Corporate Corner

Expert Systems for Evaluating Business Opportunities: Implementing the Management Advisor at Krypton Chemical

W. Reitman and S.J. Shim
Decision Sciences and Engineering Systems Department, School of Management, RPI, Troy, NY 12180, USA

ABSTRACT **This paper investigates issues arising in procuring and implementing a commercial high-end expert system, Palladian's Management Advisor. This system is intended for improving corporate decisions about financial investments in new business opportunities, e.g. projects involving new plants, equipment, products, or marketing expenditures. We look at Krypton Chemical's experiences with the Management Advisor from both the customer's and the vendor's viewpoints. The discussion raises some difficult questions about what it takes to make high-end financial expert systems successful. The paper then concludes with recommendations intended to improve design, marketing, procurement and implementation practices with respect to such systems.**

INTRODUCTION

This paper analyzes a case involving the procurement and implementation of a commercial high-end expert system product intended for multiple business users and organizations. The generic products developed by Applied Expert Systems (Stansfield and Greenfeld, 1987), Palladian Software (Bitran and Papageorge, 1987; Reitman, 1990) and Syntelligence (Duda *et al.*, 1987) are examples of such systems. Although these systems all have different domains of application, they face the same issue: how to apply effectively a common body of generic domain knowledge across a broad range of user problems.

Those involved in procuring and implementing high-end commercial expert systems face a number of problems largely specific to this class of software products. Like other expert systems, they are, of course, a relatively new technology, and may pose acceptance problems in organizations unfamiliar with them. But high-end commercial expert systems also differ from other expert systems in important ways. They typically are expensive: a single system may cost tens or even hundreds of thousands of dollars. Typically, also, high-end expert systems have a relatively broad organizational scope. For these reasons as well as others, procuring and implementing a high-end commercial expert system may create organizational problems and call for levels of organizational skill and commitment well beyond anything required for individual purchases of small expert systems or shells.

In studying high-end expert system procurement and implementation it would seem desir-

1055-615X/93/030191-14$12.00

Received 11 March 1992
Revised 6 October 1992

able to build up a general picture of the process by examining at a given instance from both the vendor's and the customer's viewpoints. In this paper we have attempted to approximate this ideal by taking an existing case study (Hertenstein and Addonizio, 1989a,b) based upon interviews with the customer, and complementing it with an analysis based upon interviews with the vendor. Hertenstein and Addonizio's study of Krypton Chemical Company's experience with Palladian's Management Advisor is one of the Harvard Business School's case studies in accounting and control. However, it also includes extensive and thoughtful discussions of the technical and organizational issues involved in Krypton's attempt to make use of the Advisor, and so it will serve our purpose very well.

THE MANAGEMENT ADVISOR[1]

The Management Advisor (MA) is a generic, knowledge-based application package designed to improve corporate decisions about new business investments in plants, equipment, products, etc. It is intended for executives, managers and other business end-users. In addition to financial value, the system looks at risk, timing, competition and overall business impact. For each analysis, the MA also manages the relations among prices, market share, costs, depreciation, taxes, and the numerous other variables in a business decision that affect a company's bottom line.

Adapting the MA to the User

To match the MA's generic expert knowledge about business to the needs of a specific user and problem, the system includes a tailoring module. Basically, this adapts the system to users' preferences, financial assumptions and terminology. Users can deviate from the default corporate assumptions, but all such deviations are flagged and must be justified. In sum, the tailoring options are intended to enable users to employ their usual ways of thinking about their individual business valuation problems, while at the same time ensuring a coherent analytic framework across the organization as a whole.

At the most abstract level, the MA is designed to work with discounted cash flows. This is what enables it to evaluate business opportunities of all sorts. Because organizations vary in their preferred valuation methods, the MA has the ability to present its results in other terms as well, including internal rate of return, undiscounted and discounted payback periods, profitability index, average returns on investment and assets, and average return on sales. Thus, users get the benefits both of a single coherent conceptual framework and of a wide choice of measures matching their own ways of thinking about business opportunity valuation.

Working with the MA

Figure 1 shows the normal analytic flow of work through the MA. For a more intuitive understanding of why the MA does what it does, however, it is helpful to look at an example of how the expert upon which it is modeled works.

An expert may be called in when a company is considering a major strategic move such as developing a new product, and wants to be sure that the evaluation of the proposal is as good as it can be. Typically, the people in charge of the evaluation are perfectly capable of carrying out the mechanics of a net present value computation. What they are paying the expert for is help in improving their problem formulation and in applying the formal analytic methods in the most appropriate fashion.

With this in mind, the expert sets about analyzing the initial formulation. The expert has several groups of questions to ask the managers and analysts. For example, has the company taken adequate account of the competition? After all, if initial evaluation indicates that the project is going to be highly profitable, sooner or later the competition will pick that up and respond. What kind of response can the competition make? When can they make that response? How will that response affect the profitability of the project? As a result of these questions and the discussions they generate, the initial problem statement is iteratively refined and redefined, thereby sharpening

1 The Management Advisor was on the market for about three years, from early 1986 through late 1988. We have not been able to determine whether it is still in use anywhere.

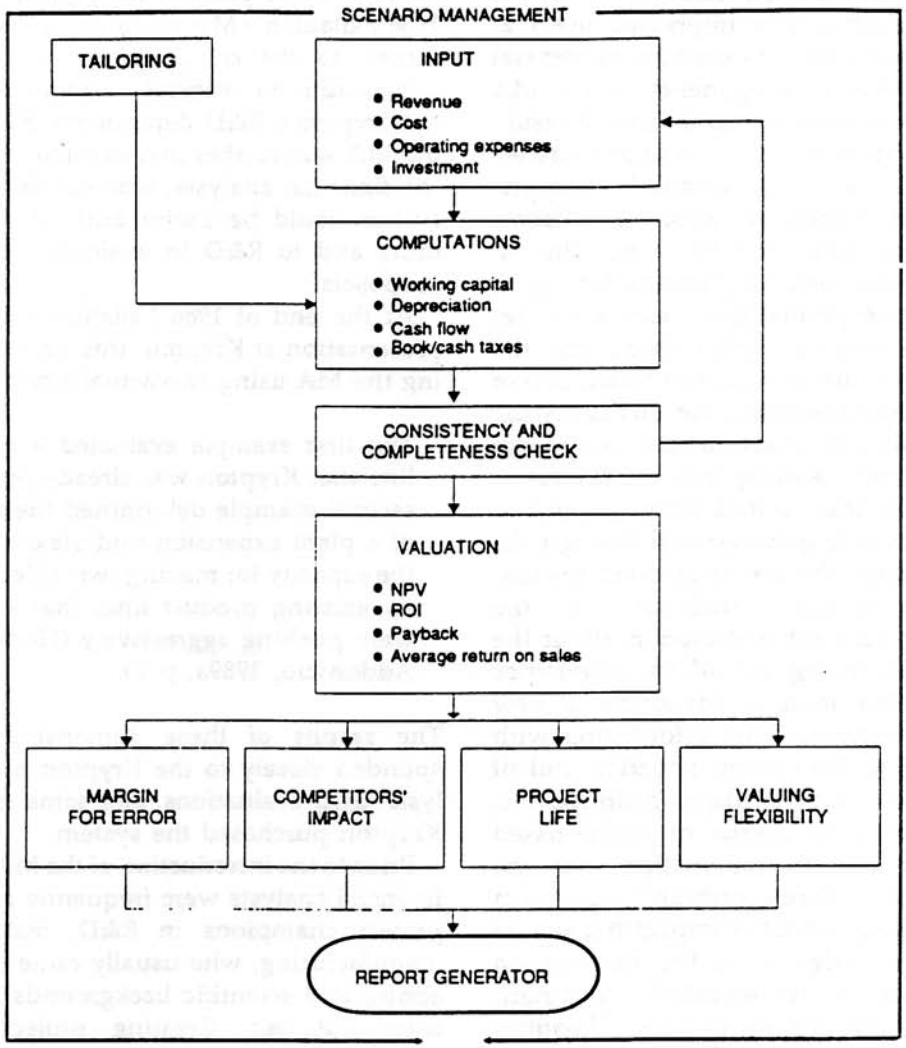

Figure 1 The Management Advisor™

up and improving the project evaluation. As we will see, the MA's Competitors' Impact module goes through much the same process with the user of the MA.

In summary, the MA proceeds from the assumption that financial value and valuation are widely useful concepts, and it attempts to map the generic knowledge required for ascertaining value onto a broad range of specific business applications.

Technical Basis of the MA

The MA is a composite expert system implemented in the Palladian Software Language (PSL), a LISP-based artificial intelligence language specifically designed for efficient implementation of management expert systems. PSL utilizes a number of artificial intelligence representation and control concepts, including contingent inheritance hierarchies, object-oriented graphics, production rules, model-based reasoning, goal-directed control and constraint-based inference (Steele, 1980). The MA consists of about 200,000 lines of code, and runs on LISP machines and high-end workstations. At the time Krypton purchased the system, it cost about $65,000 for the software, plus $50,000 for the specialized workstation.

We mentioned earlier using the MA to assess the likely effect of competitive entry upon the expected value of a proposed project. The

MA's Competitors' Impact module, one of the modules that particularly impressed users at Krypton, is a good place to see how the several artificial intelligence components in the MA work together to achieve a co-ordinated result. Conventional 'if–then' rules are used to ascertain situations in which the module is appropriate. These rules operate only after the valuation component has completed its work. The 'if' parts of the rules look for patterns having to do with absolute profitability, relative market share, growth rate of market share, and the like. Should the current situation match one of the entry condition patterns, the user is invited to call the module to determine the likely price effects of competitive entry into the market.

If the user decides to look into competitive impact, that module gains control through the goal system. First, the object-oriented presentation system creates a dialogue with the user to collect relevant information about the situation. Next, using the MA's inheritance mechanisms, the module constructs a new scenario that combines that information with its knowledge of the current situation and of how competitive entry is likely to affect price. This is achieved by means of model-based reasoning, working in conjunction with the constraint system. More specifically, the system uses its model of competitive impact to integrate situational knowledge with the information provided by the user. It then calls the constraint system to work through the network of financial relations, to determine equilibrium price under this particular combination of conditions. Finally, the results are presented to the user in an appropriate graph created by the object-oriented presentation system.

KRYPTON CHEMICAL'S EXPERIENCE WITH THE MANAGEMENT ADVISOR

Krypton Chemical[2] is a large producer of diverse chemicals and plastics products. In 1983, Krypton introduced a business unit structure to manage its 30 product categories. In February 1986 Palladian's Management Advisor was presented to the company in the course of a symposium on artificial intelligence sponsored by Krypton's R&D department. Subsequently, the MA was further investigated by two Krypton financial analysts, who concluded that the system could be useful both to the business units and to R&D in evaluating new project proposals.

At the end of 1986 Palladian made another presentation at Krypton, this time demonstrating the MA using two actual Krypton projects:

> The first example evaluated a new product line that Krypton was already pursuing. The second example determined the profitability of a plant expansion undertaken to increase the capacity for making two-sided adhesives, an existing product line, that some people were pushing aggressively (Hertenstein and Addonizio, 1989a, p. 7).

The results of these demonstrations corresponded closely to the Krypton financial analysts' own evaluations, and some months later Krypton purchased the system.

Prior to the introduction of the MA, Krypton's financial analysts were frequently at odds with project champions in R&D, marketing and manufacturing, who usually came from engineering and scientific backgrounds rather than from business. Creating project proposals involved substantial time, money and effort. Yet nine out of ten proposals were rejected by the financial analysts, who felt that the proposals they received typically focused on technical merit and slighted business and financial considerations.

The financial analysts, in turn, were regarded by much of the rest of the company as technically ill-trained number crunchers. Part of the reason the financial analysts recommended purchasing the MA was that they hoped it would find use as a training and screening device. This would improve the financial analyses in the proposals submitted to them, reducing the resources the business units and R&D invested in preparing financially unjustified plans, and allowing the financial analysts to improve their relations with other parts of the company.

2 This description of the Krypton Chemical Company and its experience with the Management Advisor is based upon Hertenstein and Addonizio (1989a,b). 'Krypton Chemical' is the pseudonym used by these authors.

Krypton's experience with the MA turned out to be a disappointment both to the company and to Palladian. Hertenstein and Addonizio, looking at the story from Krypton's perspective, point out several reasons for the outcome.

Technical Issues: Interfacing, Flexibility and Modifiability

The MA lacked good spreadsheet capabilities, and could not interface to spreadsheets and databases on other computers. In addition, it could not readily accommodate all the variables, e.g. 'bolts per machine', that R&D normally used in its analyses. Thus, using the MA required recalculating, consolidating and re-entering data, a considerable inconvenience at best. At worst, one research chemical engineer claimed, it required inputs such as manufacturing cost that could take a week to develop.

Logistics and Ownership

When Krypton bought the MA, they had the system installed at company headquarters, one floor below the financial analysts' offices, but half a mile from manufacturing and some two miles from R&D's facilities. Six months after installation, no one from R&D had yet used the system. According to the R&D director, there was:

> a feeling among some people here that anything in another area is owned by that area. In this case, it is finance. Using the Management Advisor over there would be like baking a cake in someone else's kitchen (Hertenstein and Addonizio, 1989a, p. 9).

Disagreements about the Role of Financial Analysis at Krypton

Some project champions felt that, used to the exclusion of other considerations, the kinds of financial analyses the MA carried out were a narrow and inappropriate way to evaluate the long-term merits of new R&D ideas (a point the expert the MA had been modelled upon would have agreed with). Others were concerned that the success rate for Krypton's project proposals could drop if they started using the MA to screen out projects at an early stage purely on financial grounds. Yet that, of course, was precisely what the financial analysts wanted to do.

Disagreements about Who Should Use the Management Advisor

Finance, as we have seen, hoped the system would be used by business unit managers and R&D people as well as by themselves. They considered it a 'user-friendly financial expert system that presented information in a straightforward manner' (Hertenstein and Addonizio, 1989a, p. 7). Krypton's controller agreed.

> The software is friendly enough for people in the business units to come in, sit down, and structure their proposals ... Managers can learn to ask the right questions, and develop the mindset to think naturally about economic issues. (Hertenstein and Addonizio, 1989, p. 8).

As it turned out, the Director of R&D had reservations about the appropriateness of the MA for his people. He suggested the system might be suitable for middle-level business unit managers to use. Other people in R&D also thought there might be uses for the system, but the users again would be some place else. One, for example, thought the MA would be well suited for someone at the Chief Executive Officer level. In the end, only one of the business units ever used the product on a serious basis.

Problems in Introducing the System

We noted earlier that, at Krypton's request, a key Palladian demonstration of the Management Advisor had been built around the issue of expanding capacity for making two-sided adhesives. In retrospect, even in the eyes of the manager of financial analysis, this seems to have been a particularly poor choice:

> We invited the division supervisors and all those above that level to attend ... Representatives came from manufacturing,

marketing, and research and development ... There were some hard feelings between the guy from R&D who pushed the proposal and the financial analysts in our department who evaluated it. The results of Management Advisor confirmed the financial analysts' results, exactly. The project died that afternoon; there still are some bad feelings about it (Hertenstein and Addonizio, 1989a, p. 7).

Organizational Issues

Having detailed the technical, logistic and positioning issues contributing to Krypton's lack of success with the MA, Hertenstein and Addonizio also underscore some more fundamental organizational problems. We have already noted the results of the two-sided adhesives demonstration, which Hertenstein and Addonizio (1989b) describe as a public humiliation, and an organizational behavior nightmare.

> This demonstration convinced managers that Management Advisor provided yet another way for proposals to be rejected by the financial analysts ... Public rejection only fueled their humiliation. Instead of ameliorating hostility ... , the analysts alienated the managers further (p. 7).

Additionally, as we have seen, there appears to have been a significant failure of communication within the company. Financial analysts assumed the system would be used by business unit and R&D managers, and evidently were not aware that those managers preferred to leave financial analysis to the analysts: 'Since past efforts to teach financial analysis concepts to the managers met with little success, simply making a new method available did little to change their underlying reluctance to perform analysis'. (Hertenstein and Addonizio, 1989b, p. 6).

Hertenstein and Addonizio conclude that for the Management Advisor to have been successful at Krypton, at least two conditions would have to have been met: the company would first have had to have reached a consensus regarding the relative roles of technical and business strategy, on the one hand, and financial analysis, on the other, in making decisions about new projects. It would also have needed to have found better organizational forms for financial analysis, perhaps by integrating analysts into the business units in order to reduce conflict and to encourage a more positive role in structuring and developing acceptable proposals.

PALLADIAN'S PERSPECTIVE

Having looked at Krypton's experience with the Management Advisor from the viewpoint of the user organization, let's see what more we can learn by augmenting the picture from Palladian's perspective.

Palladian staff gave an account of the basic history of the Management Advisor at Krypton that was substantially similar to what Hertenstein and Addonizio found out at Krypton. But there also are some interesting differences. For example, Susan Berner, a member of Palladian's MA sales and product management group, recalls that at the beginning of the sales process Phil Cooper, Palladian's Chief Executive Officer, played a significant role in getting Krypton to sign up with Palladian. According to Berner, Cooper helped develop Krypton's initial interest in the Management Advisor through top-level contacts with Krypton's management.

Cooper played a similar role in generating a good number of Palladian's initial sales. Cooper's background was in marketing, and he was a superb salesman. Furthermore, he also was the originator of the MA concept. Thus he thoroughly understood its appeal to top level management, and he carefully positioned himself to create and develop contacts with senior executives of potential client companies in order to build their interest in the product.

We noted in the introductory discussion the financial costs and the levels of organizational involvement that procuring and implementing high end expert systems typically entail. Thus it will be understandable that during Palladian's initial efforts to introduce the MA, having a senior management champion in the client company often proved critical to securing a

decision to purchase and try out the new technology. This seems to have also been the case at Krypton Chemical.

Technical Issues: Absence of Hardware and Software Interfaces

Palladian was aware of the user reactions described above to the technical limitations of the MA as initially released. Susan Berner, Linda Tyler and Cindy Southerly, each of whom had participated in Palladian's sales and implementation effort at Krypton, all described in detail the distress and inconvenience Krypton MA users experienced at not being able to interface with the spreadsheets and databases they ran from their network of PCs.

The initial release also fell well short of the MA Cooper had envisaged in other important respects. In designing large systems based upon rapidly developing technology it is necessary to make assumptions about what the state of that technology will be at the time the system is ready for delivery. The original MA concept, developed in 1984, had called for a multi-user expert system implemented on IBM PCs networked to an IBM mainframe. As a backup plan, in the event suitable IBM hardware and software did not become available soon enough, Palladian was prepared to deliver the MA on high-performance AI workstations.

Well before the target date for delivering the MA it became clear that the most advanced IBM PC then available, the IBM 286, would not be able to handle the MA's technical requirements. In particular, the MA would require a computer with virtual memory and a sophisticated LISP. At that time, the only prospective delivery vehicles satisfying those requirements were those special-purpose AI workstations. Unfortunately, there was neither software nor hardware available that would interface such workstations with user databases and spreadsheets, or network them with IBM mainframes or PCs. Accordingly, while continuing to hold to the original concept as its long-term plan, Palladian decided to make initial deliveries of the MA on the AI workstations on which it had been developed.

As we have seen, the MA was originally designed to allow multiple users to pass project proposal scenarios back and forth across networked individual workstations. Users were to have the capability of making quick and easy modifications in their copies of one another's proposals, thus allowing individuals in one corporate area to explore and react to the impact of particular assumptions in projects prepared by other areas. This would make it easy to assess the consequences of replacing those assumptions with others. The MA was also designed to maintain a tree of proposal derivations in the mainframe integrated into the network of user workstations. As a result, anyone who wanted to would be able to trace the history of this overall corporate evaluation to see whether all the relevant alternatives had been examined. In many cases it was precisely this vision of the MA's intended long-run organizational and strategic impact that attracted the senior managers at Palladian's potential customers to the product in the first place.

All these capabilities had been programmed into the system as part of the core MA concept. All would have to be usable in order for the MA to serve its intended function as a tool for a corporate-wide dialogue, a dialogue that would evaluate a proposal in depth, and insure that all involved parties, including top corporate management, had an opportunity to shape it to the corporation's overall best advantage. Thus, when the decision was made to deliver initial versions of the MA as a single user stand-alone product, it meant that much of the capability for corporate-wide analysis and evaluation of proposals that had been built into the product would not be usable in the first release.

Palladian was well aware that releasing the MA on isolated individual workstations would entail major drawbacks. Users would have to do substantial recalculating, consolidating and re-entering of data. Furthermore, only those capabilities that could be exercised by an individual user were available to the customer. Nonetheless, if the MA was to go out on the schedule Palladian had committed to with its financial backers, there was no choice but to deliver on isolated workstations. In doing this, Palladian counted on the unique capabilities of the product to overcome the user resistance created by these difficulties. As the reaction to the product at Krypton indicates, however, that

assumption in most cases proved to be overly optimistic.

Technical Issues: Flexibility and Modifiability

As we have seen, some Krypton users faulted the MA for what they perceived as its lack of flexibility and modifiability. Palladian was aware of these reactions. However, since users were responding here to fundamental design decisions that had been built into the core of the system, not much could be done to deal with the problem in the short run. There also was no consensus at Palladian about whether a major change in the core was desirable, and the issue was moot in any case. Palladian's resources were stretched thin, and undertaking such a redesign would have entailed major delays in its product development and delivery plans.

The design issues involved here are quite instructive. It should be noted that, except for a few explicit and well-controlled channels, the MA was designed to be a closed system. In other words, if a user encounters a situation the Advisor does not accommodate, there is generally no way he or she can modify the Advisor to address the situation, or even count on the developers' being able to provide what's needed, at least in the short run.

This lack of user modifiability may appear surprising. Aren't expert systems supposed to be modular? When an MA user needs to incorporate some new knowledge, why can't he or she just add a new rule?

If we consider matters from the vendor's viewpoint, we find that there are several reasons for these limits on flexibility:

(1) For reasons of computational power and efficiency, the MA uses a composite AI base, as we have seen; so only some of the knowledge is represented in rules. Much of the rest is represented within the constraint system.

(2) Although early expert system marketing efforts emphasized rule modularity, even in a purely rule-based expert system the situation is more complex than it might appear. Rules certainly are *syntactically* modular. That is, if your new or modified rule is syntactically correct, you can insert the rule without modifying anything else in the system, and the system will *run*. Once you've added your new rules, however, getting the system to do what you *want*, and only what you want, can be another matter entirely. Syntactic modularity doesn't guarantee semantic modularity: you still could have a lot of testing to do to be sure that, taken in combination with the rest of the rule set, the new additions *mean* operationally what you want them to.

(3) A large-scale generic commercial expert system product is not like the small rule-based system you might build for your own use. If the system doesn't work right, and the user can't tell why, he or she assumes the developer is at fault. That being the case, even if it is technically feasible to make the system generally open, most developers will think twice before doing so, because to give users the ability to make unanticipated and unrestricted changes in the knowledge base is to give them the ability to corrupt the system, with unpredictable results. Making a system generally open can be equivalent to making it difficult, if not impossible, for the developer to guarantee robustness and functional correctness. With this in mind, many developers will tend to restrict openness within manageable bounds, e.g. through the use of something analogous to the MA's tailoring module.

(4) There is still another source of restrictions on what the MA user could modify. As we have seen, the MA makes extensive use of constraint-based inference. Because the representation of financial operations in these terms is still a relatively new enterprise, the specific formulations at the base of the Advisor occasionally required extensions in order to capture important features of a given user's financial variables and flows not yet representable in the model. Note that the issue here is not the equivalent of writing a new rule, but of augmenting the basic representation itself—the equivalent, in other words, of modifying or adding to the structure of the language

in which all rules for all users are written. Whether we are representing knowledge in rules or in networks of constraints, restructuring portions of the basic knowledge-representation language clearly is not something a vendor can leave to the user.

Palladian's Perspective on the Organizational Issues at Krypton

Among the most striking results we get from comparing Palladian's perspective on the MA at Krypton with Krypton's own experience are those obtained when we focus upon organizational issues.

Palladian was well aware that getting prospective users to accept a conceptually new product like the MA would not be like selling soap. The market for soap exists. It is well defined. It doesn't need to be educated about the value of soap. If you come up with a new and better soap, the market is likely to be able to understand and evaluate the claimed advantages, and to decide whether they justify their cost. By contrast, at the time the MA appeared, a market for expert systems that could evaluate financially new business opportunities simply did not exist in this well-defined sense. Business organizations, of course, had people evaluating such plans, but these people were not using comparable products or technologies, and they had no ready framework for evaluating and cost-justifying an expert financial evaluation system. Palladian could only guess at the goals, backgrounds and functional orientations of the people who would be using the MA, and, as Krypton's experience illustrates, the companies buying it weren't that sure who the users would be, either.

The MA's product management team had done what it could to anticipate and deal with these problems. Almost a year before the MA was installed, Leslie Watson of the MA product marketing group distributed the first of a series of memos that attempted to characterize functionally the different categories of individuals Palladian people might encounter in the course of dealing with MA prospects and users. These included three types of financial analysts: business unit managers, senior corporate management, strategic planners and advanced technology people. These memos presented background information on the organizational positions, reporting relations, and responsibilities and concerns of each of these functional groups. In addition, they included hints on how best to present and position the MA in talking with individuals in each of these groups.

Understandably, however, there were limits to what such training could do in preparing MA salespeople to analyze and deal with organizational and situational complexities of the sort Palladian encountered at Krypton. Susan Berner, who participated in the two-sided adhesive project demonstration, was grateful to the financial staff for providing real data from a project currently under discussion. To her this seemed a critical factor in getting the high-level managers who attended the demonstration to support the decision to buy the MA. Linda Tyler, of the Palladian sales staff, was of much the same opinion:

> It was like with any other sales situation. By presenting the system using problems or challenges that were real to the customers, Palladian could make them perceive a lot more value in the system. The Palladian people were provided with an excellent forum in which to do that.

The view of the Palladian people is understandable when we realize that all the Krypton personnel Berner and Tyler mention by name were financial analysts. Similarly, those trained by Cindy Southerly, the MA customer support person Palladian sent to Krypton, were described as people in the finance department. Given Palladian's limited contacts at Krypton, it is not surprising that the MA demonstration Hertenstein and Addonizio characterize as a public humiliation and an organizational behavior nightmare for some of the key players at Krypton was taken by the Palladian team to be a great success. Locally, it was: Krypton bought the MA. That's the step the Palladian salespeople would have been focusing upon at that stage. Globally, however, it contributed, as we have seen, to the longer-term failure to implement the MA successfully at Krypton.

In sum, Palladian was aware in principle that market acceptance and successful

implementation of the MA would require a significant change in how prospective customers conducted their new project planning and evaluation, and that this might raise difficult organizational issues. But a sales and implementation team totally prepared to deal with the full spectrum of issues it might encounter in marketing and implementing the MA would have to include people knowledgeable about expert systems technology, financial analysis, the specific industry involved and all the organizational complexities modern business organizations give rise to. As its experience with the MA developed, Palladian increasingly tried to include all these elements on its sales and implementation teams, but its resources were limited, and, ultimately, they counted on what they perceived to be the obvious advantages of the MA to carry the day. Also, if Hertenstein and Addonizio, as quoted above, are correct in their assessment of the situation, perhaps not even such a full court press could have overcome the organizational difficulties between Krypton's financial analysts and the rest of the company.

RECOMMENDATIONS AND DISCUSSION

We have now considered the Krypton–Palladian experience from both customer and vendor viewpoints. This combined perspective demonstrates that commercial expert systems, having made it into the high-end financial information systems arena, are prey to the problems that afflict implementations of conventional MIS (see, for example, Lucas, 1981; Lucas *et al.*, 1990), in addition to those problems more specifically their own (Coats, 1988).

What can we draw from these two accounts that might help to make similar efforts in the future more successful? Two general sets of recommendations seem appropriate. The first deals with buying and implementing such systems, the second with designing, developing and marketing them effectively. Each set of recommendations is followed by a more detailed examination of the issues involved.

(1) Companies wishing to implement financial expert systems that are strategic in intent and corporate-wide in scope must first have in hand an adequate description of the strategic and/or organizational problems they are trying to deal with. That description is a necessary precondition for successfully accomplishing the subsequent steps: developing an explicit strategic interface, and designing and carrying out a corporate-wide implementation plan.

We have seen that Palladian's original vision for the MA had a strong strategic component. We need to understand, however, how this strategic element fits into today's business environments, which are marked by two strong and apparently contradictory trends. On the one hand, there is the effort to improve the quality and increase the extent of top-level management control over a firm's business decisions and business activities. Xerox's implementation of its integrated executive support system (Osborn and Applegate, 1989) is a good example of these efforts to gain better operational control and to improve overall efficiency in allocating corporate resources. On the other hand, as illustrated by recent developments at IBM (Carroll, 1991), we also are witnessing systematic efforts to push key business decision-making responsibilities down to lower levels in the organization.

Where would a high-end expert system implemented according to the original intentions for the Management Advisor have fitted into this picture? Like Xerox's executive support system, this hypothetical version of the MA would certainly have enabled top-level management to exercise closer supervision of the proposal creation and evaluation process. To the extent top management desired to do so, it would have been able to participate more actively in directing the course of that process. Furthermore, thanks to the network interface enabling users to access and modify the proposal derivation tree, this gain in control would have been achievable at little or no cost to the corporation's ability to deal expeditiously with proposals.

On the other hand, this hypothetical MA could also have been used in ways that would have contributed to decentralized decision making, and to pushing decision making down to lower levels in the organization. After all, one primary reason for centralized control is

uncertainty about the judgment or corporate-wide perspective of lower echelons in the organization. The MA would have allowed those echelons to prepare and justify their proposals as they saw fit, but in a medium that made for easy monitoring and review by top-level management to whatever extent it wished to exercise those prerogatives. Thus one can argue that this hypothetical MA could have been used in the interest of either centralized control or decentralized decision making. In both cases, to the extent that the financial expertise built into the product was consistent with a corporation's framework for financial analysis, a full-scale MA implementation of the form described above should have facilitated more expeditious and efficient corporate resource allocations.

There is little in Hertenstein and Addonizio's account to suggest that the financial managers and analysts at Krypton who bought the Management Advisor had considered either these strategic interface issues, or how the rest of the company would react to the introduction of the product. There also is little evidence of any detailed plan for its effective implementation outside the financial analysis group.

Hertenstein and Addonizio (1989b, p. 10) conclude that the company was misperceiving the nature of its business and organizational problems, and hence overestimated the benefits it could derive from the MA. When Krypton bought the Management Advisor, the company's idea was to treat a symptom:

> it failed because it did not address the true source of the difficulties . . . Krypton must develop a companywide strategy and use it as a guide for project selection, especially in research and development. It must also integrate its financial analysts into the business units and insure that they participate in structuring and developing proposals.

Hertenstein and Addonizio's prescription for Krypton is sound advice for any company in a similar situation. If the organizational and strategic analyses required are beyond the resources of a given company, the company may want to secure the services of one of the outside consulting organizations that specialize in this work. One way or another, however, these analyses are a precondition for an effective strategic interface and a realistic corporate-wide implementation plan.

Once a company has in hand adequate analyses of its organizational and strategic objectives, how should it proceed? As yet, we have few (if any) detailed accounts of successful implementations of commercial high-end financial expert systems having corporate-wide strategic implications and scope. There are, however, a number of excellent, detailed and up-to-date case discussions of MIS planning and implementation that may be consulted to good advantage, notably Balaguer and Addonizio's (1991) study of CAD/CAM planning and implementation at Pratt & Whitney, and Garvin and Simpson's (1988) account of Digital Equipment Corporation's introduction of MRPII. The Pratt & Whitney study gives very good descriptions of strategic interface development, and of ways of selecting and testing candidate hardware and software systems so as to be sure that they really do satisfy the strategic requirements that have been decided upon. The Digital study is notable for its superb account of what it takes to plan and achieve an effective implementation of a complex corporate-wide system.

(2) Developing a high-end commercial expert system for a strategic financial application entails both technical and market risks. The more novel the system, the broader its scope, and the greater the discrepancy between the task interactions implied by the system and the way client organizations currently function, the greater the risks. One good way for vendors to cope with these risks is to build close working relations during product design and development with a small, well-chosen set of prospective client companies. These development partners must be willing to subject successive evolving versions of the product to intensive testing in representative working environments. Finally, the vendor must be willing and able to accept an indefinite negative cash flow. The product cannot go on the market until these development partners have understood in detail what the product does, and have modified their

task and organizational arrangements to the point where they have satisfied themselves that they can use the product effectively, and that it meets their needs.

In discussing our first set of recommendations, we were focusing upon a full-scale implementation of a hypothetical version of the Management Advisor—the MA as Palladian had originally envisaged the product. What Krypton actually bought and implemented, of course, was something very different: a single copy of a single-user version running on an isolated AI workstation.

Krypton may be faulted for not having thought through its strategy sufficiently, but it can hardly be blamed for not having evaluated the original-concept MA well enough: there was no implementation of that version to evaluate. Palladian, having no hardware options available that satisfied the requirements of the networked MA, and pressing to meet its revenue growth schedules, had chosen to release the Management Advisor only as a stand-alone product on an unintegrated workstation. It thereby created, as we have seen, substantial interfacing difficulties for its users, and it failed to make good on its vision of a corporate-wide tool that could help with the fundamental problem of finance: achieving the most efficient allocation of an organization's resources.

This same concern with rapid revenue growth also led Palladian to adopt what, in retrospect, proved to be an overly ambitious development and marketing plan. As a result, its sales, product management and customer-support resources, spread thin over several product ideas, and many prospects in many industries, were inadequate to the task of grasping and dealing with the organizational complexities of a company like Krypton.

It's important to recognize that purchase of the hypothetical MA by a company like Krypton would have been very expensive, and full-scale implementation would have required significant restructuring of the company's decision-making process. Such a major investment presumes a corporate management convinced of the product's value and committed to its successful implementation throughout the company. Palladian, limited to isolated single-user delivery vehicles, never achieved such full-scale implementation of the MA anywhere. Thus new prospects had no instance they could inspect to satisfy themselves that the product envisaged by Cooper and his top-level customer contacts would in fact eventually be able to do for the client company what he claimed it would. Consequently, even with top-level management support for the MA, most customers tended to proceed as Krypton did, 'getting their feet wet' by installing a single copy of the product to try out somewhere.

That is the problem, not only for the MA as originally envisaged but for any commercial high-end expert system intended for financial applications that follows the development and marketing pattern we have been examining here. These are products with corporate-wide scope and implications. They require active corporate-wide support if they are to fulfil their promise. Yet if there are no functioning implementations for a prospective customer to ask about and study, the high up-front cost and the uncertainty about functionality and fit effectively preclude the customer's purchasing the product. The customer cannot justify the risks involved in purchasing, installing and providing technical and organizational support for the product on a basis broad enough to exercise and demonstrate its capabilities.

Palladian's aggressive development and marketing schedule was, in effect, a gamble that the company could come up with a product the market would buy without building substantial time into the development schedule for multiple iterations of the basic design. That gamble is understandable in the light of the realities of venture capital funding for such entrepreneurial startup companies. The quicker the projected return on the venture capital investment, the higher the valuation of the company, and the greater the share of the company the entrepreneurs get to keep. If the gamble works out, everyone wins. But these projected returns are illusory if the provisions the company makes for dealing with its technical and marketing risks turn out to be inadequate.

Those developing, marketing or implementing products as novel as a high-end financial expert system must understand that these processes are very different from those involved in developing, marketing or trying out new

instances of existing product categories (e.g., as mentioned earlier, a new soap). People already know what products like soap are for, and what kinds of soap they like. Effective marketing can determine these preferences and likely customer reactions to a new soap beforehand. But nobody could have known how they would feel about the MA until it had been designed, developed and implemented on their desks, and they could begin trying it out on real problems.

That is the reason for our recommendations regarding development partner participation. However good the original idea driving creation of the new product, development of a novel high-level expert system entails a very high degree of technical and market uncertainty. The best way we know of to reduce this uncertainty is the progressive interated interaction around the evolving product that takes place among the designers, developers and prospective users participating in a development partner relationship. In this sense, the development partner relationship is simply a systematic way of learning by doing, a strategy for new-product development made famous by the Japanese, but one that also has been used effectively by successful Western firms (Morone, 1993).

Those who would like to investigate the development partner relationship further may want to consult Konsynski and McGee's (1989) study of Bachman Information Systems. This is an excellent example of the pivotal role a well-implemented development partner relationship can play in the successful creation of a novel software system.

It is worth noting that Palladian also had development partners for the MA. However, it used the concept more as a marketing tool than as an aid to design and development. For example, the MA was released to the public only three months after development partners received the beta version of the product. The same stage in the development of the Bachman product was allowed to go on for about seven months. Three months was barely enough time for the MA's development partners to begin to discover what they would want in a fully satisfactory version of product, let alone for the designers and developers to incorporate these additions and modifications in iterated revisions of the software. It is for this reason that our recommendation stresses the open-endedness of the development partner interaction. Not everything the development partners want will be incorporated into the released product. But if the development partners have not had time to satisfy themselves that key technical and functional uncertainties have been resolved adequately, we ought not be surprised if the product does poorly in the market.

To conclude, the kinds of problems we have been describing in this account of the Krypton–Palladian experience with the Management Advisor are well worth trying to understand and avoid. We believe the recommendations we have presented are likely to be useful in avoiding these problems, and we hope they will prove helpful to those developing or implementing commercial high-end expert systems for financial applications in the future.

Acknowledgements

We acknowledge with thanks the many very helpful suggestions made by the editor and the reviewers of the initial draft of this paper. We also thank Prentice Hall for permission to include in this paper material, in particular Figure 1 and the overview of the Management Advisor, from Reitman (1990).

References

Balaguer, N. S. and Addonizio, M., Pratt & Whitney: CAD/CAM strategy. Harvard Business School case no. 9-191-019, revised 5 June 1991.

Bitran, G. and Papageorge, T., Integration of manufacturing policy and corporate strategy with the aid of expert systems, in Oliff, M. (ed.), *Proceedings of the International Conference on Expert Systems and the Leading Edge in Production Planning and Control*, Addison Wesley–Benjamin/Cummings, Menlo Park, CA, 1987.

Carroll, P. B., IBM announces details of plan to break its business into more autonomous pieces. *Wall Street Journal*, 6 December, 1991, p. B3.

Coats, P., Why expert systems fail, *Financial Management*, 17, No. 3, 1988, pp. 77–86.

Duda, R. O., Hart, P. E., Reboh, R., Reiter, J. and Risch, T., Syntel: using a functional language for financial risk assessment. *IEEE Expert*, 2, No. 3, 1987, pp. 18–31.

Garvin, D. A. and Simpson, J. L., Digital Equipment Corporation: The Endpoint Model (A), Harvard Business School case, No.9-688-059, 1988.

Hertenstein, J. H. and Addonizio, M., Krypton Chemical Company, Harvard Business School case, No. 5-188-030 (revised 5/89), 1989a.

Hertenstein, J. H. and Addonizio, M., Krypton Chemical Company: Teaching Note, No. 5-189-095, 1989b.

Konsynski, B. and McGee, J. V., Bachman Information systems, Harvard Business School case, No. 9-189-027, 1989.

Lucas, H. C., Jr, *Implementation: the Key to Successful Information Systems*, Columbia University Press, New York, 1981.

Lucas, H. C., Jr, Ginzberg, M. J. and Schultz, R. L., *Information Systems Implementation*, Ablex, Norwood, NJ, 1990.

Morone, J. G., *Winning in High Tech Markets*, Harvard Business School Press, Boston, 1993.

Osborn, C. S. and Applegate, L. M., Xerox Corporation: executive support systems, Harvard Business School case, 1989, No. 9-189-134.

Reitman, W., Managing the development of generic expert system products, In DeSalvo, D. A. and Liebowitz, J., *Managing Artificial Intelligence and Expert Systems*, Yourdon, Englewood Cliffs, NJ, 1990.

Stansfield, J. L. and N. R. Greenfeld, N. R., Planpower: a comprehensive financial planner. *IEEE Expert*, 2, No. 3, 1987, pp. 51–60.

Steele, G. L. Jr, *The Definition and Implementation of a Computer Programming Language Based on Constraints*, PhD dissertation, MIT, Cambridge, MA, 1980.