanteau (1987) also finds that lack of feedback, low structure, or noisy environments, tend to predict poor performance. Examples of the latter groups or decision-makers include clinical psychologists, stockbrokers, physicians, and court judges.

The process-performance paradox sets the stage for the assumptions in the decision-theory tradition. Decision theorists assume that even seasoned experts are sub-optimal, and that data-based

statistical models should be provided to aid or replace human decision-makers. An implicit assumption that guides the types of systems that are developed, and are developable, in this tradition is that problems are sufficiently structured that we can build models, collect data, and have an appropriate idea of the functional form of the underlying relationship.

The Need for Decision Support

Researchers are in agreement that some experts are better than others, and therefore that the decision-making performance of many seasoned decision-makers can be improved (theoretically, decisions can at least be improved to the level of the best experts). Hence, formal decision support can have utility. Researchers disagree on the form that formal decision support should take. The computer-science tradition holds the top experts in a field to be the normative standard, leading to a knowledge-based approach. The decision-theoretic tradition holds experts to be fallible and places greater stock in statistical models – a data based view that leads to an emphasis on statistics, models, and optimization.

Scope of Consideration

Nothing described thus far requires a computer. Early techniques for decision support predated the widespread availability of computers, and included decision trees, structured decision aids, and diagnostic flowcharts, many of which had their origins in operations research, economics, or in engineering. With the advent of the computer, both more complex and more integrative systems could be created. Only these interactive, computer-based models are usually considered to be DSS. Hence, we will not review marketing management support systems that are non-computer-based (e.g., marketing models).

Although we concentrate our review of decision support systems on those developed for marketing, it is clear that marketing scientists owe a debt to researchers in other fields who have investigated similar problems. Throughout this chapter, we follow the terminology of Little (1979) and use the phrase *marketing decision support systems* (MDSS) as an umbrella category that includes simple, robust, judgmentally calibrated decision calculus models (Little & Lodish, 1969), such as CALLPLAN (Lodish, 1971) and ADBUDG (Little, 1970); marketing decision support systems such as ASSESSOR (Silk & Urban, 1978); expert systems for marketing such as ADCAD (Burke et al., 1990) and NEGOTEX (Rangaswamy et al., 1989); and

other computer-based, interactive systems whose purpose is to improve decision-making.

ANTECEDENTS TO MDSS

In this section we define the necessary antecedents to thinking about MDSS success. First, we define what we mean by a decision, and then what we mean by a MDSS. Third, we create a taxonomy of decision support systems. The disparate traditions in computer science and decision theory have been largely preserved in the marketing domain. Hence, the taxonomy delineates the boundaries of DSS from other types of management science models and decision aids. It also differentiates knowledge-based systems that stem from the computer-science tradition from the marketing decision support systems that are based on decision theory.

What is a Decision?

Because decision-making has been investigated in many disciplines and across a broad span of time, there is no standardized definition of a decision. What is a decision? For our purposes, we will adopt a consequentialist definition, noting: ‘all the knowledge we have learned and the information we have acquired only add value when the decision is made and the chosen action is executed ... a decision is the identification of and commitment to a course of action’ (Russo and Carlson, this volume). This definition fits nicely into the spirit of MDSS research, in that it is focused on actions that managers need to take. It also emphasizes the economic nature of marketing by making it clear that decisions are bound to economic actions that the manager and the firm then take.

What is a Marketing Decision Support System (MDSS)?

Little (1979) defines an MDSS with direct reference to an action-based definition of a decision. He defines an MDSS 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*’ [emphasis added]. Little also notes that decision support systems should possess other characteristics, such as interactivity, robustness, and completeness on important issues. There is widespread agreement that decision support

systems are meant to support and not to replace the decision-maker, and that the systems should improve decisions (cf. Alter, 1980; Keen & Scott-Morton, 1978). Numerous other definitions of MDSS exist in the literature, but since Little's definition seems to incorporate the most commonly mentioned characteristics, we will use it. In keeping with these definitions, we will not discuss systems that do not support decisions as we have defined the term (e.g., marketing creativity-enhancement programs, and pure forecasting methodologies). One other common part of the definition of MDSS is that the decision is repeated. Thousands of computerized analyses are developed every year in support of one-time decisions. Although these analyses are a form of decision support, they are not considered decision support *systems*, and we will not discuss them.

TAXONOMY OF DECISION SUPPORT SYSTEMS

Researchers in marketing have paralleled the broader decision-support literature. Some researchers have adopted the assumption set of computer science, and others the assumptions of decision theory. Within marketing, there are prominent examples of the two major types of decision support system: knowledge-based systems (also called expert systems, intelligent management systems, and artificial intelligence), and 'plain vanilla' DSS. Not every system can be neatly characterized as either vanilla or knowledge-based, because the systems form a continuum and hybrid systems have been created (Rangaswamy et al., 1987). The taxonomy that follows differentiates the various types of DSS that marketers have created.

Vanilla MDSS play a passive role in the human-machine interaction. They may execute computations, present data, and respond to queries. But they cannot explain their logic, deal with incomplete information, or make logical inferences. Vanilla systems may have a great deal of knowledge built into them, but they are incapable of even simple reasoning. Hence they do not serve as intelligent assistants to a decision-maker.

Some decision environments benefit from having an intelligent computerized assistant. This is the realm of knowledge-based systems. These systems are designed to substitute, in whole or in part, for human expertise. They are sophisticated, specialized computer programs whose goal is to capture and reproduce experts' decision-making processes and to achieve expert-level performance (Naylor, 1983; Turban & Aronson, 1998). Typically, they are restricted to a narrow domain, in which they follow the type of heuristic reasoning used by experts.

Another useful taxonomy can be found in Wierenga and van Bruggen (1997).

Contrasting Vanilla and Knowledge-based MDSS

Three essential features differentiate vanilla MDSS and knowledge-based systems: the use of symbolic processing, the use of heuristics rather than algorithms, and the use of inference techniques which are usually based on logical relationships. This allows the two types of systems to be used under different circumstances. Rangaswamy et al. (1987) recommend that knowledge-based systems should be used when: (1) the key relationships are logical rather than arithmetical; (2) the problem is structured, but not to a level that would allow algorithmic solution; (3) knowledge in the domain is incomplete; and (4) problem solving in the domain requires a direct interface between the manager and the computer system. Knowledge-based systems also differ from vanilla DSS in their ability to offer an explanation of how they arrived at their output, and to explain the logic behind their advice (Awad, 1996).

Vanilla MDSS are based on the tradition of decision theory, management science, and statistics. The emphasis of their development has been on the statistical quality of their predictions. Knowledge-based systems originate from the computer-science tradition. The emphasis of their development has been less on improving decision-making and more on mimicking the process of reasoning and the performance of a recognized expert (or experts) in the domain area. This means that another characterization of knowledge-based systems is to consider them to be descriptive models of the current state of knowledge in a domain. Knowledge-based systems will improve decision-making to the extent that (a) the expert(s) upon whom the system is based are better than the user (and are better than a statistical baseline prediction), and (b) the extent to which the system replicates the experts' knowledge and reasoning processes. Vanilla MDSS can improve decision-making to the extent that the output from the model is superior to that of the decision-maker. Obviously neither type of system will improve decisions if their recommendations are not accepted and implemented by decision-makers.

It should be clear that there is a continuum of MDSS architectures that range from simple facilitation of what-if questions (the simplest vanilla systems) to built-in intelligence, structured knowledge, and intelligent enabling (the most advanced knowledge-based systems). This range includes advanced hybrid systems that combine elements of both traditional DSS as well as knowledge-based reasoning, and it is this continuum of systems that is the focus of the remainder of this chapter.

INTEGRATIVE FRAMEWORK FOR MARKETING DECISION SUPPORT SYSTEM SUCCESS

In this section, we review previous frameworks, contrast different metrics of success, propose a metric of success that we believe is the proper metric, and create an integrative framework that identifies the factors that are most likely to contribute to the success of an MDSS. In this we are indebted to Wierenga et al. (1999), who have also created an excellent overarching framework. Finally, we review the factors that most directly contribute to the success of MDSS.

Review of Existing Frameworks

The classic framework for decision support systems was proposed by Gorry and Scott-Morton (1971), later updated by Keen and Scott-Morton (1978), who combined the work of Simon (1977) and Anthony (1965). Simon argued that decision-making processes fall along a continuum that ranges from highly structured (programmed) to unstructured decisions. Structured processes are routine, repetitive problems for which standard solutions exist. Unstructured processes are fuzzy, complex problems for which there are no known algorithmic solutions. Structure also incorporates the noisiness (predictability) of the system, since adding noise frequently results in a breakdown of standardized solutions. Anthony (1965) defined three broad categories that encompass all managerial activities: strategic planning – the long-range goals and policies for resource allocation; management control the acquisition and efficient use of resources in the accomplishment of organizational goals; and operational control – the efficient and effective execution of specific tasks. Keen and Scott-Morton (1978) divide the structure dimension into three categories – structured, semi-structured, and unstructured, and then create a nine cell matrix with the three Anthony categories, which can be used to classify problems and to select appropriate decision support alternatives. This framework is used with minor modifications by many other authors (e.g., Turban & Aronson, 1998; Wierenga & Ophuis, 1997; Wierenga et al., 1999).

Wierenga and coworkers (Wierenga & Ophuis, 1997; Wierenga et al., 1999) position their integrative framework for MDSS success as a generalization of the Keen and Scott-Morton framework. Their framework consists of five factors: decision situation characteristics, characteristics of the MDSS, the match between demand and supply of decision support, design characteristics, and implementation process. These factors contribute to the success of an MDSS. The framework also properly

recognizes that there are important factors that contribute to the success of a MDSS that are not captured by the Keen and Scott-Morton framework. These omitted factors include characteristics of the implementation process, problem characteristics other than structure, and decision-maker characteristics.

We agree with Wierenga et al. that the classic framework does not fully characterize the factors contributing to the success of an MDSS, but we feel that it is possible to create an even simpler framework. Our framework has only three major stages. Problem characteristics includes the problem definition and the constraints on our knowledge. Adoption and use includes those factors that change the likelihood that an MDSS will be implemented and used, including characteristics of the system itself. These factors lead to success, which we argue should be measured by increased profit to the firm. A further review of the controversy surrounding the choice of a measure of success is discussed in more detail later. We believe that this framework is more general than either the Keen and Scott-Morton or the Wierenga et al. frameworks. Our integrative framework appears below.

Integrative Framework of the Factors that Determine MDSS Success

1. *Problem characteristics*
 - (a) Structuredness includes amount of noise, knowledge of causal drivers and relationships among the drivers
 - (b) Availability of Data
 - (c) Stationarity of the process
 - (d) Type of answer required recognizes that choosing a discrete option (such as acquiring or not acquiring a competitor) is fundamentally different from predicting a continuous measure (such as how much money to allocate to advertising).
2. *System design*
 - (a) Technical validity is the quintessential component of the system. Systems cannot add value without giving good answers for the right reasons.
 - (b) Other system characteristics include ease of use, validity, accuracy, and the quality of alternative systems.
3. *Adoption and use*
 - (a) Decision-maker characteristics relate to the experiences and abilities that the users possess, which change the likelihood of using the system.
 - (b) Organizational factor characteristics relate to the way that the organization

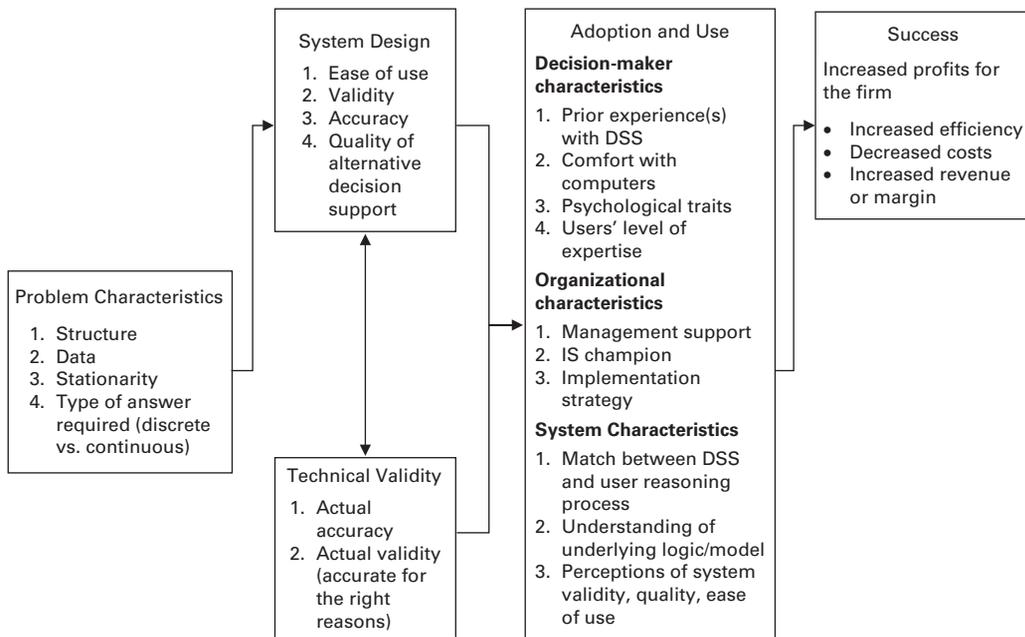


Figure 17.1 Framework for MDSS Success

chooses to implement the system, and how the system is developed.

- (c) System characteristics reflect user perceptions of the validity, accuracy, ease of use, and quality of alternative systems, as well as user understanding of how the system generates its output.
4. *Success* is defined as increased profit for the corporation, or perhaps a subsidiary goal such as increased market share or sales.

What is Success?

Success in the Plain Vanilla Tradition

The question of the dependent variable in DSS research has been a matter of considerable debate. As Wierenga et al. (1999) point out, 'DeLone and McLean (1992), who examined dependent variables in 100 empirical DSS/IS studies, find that "there are nearly as many measures as there are studies."' In marketing, the main discussion of measures of success is due to Wierenga and his colleagues (Wierenga & van Bruggen, 1997; Wierenga et al., 1999). They distinguish four different measures of MDSS success: (a) technical validity, which is the extent to which the underlying MDSS model is valid and the makes statistically accurate

predictions; (b) the adoption and use of the MDSS; (c) impact for the user, usually measured by user satisfaction; and (d) impact for the organization, which are variables such as profit, sales, or market share. These output measures are frequently used by researchers who are studying MDSS success.

Comparison of Metrics

We feel that the criterion for MDSS success must be increased profits for the corporation. This is not to say that the other measures posited by Wierenga et al. (and other researchers) are not relevant. Technical validity is the *sine qua non* of success. Decisions cannot reliably be improved by a model that does not make accurate predictions for the right reasons, which is what is meant by technical validity. Another advantage of technical validity is that it can be tested in the absence of being implemented. Since implementation is costly, technical validity serves as a useful screen – we do not need to implement models that are not technically valid. Technical validity is not viewed as a metric of success in our framework. Instead, it is a necessary system characteristic, and a consequence of system design. This positioning highlights that technical validity is not an end, but a means to achieve the end of better decisions and greater profits.

The two frameworks agree that adoption and use by individual decision-makers must occur if the

system is to affect decision performance. Much of the research in the DSS literature on adoption concentrates on users' satisfaction or perceptions of usefulness of a system. These perceptions are called user impact variables. Our framework subsumes user impact into adoption and use. We do this because it seems logical to assume that systems that are perceived to perform well by users are more likely to be adopted and implemented. On the other hand, there is no logical reason to believe that systems do in fact perform well just because users think they do, or because users perceive them as useful. Therefore, the use of user impact as a measure of success is problematic (Ives & Olson, 1984), and we look at user impact as part of adoption.

Success in the Knowledge-based System Tradition

The vanilla DSS tradition acknowledges that success ideally ought to be measured by better performance on objective, economically relevant measures. By contrast, the literature on knowledge-based systems and expert systems tends to de-emphasize many of the measures of success that are used in vanilla DSS research. In some cases, this lack of emphasis stems from the fact that external validation is difficult to obtain. Such situations arise because even *post hoc* it may not be clear what the proper answer should have been (e.g., what negotiation strategy should be adopted, what clause should be inserted into a contract).

The most apparent example of differences in success metrics is that knowledge-based systems research tends virtually to ignore technical validity (Carroll, 1987; Turban & Aronson, 1998). This is a consequence of the assumption that human experts are the normative standard and that the experts are able to achieve an outstanding level of decision-making performance (Hayes-Roth et al., 1983), and also reflects the difficulty of validating output in knowledge-based environments. In addition, researchers in knowledge-based systems use other criteria that are not used by vanilla DSS researchers. For example, a major goal of knowledge-based systems is to mimic the reasoning process that an expert would have used (Carroll, 1987; Davis & Lenat, 1980; Rangaswamy et al., 1987, 1989; Turban & Aronson 1998).

Our framework incorporates these additional measures for both vanilla and knowledge-based systems. However, we incorporate them as process measures rather than measures of success. The reason is that a good match to the reasoning process of experts is not a substitute for technical validity, increased profit, or for improving decisions. But the match between reasoning processes is relevant if it can be shown that a better match results in greater likelihood of use or implementation. The potential for increased adoption through matched reasoning

holds for both vanilla and knowledge-based systems. Thus we place this factor into the system characteristics dimension of the framework, rather than as a success factor.

Success in the Real World

Carroll (1987) points out that success depends critically on who the users will be. This is because 'improving' decisions requires a baseline for comparison. If the system users are themselves seasoned experts (such as experienced managers or salespeople) then the hurdle is raised. The system must do sufficiently better than those experts to justify the investment and ongoing costs of the MDSS. If the users are novices, then even a substantially non-optimal system might improve performance and generate greater profits for the firm. This is especially true if users are more likely to use a less complex, non-optimal, but easier to understand MDSS. Little (1970) made a similar point, noting that models should be 'simple, robust, easy to communicate with, adaptive, and complete on important issues.' If the definitions of 'simple,' 'complete,' etc. are understood to be defined relatively rather than absolutely, then we arrive at a vision of MDSS success that is quite similar to Urban and Karash's (1971) view of evolutionary model building (Lilien et al., 1992).

In summary, success in MDSS research requires: (1) technical validity; (2) adoption and use; and (3) generation of positive profit for the organization. Technical validity is necessary but not sufficient. Adoption and use is also necessary but not sufficient, because users may adopt and use invalid systems and believe that they are beneficial. A more granular metric could differentiate the impact on individuals' decisions from impact at the firm level. This distinction recognizes that implementation failures due to organizational factors beyond the control of an individual manager (e.g., lack of capital or constraints on production) may reduce the impact of an MDSS for the corporation, even though decisions would be improved by following the recommendations of the MDSS.

INCREASING THE LIKELIHOOD OF MDSS SUCCESS

In this section we examine the factors of our framework that contribute to the success of MDSS. We review the literatures that have been written about the factors and attempt to provide guidance to researchers and practitioners on how to make MDSS development successful. We follow our integrative framework, first focusing on problem characteristics, then on the dimensions

of adoption and use – decision-maker characteristics, organizational characteristics, and system characteristics.

Problem Characteristics

The integrative framework characterizes decision problems along four dimensions: structure, availability of data, the stationarity of the system, and the type of decision to be made.

Structure, Data, and Stationarity

The most important characteristic of the decision problem in terms of the impact on the development of MDSS is the structuredness of the problem (Keen & Scott-Morton, 1978; Russell & Norvig, 1995; Simon, 1977; Turban & Aronson, 1998; Wierenga & Ophuis, 1997; Wierenga et al., 1999). Structure incorporates both the predictability (amount of noise) of the system as well as our knowledge of the underlying causal factors and the relationship among these causal elements. A closely related concept is depth of knowledge, which includes the degree to which our understanding of the structure has been codified by scientific research. The precision with which we know the functional form of the response is an example of depth of knowledge. High structure problems are those with standardized solutions; low structure problems lack standard solutions (Simon, 1977). Availability of data is obviously necessary for developing statistically based MDSS, such as regression or other probability models. Data may not be as helpful if the environment is non-stationary. At a minimum, the underlying analytic models in the MDSS must model the nature of the nonstationarity. Nonstationarity may also impugn the validity of many types of data-based models, and the technical modeling requirements increase. Nonstationarity is less well defined for knowledge-based systems, since these systems are not fundamentally based on data. However, maintaining the knowledge base of the system and keeping the rules and reasoning processes up-to-date is costly and time consuming (Gill, 1995), and is the analogue of nonstationarity in a data-based system.

Type of Output

The type of output that the decision support system is expected to produce is closely related to the type of decision that a manager is required to make. Decisions are differentiated from non-decisions in that an action accompanies a decision. These actions vary on many dimensions, but one central difference among them is whether the action to be taken is discrete or continuous. Conceptually, the

spectrum of actions is a continuum that ranges from binary actions, through choose k of n , to continuous or near continuous allocations. This dimension is frequently ignored in the literature – an oversight in our view. The requirements for decision support depend critically on this dimension, because in the case of a binary action we may need only directional information from the MDSS. Requiring only directional advice may allow for the construction of helpful models even in low structure, low depth of knowledge, low data arenas. Conversely, continuous outputs, such as exact dollar allocations across many marketing projects, will require greater predictability, structure, and depth of knowledge. It is obviously a contradiction to create exact models under conditions of very low structure. As Wierenga et al. (1999:1998) state ‘... we find no papers addressing issues where there is low structure and the environment is turbulent. One reason for the lack of papers may be the focus on continuous rather than categorical decisions. Hence, an area for additional research in low structure environments is to search for directional (categorical), ‘vaguely right,’ models, rather than having no models or ‘precisely wrong’ specifications (Lodish, 1974; Lodish & Reibstein, 1986).

After examining the literature, it appears to us that the most important problem characteristics that contribute to the eventual success of an MDSS are the structuredness of a problem and the type of output that is required. Marketing problems exhibit enormous variation along both dimensions (Wierenga et al., 1999).

Factors Relating to Adoption and Use of MDSS

MDSS adoption and use is primarily affected by decision-maker characteristics, organizational characteristics, and characteristics of the system. We discuss these factors below.

Decision-maker Characteristics

Individual-specific factors influence both the use of MDSS as well as the benefits derived from such use. These factors can be divided into psychological characteristics of users (e.g., cognitive style, problem solving or decision style, personality, etc.), and user-situational characteristics (e.g., user involvement in design and implementation, prior experience with DSS or computers, expertise in the decision domain). There is only a small amount of research within marketing on the interactions between decision-maker characteristics and MDSS success. Where the effects have been studied in a marketing-specific setting, the results have been similar to those in the general DSS literature (Wierenga et al., 1999; Zinkhan et al., 1987). We

thus review both marketing-specific results as well as results from the general literature.

Psychological Characteristics of Decision-makers

The most commonly mentioned decision-maker characteristics in the general DSS literature are psychological traits. However, most articles focus on one trait, called cognitive style (Alavi & Joachimsthaler, 1992; Turban 1995; Turban & Aronson, 1998; Wierenga & Ophuis, 1997; Wierenga et al., 1999). As a construct, cognitive style is not as well-defined as it might be. A consensus definition is that cognitive style is a multidimensional construct that describes the characteristic ways that individuals process and use information in the course of solving problems and making decisions (Huysmans, 1970; van Bruggen et al., 1998). It seems reasonable to believe that such a trait would influence MDSS use, but there are numerous problems with using it to predict MDSS success. First, most research in the DSS field focuses on only one subdimension of cognitive style, the analytic/heuristic subdimension (Huber, 1983; Wierenga & Ophuis, 1997; Wierenga et al., 1999; Zinkhan et al., 1987; Zmud, 1979). This subdimension is usually further simplified to a binary classification with opposite types of decision-makers at the extremes. At one end of the spectrum are presumed to be high analytical decision-makers, who prefer reductionist reasoning to a core set of underlying relationships; at the other end are low analytical decision-makers, who tend to look for heuristic solutions or to solve problems by analogy, frequently in a more holistic manner (Huysmans, 1970; van Bruggen et al., 1998).

The problem with this approach is pointed out by Huber (1983), who notes that the dichotomization is an oversimplification, since cognitive style is actually a continuous variable, and most people are neither completely analytic nor completely non-analytic. Furthermore, no research appears to exist that quantifies the relationship between a given score on a personality test and the degree to which heuristic (low analytical) reasoning will be used (a result also found by van Bruggen et al., 1998: 647). Worse, much of the research on cognitive style has used personality measures that were not developed to accurately measure the type of specific analytic or analogical reasoning that might interact with decision support system use (Alavi & Joachimsthaler, 1992). Finally, substantial evidence from the literature on expertise suggests that expertise is domain specific (Chi et al., 1988; Ericsson, 1996; Ericsson & Smith, 1991). This means that decision-making within the area of expertise can use a different cognitive style than ordinary day-to-day thinking (Shanteau, 1992a, 1992b). It is day-to-day thinking (general patterns of thought) that are measured by the personality test. However, it is usually decision-making within the area of expertise that we measure in MDSS research.

The most damning evidence against the use of cognitive style as a basis for recommendations in MDSS design is that the psychological factors that have been measured, including cognitive style, appear to be poor predictors of DSS success. Alavi and Joachimsthaler (1992) note in their meta-analysis of DSS implementation research:

overall, these results suggest that the relationship between cognitive style and DSS performance is small ... [it] translates to a correlation of .122, which implies that, on average, less than 2 percent of DSS performance can be explained by this dimension ... [the] meta-analytic results indicate that [general] psychological factors have only a small to moderate effect on DSS performance and user attitudes. (1992: 103, 109)

Huber (1983) reaches similar conclusions about the use of cognitive style in DSS research.

Despite the small effect sizes associated with cognitive style, researchers seem to agree that high analytical decision-makers outperform heuristic decision-makers, and that they also have more positive attitudes towards DSS (Alavi & Joachimsthaler, 1992; Larreche, 1979; Leonard et al., 1999; Ramaprasad, 1987; Van Bruggen et al., 1998; Zinkhan et al., 1987). Lusk and Kersnick (1979) and Cole and Gaeth (1990) have both shown that high analytical decision-makers structure and solve problems in a manner that positively affects decision performance. Whether structuring problems in a manner which increases performance interacts with better use of decision support systems remains a contentious topic. Research on the interaction of psychological characteristics with preference for the use of decision support systems (De Waele, 1978; Hunt et al., 1989) as well as performance improvements/declines (Benbasat & Dexter, 1982, 1985) is inconclusive.

We believe that these inconclusive results are likely to have been caused by heterogeneity in the user population. For example, some studies have found that decision-makers prefer a DSS that will complement their weaker cognitive style, but other studies have found that decision-makers prefer to match their stronger style. Similar ambiguities exist when looking at the interaction of cognitive style DSS use and performance improvements. These reversals and small effect sizes could easily be due to unobserved heterogeneity of preferences (Hutchinson et al., 2000): some high analytical decision-makers prefer to complement their weaker style, others their stronger style, and the same is true for low-analyticals. This is a critical area for further research, and an example of an area where greater experimental rigor would yield valuable theoretical and practical results.

Other User Characteristics Another frequently studied user characteristic is the expertise of the MDSS user. The use of decision support systems

and decision aids is generally believed to help novices more than experts, and low-analyticals more than high (Benbasat & Dexter, 1982, 1985; Van Bruggen et al., 1998). Spence and Brucks (1997) demonstrate the expert–novice comparison using a (non-computerized) decision aid. They found that novices especially benefited from the use of a decision aid in a structured housing valuation task. Experts performed equally well with and without the aid. This result underscores the need to specify the user pool for the MDSS. Depending on whether the primary users will be experts or novices, the design of the MDSS may change, and the system must also work better than the alternative option of non-use if decisions are to improve.

A final user characteristic is prior experience with a decision support system. Not surprisingly, prior positive experience increases likelihood of adoption and use (Alavi & Joachimsthaler, 1992). We believe that prior experience with computer technology in general might also be predictive. This is not a user characteristic that is often studied. But it is not difficult to believe that a manager who does not know how to read his email is unlikely to use a decision support system, or that a technological wizard is more likely to do so. Over the last several years the authors of this chapter have seen enormous increases in the computer sophistication of salesforce managers in executive education. This type of increased familiarity with computers, spreadsheets, and financial models can only increase the likelihood of MDSS use. One troubling aspect of increasing familiarity is that the nonstationary nature of the ‘familiarity base’ limits our ability to draw inferences from research that may be 20 to 30 years old, as much of it is.

Organizational Characteristics

Organizational characteristics include the support of top management, implementation strategy, and user involvement in design and implementation.

Management Support ‘The evidence for the need for management support is so strong that any attempt to implement a system without it or without the related conditions of commitment and authority to implement will probably result in failure’ (Hanssens et al., 1990: 324). The need for top management support appears to be especially important for knowledge-based or expert systems (Tyran & George, 1993). It seems equally intuitive that the existence of an IS champion increases implementation success, though the literature is unclear how important this characteristic is.

Implementation Strategy Implementation strategy encompasses a wide variety of activities that help users learn and use a new MDSS. Training and user involvement in the design process are two

of the most frequently mentioned aspects of implementation strategy.

Many decision-makers must justify their decisions to more senior management, or to outside parties such as analysts, courts, or clients. Adoption and use become much less likely when managers do not understand the workings of the MDSS, or the logic behind the output. We strongly feel that MDSS cannot be mere black-boxes. In their meta-analysis, Alavi and Joachimsthaler (1992) find that training in the use of the DSS aids implementation. They also point out that most studies interpret training very narrowly. These studies tend to train very specific aspects of use (the specific hardware and software skills needed to interact with the DSS). They point out that adoption and use might increase more with training in how the model works. Although the authors note wide variance in the reported effect sizes of the component studies, the larger (better) effect sizes are related to the training variable in field studies rather than laboratory studies, which lends support to the importance of training in real-world environments.

User involvement is generally defined as the participation in the system development process by the users. It is stridently advocated throughout the process of DSS development, and is viewed as essential for success (Lilien et al., 1992; Sprague & Watson, 1993; Turban & Aronson, 1998). In their meta-analysis, Alavi and Joachimsthaler (1992) find support for user involvement increasing the success of implementation, but they also find large measurement error. They attribute this error to the construct being poorly defined. A more critical view is espoused by Ives and Olson (1984), who also examined the relationship between user involvement and MDSS success. Their review of the role of involvement in MIS success concludes that ‘much of the existing research is poorly grounded in theory and methodologically flawed; as a result the benefits of user involvement have not been convincingly demonstrated’ (Ives & Olson, 1984: 586). They note that involvement in the design process increases user satisfaction and perceived usefulness of the system (and may therefore help implementation). But since these constructs are not logically related to technical validity or to the appropriate use of the system (e.g., users might like the system because its recommendations are easy to override), the link between involvement and better decisions is weak.

DECISION SUPPORT SYSTEM IN MARKETING – DO THEY WORK?

As we remarked in the introduction, Wierenga et al. (1999) is also an excellent review of decision support system research in marketing.

Field Studies

Many authors have lamented the lack of studies that investigate the effects of MDSS in the field rather than in the lab (Sharda et al., 1988; Wierenga et al., 1999). Field studies with control groups are the gold-standard, since the magnitudes and directions of the performance changes are most likely to mirror what practitioners would achieve. Perhaps the best-known field study is Fudge and Lodish (1977), who implemented the CALLPLAN model (Lodish, 1971) in a field test with matched pairs of sales representatives. The result was that the average salesperson with access to the model had 8.1% greater sales than those without model support (though not every model-user performed better than every non-user). The authors achieved this result in spite of the fact that the control group judgmentally estimated the parameters of the model, they just did not receive the model's recommendations. The carefulness of this study, combined with its matched pair comparison, make it the gold-standard against which other research can be judged. Furthermore, this study fits directly into our integrative framework. The users were trained on the use of CALLPLAN, adopted and used it, and the output measure was profit. There were no claims that CALLPLAN generated completely optimal allocations. CALLPLAN is a decision-calculus model – its parameters are judgmentally calibrated, users have the flexibility to arrive at nearly any answer they desire, and the model is not 100% complete. Nevertheless, the model works better than the average salesperson's intuition, is simple enough to be understood, and was therefore more likely to be adopted by the test users.

Other field studies include Lodish et al. (1988) and Gensch et al (1990). The Gensch study compared two districts that used the MDSS against control districts. The MDSS-using districts had large increases in sales at a time when the market as a whole was going down; the control districts showed declines in sales. In this case, the model was rolled out company-wide on the basis of its success, making it difficult to quantify exactly what the gains due to the model were, but they were very likely positive. Lodish et al. implemented a salesforce size and deployment model, with parameters calibrated by a modified Delphi technique. Although the full recommendations of the model were not implemented, substantial gains in sales were documented even with the more limited implementation. These successes in the field support our claim that for many problem types it is only necessary to get directionally correct, order of magnitude information, rather than 'precisely wrong' optimized output from a technically deficient model. Directional output was sufficient to increase profits and sales in these studies, and made the model and recommendations more likely to be accepted by the users.

Laboratory Studies

Most other studies of MDSS success have been performed in the laboratory. Chakravarti et al. (1979) conducted a laboratory study using a simplified version of ADBUDG (Little, 1970), the original decision calculus model, in the context of a game simulation. Participants in an executive education program played the game, making advertising allocation decisions over several rounds. After participants received some experience (training), a portion of the participants were given access to ADBUDG. If the model were properly calibrated, it would have generated the optimal solution, but the executives did not know this. Oddly, those who used the model earned *less* profit than those in the control group, implying that the MDSS had actually led to worse accuracy in predicting share.

In a very similar experiment, however, McIntyre (1982) used CALLPLAN with a group of MBA students in a game setting. He found that access to the MDSS enhanced decisions along multiple dimensions, including better average profit earned, fewer large errors, and faster learning. The number of allocation units and the noise level in the environment had little interaction with the benefits – the model helped across problem sizes and noise levels.

Several different commentators (Little & Lodish, 1981; McIntyre, 1982; Wierenga et al., 1999) attempt to reconcile the results of these two studies by looking at the study characteristics. These reconciliations generally point out that Chakravarti et al. use a more complicated functional form for the advertising response that incorporated lagged effects from the previous period. These carryover effects induce a nonstationarity in the response that is then confounded with previous treatments. This explanation squares with the research of Serman (1987, 1989), who found that people have a very difficult time dealing with lagged effects, even within the setting of a simple game. Little and Lodish point out that using ADBUDG rather than BRANDAID was unfortunate, since BRANDAID removes the confound between carryover and current advertising changes. Similar changes to the parameter input structure are found in Lodish (1980), who modified DETAILER (Montgomery et al., 1971) to simplify the dynamics, with corresponding success on implementation.

Van Bruggen et al. (1998) found positive effects of MDSS use on market share and profit in the context of the MARKSTRAT game. They believe that their MDSS is effective because it assists users in identifying the important variables and by aiding decisions that are based on those variables. They also found that decision-makers using an MDSS are less susceptible to applying the anchoring and adjustment heuristic, which often results in poor outcomes.

Blattberg and Hoch (1990) demonstrated that a combination of managerial judgment and MDSS

output resulted in superior performance compared with using either by itself. They attribute the result to the fact that managers and models have complementary skills. Models are more consistent and better at integrating information; humans are better at identifying 'broken-leg cues' (Meehl, 1954) – diagnostic variables that have not yet been incorporated into the model. They recommend placing 50% weight on manager and model judgment as an heuristic.

Hoch and Schkade (1996) performed a laboratory experiment that attempted to determine whether it is better to design a MDSS that will capitalize on managerial strengths or compensate for weaknesses. There is a tradeoff because designing to capitalize on human strengths may exacerbate the weaknesses. They varied both the predictability of the environment and the type of decision aid. One decision aid was a pure database lookup, which capitalizes on humans' strength in pattern matching. The other MDSS was a statistical model, which compensates for human inconsistency. A third group had access to both models. The authors conclude that in a highly predictive environment, the various MDSS improve decision performance about equally. In low predictive environments, database lookup was the worst possible support strategy. Supporting decision-makers' strengths will not necessarily improve decisions.

Hybrid and Problem-oriented MDSS

A whole class of marketing decision support systems fall in between the more classic DSS, expert systems, and test marketing. Systems such as ASSESSOR (Silk & Urban, 1978), Promotionscan (Abraham & Lodish, 1987), the segment selection methodology developed by Montoya-Weiss and Calantone (Montoya-Weiss & Calantone, 1999), and various product design and positioning optimizers (e.g., Green & Krieger, 1985, 1992) are hybrid systems (also called 'problem-oriented systems,' Rangaswamy et al., 1987). These systems do not make explicit the reasons for their recommendations, operate in a higher structure world where algorithmic solutions exist, and require all data inputs to be present, making them superficially similar to vanilla DSS. But these systems also have logical structures that are easy to follow, robust and complex models of the world, and 'what if' capabilities for sensitivity analysis. These characteristics make them more similar to knowledge-based systems because: (1) the easily understandable logical structure makes it possible to make explicit the reason for the recommendation; (2) the complex and more complete underlying models incorporate a great deal of accumulated knowledge; and

(3) 'what-if' and sensitivity analysis reduces the need for every input to be filled in precisely, and also allows recommendations to be accompanied by a measure of confidence.

ASSESSOR is a pre-test evaluation system for new packaged goods. Using it, a sample of consumers are surveyed about current category usage. They are then shown advertising for a new product and its major competitors, and participate in a simulated shopping experience. At a later time, subjects report their repeat purchase intentions. Measures taken during the process are used to forecast the product's expected market share and to diagnose product problems. The MDSS has been implemented for hundreds of product evaluations across many companies, and has helped lower the failure rate of new products in test market by almost half, and thus significantly increases corporate profits (Urban & Katz, 1983). Within our framework, it is technically valid, easy to use and understand, and implementable with the support of top management. This makes it very successful.

Promotionscan (Abraham & Lodish, 1993) is an automated system for measuring short-term incremental volume due to promotions by automatically computing an appropriate baseline sales volume and by adjusting for other variables that might be confounded with the promotional effect. The output of the system is used by managers to determine which retail promotional options should be chosen or negotiated. The system (or similar systems based on it), are in use in firms that sell over 50% of the frequently purchased packaged goods revenue in the US (Lodish, pers. comm.) The initial sample application showed a 15% increase in incremental sales with the use of Promotionscan. Anecdotally, one manager says:

The system provides answers to questions such as what the competition is doing, which distribution outlet is most effective, what merchandising strategy would prove productive and other questions plaguing our sales force. To get the same information from the paper reports would take three to four times longer. (Progressive Grocer, 1994)

A recent advance on Promotionscan is due to Silva-Risso et al. (1999), who disaggregate the store level estimates used in Promotionscan so that a more precise measure of incremental vs. borrowed sales can be estimated. A constrained model's recommendations were implemented by a firm, who reported positive results from the implementation.

Montoya-Weiss & Calantone (1999) used a company-wide implementation of a 'problem-oriented' MDSS for the selection of industrial product markets. The DSS is actually a complete system of methodologies that includes problem structuring, segment formation, segment evaluation and selection, and segment strategy description. These steps are accomplished by a conjoint analysis, cluster

analysis, product design optimization, and a multi-objective integer programming model. They implemented the entire MDSS at an automotive supply company, and the company adopted most of the model recommendations, blended with managerial judgment. In the two years after implementation, the company realized a 5% savings in operating expenses, a 4.5% increase in sales revenue, a 3% decrease in cost of goods sold, and a 15.8% increase in net profit. The authors also simulated the company's performance without the system as a comparison metric, and compared the model's estimates to actual performance. Compared with their simulation, actual performance represented a 36% increase in net profit, and a 41% reduction in communication costs. A controlled implementation would have made a stronger test, but these results are impressive nonetheless.

Product line optimizers such as SIMOPT (Green & Krieger, 1992) use a combination of conjoint and optimization techniques to aid in the optimal design and positioning of products. They have not been extensively validated in the field, but it is possible to look to the success of conjoint in general for the likely performance of these MDSS (Wind et al., 1989).

KNOWLEDGE-BASED SYSTEM IN MARKETING – DO THEY WORK?

Knowledge-based systems enjoyed their heyday in marketing from the mid-1980s to the early 1990s, paralleling the golden age of expert systems in computer science. Two of the most significant attempts at system design are NEGOTEX (Rangaswamy et al., 1989) and ADCAD (Burke et al., 1990), which help to prepare strategies for international negotiations and advertising strategies respectively. Other systems include INNOVATOR (Ram & Ram, 1988; Ram & Ram, 1989), which helps to screen new product ideas in the financial services. Business Strategy Advisor (Schumann et al., 1989; see also Wierenga, 1990) uses a BCG-like matrix to make strategic recommendations. These systems share a number of characteristics. First, they all operate in domains for which there is no easily agreed upon algorithmic solution, making the expert systems methodology appropriate (Rangaswamy et al., 1987, 1989). Second, the knowledge base underlying each of these systems is derived from industry experts, published material, the authors' experiences, or a combination of the three. Third, in keeping with the computer-science tradition, which tends to de-emphasize technical validity measures, these systems were not tested against established experts or other systems in either field tests or laboratory settings at the time of publication.

The lack of strong tests of validity makes it difficult to determine whether the systems 'work,' especially in the context of our framework for implementation and the previous discussion. ADCAD was informally validated by expert comments. INNOVATOR's authors used an *ad hoc* comparison with an expert in their initial article, and then tested the completeness of the knowledge base, the consistency and accuracy of the decisions made by the system, and the reasoning process by which the system made decisions several years after initial publication (Ram & Ram, 1996). The most in-depth validation was performed for NEGOTEX. It included full reviews using formal questionnaires by seven leading academics who specialized in negotiations, and by MBA students taking a marketing strategy class, as well as informal feedback from several practitioners. These validations do not comprise field trials, nor do they use laboratory experiments (for example, by having a group of experts rate the recommendations of subjects with and without access to the systems). We should stress that the approach to validation employed by these authors is in keeping with the normal validation methods in the expert systems literature. Rangaswamy et al. (1989) point out that the validation of expert systems has been a topic of controversy, and quote Sheil (1987), saying: 'there is no way to check that all knowledge is "correct" and no way to prove that the system has no significant gaps in its coverage' (1989: 32). Sheil is correct, but his point fails to address the larger issue.

Knowledge-based systems should be validated better. One reason that the expert system golden age faded is that lack of validation made it difficult to justify the substantial costs of development and system maintenance. Yet, it is possible to validate a knowledge-based system. Its recommendations can be compared to recommendations of experts in the field. Alternatively, a split-pool of likely users can be constructed, and the responses of the pool with access to the system can be compared against those without access. To be careful, the procedure followed in the test of CALLPLAN should be followed, where the control group still answers all the questions that the system-enabled group answer, just without the model output.

Carroll (1987) examines expert system performance across numerous application areas. She examines the two major standards of validation: comparison of the recommendations to those of experts, and comparison of the logical process. Recall that one of the main selling points of knowledge-based systems is the similarity in the reasoning process between model and user. In outcome performance, she finds that expert systems do about as well on average as the humans on whom they are trained under ideal circumstances. But the quality of knowledge-based systems' recommendations deteriorates much faster than human experts'

at the boundary of the problem domain, or if the underlying logical relationships change slightly (Carroll, 1987; Davis, 1984; Davis & Lenat, 1980; Davis et al., 1987). Carroll also finds that 'expert systems are generally not superb [descriptive] models of human expertise' (Carroll, 1987: 285), meaning that the logic patterns followed by the system are not fully consistent with those of the human expert. This is because most knowledge-based systems are limited to rule-based inference, but rules make up only a subset of human inferential techniques, and because the rule elicitation process is necessarily inaccurate (see Nisbett & Wilson, 1977 for more discussion on this).

There is no definitive answer to the question 'do knowledge-based systems in marketing work?' By the success metrics that we have defined, we have almost no data. Evaluated against subordinate necessary conditions such as technical validity, the systems appear to perform well under ideal circumstances, but are likely to be non-robust. In general, systems only partially match the reasoning processes used by experts in the domain. Knowledge-based systems are more expensive to build (Turban, 1995; Turban & Aronson, 1998). They are also abandoned more frequently than vanilla DSS once implemented, primarily due to lack of acceptance by users, and the costs of transitioning from a development to an ongoing maintenance outlook (Gill, 1995; Turban, 1992, 1995).

Amidst these apparent disadvantages, Rangaswamy et al. (1989) provide a completely different rationale for the continued importance of knowledge-based systems. They note that 'the mere process of building an expert system can contribute to the marketing discipline independent of whether the final system is used by decision makers' (Rangaswamy et al., 1989: 33). The meta-analytic synthesis performed while developing knowledge-based systems may help to point out gaps and inconsistencies in current knowledge and may help the developers to develop empirical generalizations. Although these are both worthy goals, it then remains to be shown that expert system design is a better way to accomplish this goal than other approaches to research synthesis.

CUTTING EDGE SYSTEMS

Cutting edge systems that have been developed by computer scientists are just beginning to hit marketing. Firms are increasingly attempting to resurrect AI under the new names or 'machine learning' and 'rule discovery.' Machine learning encompasses various ways to allow the coefficients of a model to automatically update over time. Rule discovery is the process of generating symbolic rules of the form 'if <pattern> then <action>' using

discrete variables as the inputs (Cooper & Giuffrida, 2000).

The most common technique used for machine learning is neural networks (Turban & Aronson, 1998). Neural networks process information in a similar manner to biological neural systems. They accept multiple inputs into a processing element, integrate the inputs according to a weighting function, and then the processing element either produces an output or does not (binary activation). The output from the first layer of processing elements is typically fed into another layer of elements, and then to an output node or nodes (which recommends an action). When used to recommend an action, a neural network is mathematically equivalent to a nonlinear discriminant function, whose parameters are determined by the pattern of weights among the processing units.

Neural networks have been found to be particularly good in pattern recognition, generalization and abstraction of prototypical patterns, and interpretation of noisy inputs. All of these uses are of course 'trained' into the network, rather than having a modeler input and test a structural equation. Their primary successes thus far have been in speech and handwriting recognition, and in prediction of credit default and business failure (Tam & Kiang, 1992; Wilson & Sharda, 1994). Their advantages include the ability to learn from past data, to model highly complex relationships among the data, to be easily maintained and updated (just train them on new data), and to process data quickly. The major disadvantage of neural networks is the nearly complete lack of explanation or reasoning behind the output. This black-box nature is because the function is nonlinear and the connection weights between processing elements have no obvious interpretation (Turban & Aronson, 1998).

Rule discovery refers to procedures that automatically create 'if-then' types of rules from existing data. Cooper & Giuffrida (2000) provide an excellent taxonomy and summary of rule discovery methods as they apply rule discovery techniques to the residuals of a promotion forecasting system (PromoCastTM). Because they apply the rule discovery techniques to the residuals of a forecasting model, the discovered rules represent local variations that would not have been captured by a standard market response model. Furthermore, the rules include a measure of confidence so that managers can know how 'certain' they should be in applying the rule. Their extensive validation on a holdout sample demonstrates a significant improvement in forecast accuracy.

CONCLUSIONS AND FUTURE RESEARCH

We have placed MDSS in context, offered a taxonomy of decision support systems, clearly defined

a measure of success, created an integrative framework of factors contributing to MDSS success, and have examined the assumptions that underlie the relationship between our framework and existing MDSS. In summary, both the computer science and decision theory research traditions have reached the conclusion that it is possible to improve even expert decision-making. Computer science assumes that the top expert(s) in a field are the gold standard of performance. This assumption has given rise to knowledge-based decision support, which attempts to capture the knowledge and reasoning process of an expert in the form of a computer system. By contrast, decision theorists have argued that the normative standard should be an objective outcome measure, either a statistical model or an expert, depending on which is better. The decision theory and psychology literatures conclude that statistical models are frequently better than human experts. This conclusion has led to a focus on data-based, statistically grounded systems. These two types of system anchor the ends of a continuum. On one end are plain vanilla systems that almost never mimic the reasoning process of human experts, and may be just simple 'what-if' simulators. At the other end of the continuum are knowledge-based systems, which are 'sophisticated and highly specialized computer programs that try to capture and reproduce experts' decision-making processes (Naylor 1983)' (Carroll, 1987: 280). In the middle are hybrid systems that include automatic diagnostic tools such as Promotionscan, and test-forecast models such as ASSESSOR.

Regardless of what type of system is contemplated, we argue that increased profit is the proper metric to use to measure DSS success. Our insistence on profit as the proper metric of success is a departure from most previous literature. In most prior literature, system validity and implementation have been treated separately, with user opinions of system efficacy or measures of adoption and use frequently substituting for profit as the metric of success. By proposing that increased profits for the firm is the proper metric, we propose that the standard for success must include both accuracy *and* implementation. This reconceptualization of success implies that there is a valid tradeoff between factors that affect adoption and use and the accuracy of the system, if increasing accuracy has the potential to decrease probability of implementation. Such a situation may arise under a variety of conditions, notably when the increase in accuracy comes along with increased system complexity or when the increase in accuracy comes only with the use of a model that the users do not understand (a black-box model). It should be clear that optimizing accuracy of the system is not necessary, and in fact would be viewed negatively under our metric if adoption and use is negatively related to the factors that increase accuracy. Similarly, characteristics of

the system, such as the similarity of the system to human reasoning or the ability to explain why certain output has been generated, are relevant only if they increase accuracy or increase the likelihood of adoption and use. The existence of the tradeoffs outlined here should not be taken to be an abandonment of theory. Technical validity is still the *sine qua non* of DSS construction – the system must be accurate for the right reasons. But, although models should be 'simple, robust, easy to communicate with, adaptive, and complete on important issues' (Little, 1970), the definitions of 'simple,' 'robust,' etc. should be viewed relatively. Accuracy should also be viewed relatively – better to have a vaguely right and implemented solution that is better than what previously existed, than to have a precise, mis-specified model that is wrong, or an ideal model that not used.

MDSS Success

With the exception of Chakravarti et al., studies of vanilla DSS effectiveness in marketing generally find a positive main effect of the MDSS on performance. These results are somewhat anomalous within the broader field of DSS research. Sharda et al. (1988) review this general literature (including many of the marketing results) and conclude: 'field and laboratory tests investigating superiority of DSS over non-DSS decisions show inconclusive results' (Sharda et al., 1988: 144). As Sharda et al. point out, part of the reason that many DSS may not appear to improve decision-making is that most of the non-significant studies are based on a one-time measurement of performance. Furthermore, in most of the inconclusive studies the DSS was a 'black-box' – subjects did not know or understand the workings of the model. They go on to note that MDSS are likely to be used more than once, users are likely to be trained, and some understanding of how the DSS works is likely to be transmitted. These criticisms imply that the inconclusive studies confound lack of training with measures of the ability of the DSS to improve decision-making. In marketing, the MDSS used are rarely black-boxes, and most subjects receive some training on how the MDSS works. This may explain the positive effects of MDSS in marketing. However, Sharda et al. raise a critical point about the rigor of MDSS (and more broadly, DSS research). Many inconclusive results are likely to have been caused by heterogeneity in the user or test populations. We raised this point in the context of examining the evidence of the effect of individual differences on successful DSS adoption and use, but the argument applies to almost every aspect of DSS measurement. Users of DSS may be heterogeneous on a variety of dimensions, including familiarity with computers, expertise in the decision-making task, psychological traits, and

prior use of DSS. These differences will affect the success of DSS, and must be controlled for and modeled in the research. Hutchinson et al. (2000) provide extensive recommendations and procedures for diagnosing unobserved heterogeneity.

Future Research

In this section we summarize some emerging technologies and trends in DSS and knowledge-based systems. We also make recommendations to researchers. Wierenga et al. (1999) make a number of excellent recommendations as well.

One important area for future research is to engage in more validation along the continuum of validation approaches. The most important priority is to generate more controlled field studies with both existing and new models. This is critical if we are to get information about the entire process of MDSS development, implementation, and measures of success in complex environments. Secondly, we should encourage systems to be evaluated in the laboratory as thoroughly as possible if they are not going to be implemented in the field. One way to do this is to use subjects who are representative of the user pool in as realistic a decision-making environment as possible. Assuming a representative user pool, researchers should collect measures of likelihood of adoption and use, in addition to assessing improvement in decision quality (or efficiency). Where possible, simpler models should be compared to more complex models on both the accuracy and likelihood of adoption measures so that the tradeoff between complexity and adoption can be found (if it exists). Prior to extensive laboratory testing, every system (including knowledge-based systems) should be assessed for technical validity – against an appropriate array of possible other decision-making systems and human experts at a minimum.

Although field studies are the gold standard in MDSS research, simpler non-DSS solutions could also prove very helpful. One simple and informative area of research would be to just keep track of the decisions that are made within a company. Then one could at least compare the before, during, and after DSS-use performance, controlling for other variables. Many firms keep detailed records of their plans, but virtually no information about what they actually did. With the help of researchers, these firms could then at least directionally determine the effect of the MDSS on decision performance.

We also recommend an increased emphasis on solving more complex problems. In order to do this, three separate research streams must be united. The psychological literature on managerial decision-making should be employed to determine whether and under what circumstances DSS should reinforce managerial strengths versus compensating for

weaknesses (along the lines of Hoch & Schkade, 1996). Second, more hybrid systems should be constructed, with data-gathering methodologies, intuitive processing models, and a combination of logical and data-processing capabilities built into them. Third, we need to renew the emphasis on directionally correct solutions and the validation of such solutions. This is a reasonable initial goal in highly complex environments. We need more research into the process of how managers go about making directional decisions, and the level of expertise that they can achieve in complex environments.

Finally, it is the dream of most marketing managers to have automated systems that can extract relevant information from customer and transaction databases. With the explosion of transactions in the online world, data-mining, machine learning, and rule discovery will become more important. Marketers should be leading this wave of research with colleagues in computer science.

The authors would like to thank Arvind Rangaswamy, Robin Wensley, and Christopher Van den Bulte for comments on earlier drafts of this paper.

References

- Abraham, Magid M. & Lodish, Leonard M. (1987) Promoter: an automated promotion evaluation system. *Marketing Science*, 6 (2), 101–23.
- (1993) An implemented system for improving promotion productivity using store scanner data. *Marketing Science*, 12, 3 (Summer), 248–269.
- Alavi, Maryam & Joachimsthaler, Erich A. (1992) Revisiting DSS implementation research: a meta-analysis of the literature and suggestions for researchers. *MIS Quarterly*, 16 (1), 95–116.
- Alter, Steven (1980) *Decision Support Systems: Current Practice and Continuing Challenges*. Reading, MA: Addison-Wesley Publishers.
- Anthony, R.N. (1965) Planning and control systems, a framework for analysis. Boston, MA: Division of Research, Graduate School of Business Administration, Harvard University.
- Awad, Elias M. (1996) *Building Expert Systems*. Minneapolis, St. Paul: West Publishing Company.
- Benbasat, Izak & Dexter, Albert S. (1982) Individual differences in the use of decision support aids. *Journal of Accounting Research*, 20 (1), 1–12.
- (1985) An experimental evaluation of graphical and color-enhanced information. *Management Science*, 31 (11), 1348–65.
- Blattberg, Robert C. & Hoch, Stephen J. (1990) Database models and managerial intuition: 50% model + 50% manager. *Management Science*, 36 (8), 887–99.
- Burke, Raymond R., Rangaswamy, Arvind, Wind, Jerry & Eliashberg, Jehoshua (1990) A knowledge-based system for advertising design. *Marketing Science*, 9 (3), 212–29.

- Camerer, Colin F. & Johnson, Eric J. (1991) The process-performance paradox in expert judgment: how can experts know so much and predict so badly? In *Toward a General Theory of Expertise: Prospects and Limits*, edited by K. Anders Ericsson & Jacqui Smith. New York: Cambridge University Press, pp. 195–217.
- Carroll, Barbara (1987) Expert systems for clinical diagnosis: are they worth the effort? *Behavioral Science*, 32 (4), 274–92.
- Carroll, John S. & Payne, John W. (1976) The psychology of parole decision process: a joint application of attribution theory and information-processing psychology. In *Cognition and Social Behavior*, edited by John S. Carroll & John W. Payne. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Chakravarti et al. (1979)
- Chi, Michelene T.H., Glaser, Robert & Farr, Marshall J. (eds) (1988) *The Nature of Expertise*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Christensen-Szalanski, Jay J. & Bushyhead, James B. (1981) Physicians' use of probabilistic information in a real clinical setting. *Journal of Experimental Psychology: Human Perception & Performance*, 7 (4), 928–35.
- Cole, Catherine A. & Gaeth, Gary J. (1990) Cognitive and age-related differences in the ability to use nutritional information in a complex environment. *Journal of Marketing Research*, 27 (2), 175–84.
- Cooper, Lee G. & Giuffrida, Giovanni, (2000) Turning datamining into a management science tool: new algorithms and empirical results. *Management Science*, 46 (2), 249–64.
- Davis, Donald L., Davis, R. Dean & Shrode, William S. (1987) Decision support systems (DSS) design for operations managers: an empirical study of the impact of report design and decision style on effective choice. *Journal of Operations Management*, 7 (1–2), 47–62.
- Davis, R. (1984) Amplifying expertise with expert systems. In *The AI Business: The Commercial Uses of Artificial Intelligence*, edited by Karen A. Prendergast & Patrick Henry Winston. Cambridge, MA: MIT Press, pp. 17–40.
- Davis, Randall & Lenat, Douglas B. (1980) *Knowledge Based Systems in Artificial Intelligence*. New York: McGraw-Hill International Book Co.
- Dawes, Robyn M. & Corrigan, Bernard (1974) Linear models in decision making. *Psychological Bulletin*, 81 (2), 95–106.
- Dawes, Robyn M., Faust, David & Meehl, Paul E. (1989) Clinical versus actuarial judgment. *Science*, 243 (4899), 1668–74.
- DeLone, W.H. & McLean, E.R. (1992) Information systems success: the quest for the dependent variable. *Information Systems Research*, 3 (1), 60–95.
- De Waele, M. (1978) Managerial style and the design of decision aids. *OMEGA*, 6 (1), 5–13.
- Einhorn, Hillel J. (1974) Expert judgment: some necessary conditions and an example. *Journal of Applied Psychology*, 59 (5), 562–71.
- Einhorn, Hillel J. & Hogarth, Robin M. (1975) Unit weighting schemes for decision making. *Organizational Behavior & Human Decision Processes*, 13 (2), 171–92.
- Ericsson, K. Anders (ed) (1996) *The Road to Excellence: The Acquisition of Expert Performance in the Arts and Sciences, Sports, and Games*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Ericsson, K. Anders & Smith, Jacqui (eds) (1991) *Toward a General Theory of Expertise: Prospects and Limits*. New York: Cambridge University Press.
- Fudge, W.K. & Lodish, L.M. (1977) Evaluation of the effectiveness of a model based salesman's planning system by field experimentation. *Interfaces*, 8 (1, part 2), 97–106.
- Gensch et al. (1990)
- Gill, T. Grandon (1995) Early expert systems: where are they now? *MIS Quarterly*, 19 (1), 51–84.
- Green, Paul E. & Krieger, Abba M. (1985) Models and heuristics for product line selection. *Marketing Science*, 4 (1), 1–19.
- (1992) An application of a product positioning model to pharmaceutical products. *Marketing Science*, 11 (2), 117–32.
- Gorry, G.A. & Scott-Morton, M.S. (1971) A framework for management information systems. *Sloan Management Review*, 13 (1), 55–70.
- Hanssens, Dominique M., Parsons, Leonard J. & Schultz, Randall L. (1990) *Market Response Models: Econometric and Time Series Analyses*. Boston: Kluwer.
- Hayes-Roth, Frederick, Waterman, Donald A. & Lenat, Douglas B. (1983) An overview of expert systems. In *Building Expert Systems*, edited by Frederick Hayes-Roth, & Donald A. Waterman & Douglas B. Lenat. Vol. 1. Reading MA: Addison-Wesley Publishing Co., pp. 3–29.
- Hoch, S.J. & Schkade, D.A. (1996) A psychological approach to decision support systems. *Management Science*, 42 (1), 51–64.
- Huber, George P. (1983) Cognitive style as a basis for MIS and DSS designs: much ado about nothing? *Management Science*, 29 (5), 567–79.
- Hunt, Raymond G., Krzystofiak, Frank J., Meindl, James R. & Yousry Abdalla M. (1989) Cognitive style and decision making. *Organizational Behavior & Human Decision Processes*, 44 (3), 436–54.
- Hutchinson, J. Wesley, Kamakura, Wagner A. & Lynch, John G. Jr. (2000) Unobserved heterogeneity as an alternative explanation for 'reversal' effects in behavioral research. *Journal of Consumer Research*, 27 (3), 324–44.
- Huysmans, J.H.B.M. (1970) The effectiveness of the cognitive constraint in implementing operations research proposals. *Management Science*, 17 (1), 92–104.
- Ives, Blake & Olson, Margrethe H. (1984) User involvement and MIS success: a review of research. *Management Science*, 30 (5), 586–604.
- Johnson, Eric J. (1988) Expertise and decision under uncertainty: performance and process. In *The Nature of Expertise*, edited by Michelene T.H. Chi & Robert Glaser. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Keen, Peter G.W. & Scott-Morton, Michael S. (1978) *Decision Support Systems: An Organizational Perspective*. Reading, MA: Addison-Wesley.

- King, David (1993) Intelligent support systems: art, augmentation, and agents. In *Decision Support Systems: Putting Theory into Practice*, edited by Ralph H. Sprague Jr. & Hugh J. Watson 3rd edn. Englewood Cliffs, NJ: Prentice-Hall, pp. 137–59.
- Larreche, Jean C. (1979) Integrative complexity and the use of marketing models. In *The Implementation of Management Science*, edited by Robert L. Doktor, R.L. Schultz & D.P. Slevin. New York: North-Holland Pub. Co.
- Leonard, Nancy H., Scholl, Richard W. & Kowalski, Kellyann Berube (1999) Information processing style and decision making. *Journal of Organizational Behavior*, 20 (3), 407–20.
- Lilien, Gary L., Kotler, Philip & Moorthy, K. Sridhar (1992) *Marketing Models*. Englewood Cliffs, NS: Prentice-Hall.
- Little, John D.C. (1970) Models and managers: the concept of a decision calculus. *Management Science*, 16 (8), 466–85.
- (1979) Decision support systems for marketing managers. *Journal of Marketing*, 43 (3), 9–26.
- Little, John D.C. & Lodish, Leonard M. (1969) A media planning calculus. *Operations Research*, 17 (1), 1–33.
- (1981) Commentary on 'Judgment based marketing decision models'. *Journal of Marketing*, 45 (Fall), 24–9.
- Lodish, Leonard M. (1971) CALLPLAN: an interactive salesman's call planning system. *Management Science*, 18 (December), 25–40.
- (1974) A vaguely right approach to sales force allocation decisions. *Harvard Business Review*, 52 (January–February), 119–24.
- Lodish, Leonard M. (1980)
- Lodish, Leonard M. & Reibstein, David J. (1986) New gold mines and minefields in market research. *Harvard Business Review*, 64 (1), 168–82.
- Lodish, L.M., Curtis, E. et al. (1988) Sales force sizing and deployment using a decision calculus model at syntax Laboratories. *Interfaces*, 18 (1), 5–20.
- Lusk, Edward J. & Kernick, Michael (1979) The effect of cognitive style and report format on task performance: the MIS design consequences. *Management Science*, 25 (8), 787–98.
- McIntyre, Shelby H. (1982) An experimental study of the impact of judgment-based marketing models. *Management Science*, 28 (1), 17–33.
- Meehl, Paul E. (1954) *Clinical vs. Statistical Prediction: A Theoretical Analysis and A Review of The Evidence*. Minneapolis: University of Minnesota Press.
- Montgomery, D., Silk, Alvin J. & Zaragoza, C. (1971) A multiple product sales force allocation model. *Management Science*, 18 (December), 3–24.
- Montoya-Weiss, Mitzi & Calantone, Roger J. (1999) Development and implementation of a segment selection procedure for industrial product markets. *Marketing Science*, 18 (3), 373–95.
- Naylor, Chris (1983) *Build Your Own Expert System*. Wilmslow, Cheshire: Sigma Technical Press.
- Nisbett, Richard E. & Wilson, Timothy D. (1977) Telling more than we can know: verbal reports on mental processes. *Psychological Review*, 84 (3), 231–59.
- Northcroft, M.A. & Neale, G.B. (1987) Experts, amateurs and real-estate: an anchoring and adjust perspective in property pricing decisions. *Organizational Behavior & Human Decision Processes*, 39, 84–97.
- Oskamp, Stuart (1962) The relation of clinical experience and training methods to several criteria of clinical prediction. *Psychological Monographs*, 76 (28, Whole No. 547), 27.
- Phelps, Ruth H. & Shanteau, James (1978) Livestock judges: how much information can an expert use? *Organizational Behavior & Human Decision Processes*, 21 (2), 209–19.
- Progressive Grocer (1994) Case studies: Hormel Foods Corp. turning data into useful information. *PROMO: Progressive Grocer Special Report*: S14.
- Ram, Sudha & Ram, Sundaresan (1988) INNOVATOR: an expert system for new product launch decisions. *Applied Artificial Intelligence*, 2, 129–48.
- (1996) Validation of expert systems for innovation management: issues, methodology, and empirical assessment. *Journal of Product Innovation Management*, 13 (1), 54–68.
- Ram, Sundaresan & Ram, Sudha (1989) Expert systems: an emerging technology for selecting new product winners. *The Journal of Product Innovation Management*, 6 (2), 89–99.
- Ramaprasad, Arkalgud (1987) Cognitive process as a basis for MIS and DSS design. *Management Science*, 33 (2), 139–48.
- Rangaswamy, Arvind, Burke, Raymond R., Wind, Jerry & Eliashberg, Jehoshua (1987) Expert systems for marketing. Working Paper, Marketing Science Institute, Report No. 87–107 (November).
- Rangaswamy, Arvind, Eliashberg, Jehoshua, Burke, Raymond R. & Wind, Jerry (1989) Developing marketing expert systems: an application to international negotiations. *Journal of Marketing*, 53 (4), 24–39.
- Russell, Stuart J. & Norvig, Peter (1995) *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice-Hall, Inc.
- Schumann, M., Gongla, Patricia, Lee, Kyoung-Sang & Sakamoto, J. Gene (1989) Business Strategy Advisor: an expert systems implementation. *Journal of Information Sciences Management*, 6 (2), 16–24.
- Shanteau, James (1987) Psychological characteristics of expert decision makers. In *Expert Judgment and Expert Systems*, edited by Jeryl L. Mumpower, Ortwin Renn, Lawrence D. Phillips & V.R.R. Uppuluri. Berlin: Springer-Verlag.
- (1988) Psychological characteristics and strategies of expert decision makers. *Acta Psychologica*, 68 (1–3), 203–15.
- (1992a) Competence in experts: the role of task characteristics. *Organizational Behavior & Human Decision Processes*, 53 (2), 252–66.
- (1992b) The psychology of experts: an alternative view. In *Expertise and Decision Support*, edited by George Wright & Fergus Bolger. New York: Plenum Press.
- Sharda, Ramesh, Barr, Steve H. & McDonnell, James C. (1988) Decision support system effectiveness: a review

- and an empirical test. *Management Science*, 34 (2), 139–59.
- Shiel, B. (1987) Thinking about artificial intelligence. *Harvard Business Review*, 65 (July–August), 91–7.
- Silk, Alvin J. & Urban, Glen L. (1978) Pre-test-market evaluation of new packaged goods: a model and measurement methodology. *Journal of Marketing Research*, 15 (2), 171–91.
- Silva-Risso, J.M., Bucklin, R.E. et al. (1999) A decision support system for planning manufacturers' sales promotion calendars. *Marketing Science*, 18 (3), 274–300.
- Simon, Herbert A. (1977) *The New Science of Management Decision*, rev. edn. Englewood Cliffs, NJ: Prentice-Hall.
- Spence, M.T. & Brucks, M. (1997) The moderating effects of problem characteristics on experts' and novices' judgements. *Journal of Marketing Research*, 34 (2), 233–47.
- Sprague, Ralph H. Jr. & Watson, Hugh J. (1993) *Decision Support Systems: Putting Theory into Practice*, 3rd edn. Englewood Cliffs, NJ: Prentice-Hall.
- Sterman, John D. (1987) Testing behavioral simulation models by direct experiment. *Management Science*, 33 (12), 1572–92.
- (1989) Modeling managerial behavior: misperceptions of feedback in a dynamic decision making experiment. *Management Science*, 35 (3), 321–39.
- Tam, Kar Yan & Kiang, Melody Y. (1992) Managerial applications of neural networks: the case of bank failure predictions. *Management Science*, 38 (7), 926–47.
- Turban, Efraim (1992) *Expert Systems and Applied Artificial Intelligence*. New York: Macmillan Publishing.
- (1995) *Decision Support and Expert Systems: Management Support Systems*, 4th edn. Englewood Cliffs, NJ: Prentice-Hall.
- Turban, Efraim & Aronson, Jay E. (1998) *Decision Support Systems and Intelligent Systems*. Upper Saddle River, NJ: Prentice-Hall Inc.
- Tyran, Craig K. & George, Joey F. (1993) The implementation of expert systems: a survey of successful implementations. *Data Base*, 24 (1), 5–15.
- Urban, G.L. & R. Karash (1971) "Evolutionary Model Building in the Analysis of New Products." *Journal of Marketing Research* 8 (February): 62–66.
- Urban, Glen L. & Katz, Gerald M. (1983) Pre-test-market models: validation and managerial implications. *Journal of Marketing Research*, 20 (3), 221–35.
- van Bruggen, Gerrit H., Smidts, Ale & Wierenga, Berend (1998) Improving decision making by means of a marketing decision support system. *Management Science*, 44 (5), 645–58.
- Wierenga, Berend (1990) *The First Generation of Marketing Expert Systems*. Philadelphia, PA: Wharton School, University of Pennsylvania.
- Wierenga, Berend & Ophuis, Peter A.M. Oude (1997) Marketing decision support systems: adoption, use, and satisfaction. *International Journal of Research in Marketing*, 14 (3), 275–90.
- Wierenga, Berend & van Bruggen, Gerrit H. (1997) The integration of marketing problem-solving modes and marketing management support systems. *Journal of Marketing*, 61 (3), 21–37.
- Wierenga, Berend, van Bruggen, Gerrit H. & Staelin Richard (1999) The success of marketing management support systems. *Marketing Science*, 18 (3), 196–207.
- Wilson, Rick L. & Sharda, Ramesh (1994) Bankruptcy prediction using neural networks. *Decision Support Systems*, 11 (5), 545–58.
- Wind, Jerry, Green, Paul E., Shifflet, Douglas & Scarbrough, Marsha (1989) Courtyard by Marriott: designing a hotel facility with consumer-based marketing models. *Interfaces*, 19 (1), 25–48.
- Wright, George & Bolger, Fergus (1992) *Expertise and Decision Support*. New York: Plenum Press.
- Zinkhan, George M., Joachimsthaler, Erich A. & Kinnear, Thomas C. (1987) Individual differences and marketing decision support system usage and satisfaction. *Journal of Marketing Research*, 24 (2), 208–14.
- Zmud, Robert W. (1979) Individual differences and MIS success: a review of the empirical literature. *Management Science*, 25 (10), 966.