Management of R&D Projects Under Uncertainty: A Multidimensional Approach to Managerial Flexibility

Leonardo P. Santiago and Thomas G. Bifano

Abstract—**In this paper, we describe the practical application of a flexibility-based management approach to new product development, highlighting advantages, and limitations of this methodology. The model is concerned with the resolution of uncertainty over the product development life cycle and deals with technical, market, and cost factors all together. To this end, we consider a real options model, which uses multidimensional decision trees, to assess the development process of a high-technology product, namely, the Adaptive Optics Scanning Laser Ophthalmoscope. Moreover, we show how this project could be managed by estimating its value and determining optimal managerial actions to be taken at each review stage of the new product development process. Finally, we draw conclusions about this model's general utility and particular challenges associated with its use as a product development tool, and emphasize the need to consider a multidimensional model, instead of a single dimensional one.**

Index Terms—**Decision trees, managerial flexibility, new product development, R&D projects, real options, stochastic dynamic programming.**

I. Introduction

IN THIS PAPER, we introduce a multidimensional decision support model and discuss one application of it. The model is concerned with the resolution of uncertainty over the product development life cycle and deals with technical, market, and cost factors. To this end, we show how the development process of a high-technological product could be managed by estimating its value and determining optimal managerial actions to be taken at each project review stage.

We follow up the development process of an *in vivo* scanning laser ophthalmoscope (SLO) with the ability to image a living mouse's retina. This device could potentially be used to study the surface and subsurface microstructures of the retina with extremely high resolution. The high resolution is obtained specially by using microelectromechanical deformable mirrors, adaptive optics technology, which can remove the blurring of *in vivo* images of the retina tissue by measuring and correcting high order aberrations of the eye. The new product development process is being carried out by a partnership involving academic and research institutions, and a high-technology company, which develops and manufactures deformable mirrors used for adaptive optics.

Manuscript received May 1, 2004; revised September 1, 2004. Review of this manuscript was arranged by Department Editor J. K. Pinto.

L. P. Santiago is with the Department of Production Engineering, Federal University of Minas Gerais, Minas Gerais 31270-901, Brazil (e-mail: leonardo santiago@ieee.org).

T. G. Bifano is with the Department of Manufacturing Engineering, Boston University, Boston, MA 02215 USA (e-mail: tgb@bu.edu).

Digital Object Identifier 10.1109/TEM.2005.844465

An intrinsic characteristic of the adaptive optics scanning laser ophthalmoscope (AOSLO) is that a straightforward comparison with previous product development projects could not give enough information to the decision makers. The decision makers need to shed light on nontrivial questions such as: how to evaluate the market payoff (how much should a company charge for this product and what is its market potential), what the required product attributes and features performance of the product are, how to balance extra development investment and expected benefit, or how to establish abandonment criteria for the development process.

Our goal in this paper is to illustrate and discuss the application of the decision support model introduced in [1] and further developed in [2]. The methodology used in this case study aims to provide assistance to decision makers of the company that develops and manufactures the deformable mirror. In other words, the report illustrates how decision makers from this company assessed the project at critical review stages which correspond to key moments of investing in the new product development process. However, in the present report, we do not provide details regarding the company itself, nor illustrate effective or ineffective managerial practices. As mentioned before, our goal is exclusively to discuss one of the first applications of the project management methodology and to highlight lessons obtained from this experience.

The model discussed in this paper focuses primarily on the managerial flexibility during the development process of an R&D project. Its framework falls into the real options approach, which have been proposed for a better valuation of development projects (see, e.g., [3]–[5]), and considers that managers would have the option but not the obligation to act upon the development process. One important contribution of the real options approach to the evaluation of R&D projects is the emphasis it has put on the considerations of flexibility and options available to management (see, e.g., [6]–[10]). In this paper, we consider three managerial options at each review stage, namely, to continue the project, to improve it, and to abandon it.

The value of flexibility is what is known as the difference in value between the active management of a project (i.e., when considering the managerial flexibility) and the passive management of a project which assumes that decision makers would not have the option to interfere in the project throughout its development. As pointed out by [6], the traditional evaluation methods [e.g., return of investment (ROI) and discount cash flow (DCF)] need to be "business-case" adjusted in order to provide the same results as the real options approach, otherwise, these methods would reflect the passive management of a project. Thus, the motivation to compare active and passive management arises. Indeed, the passive management of a project is not a realistic baseline to deal with new product development projects, since decision makers would always have the option to interfere in the project. In spite of that, becoming acquainted of whether the value of the managerial flexibility and/or the overall value of a project increases, when uncertainty increases, is undoubtedly important, and one way to do it is to compare the active management approach with the static one (see, e.g., [2], [10], and [11]).

The multidimensional approach has also been considered in other contexts to deal with managerial flexibility. For instance, [11] evaluate optimal environmental investment decisions considering two controllable variables, namely, the rate of environmental investment and the rate of production. Another multidimensional approach considers several qualitative parameters to assess technology projects and their option value [12]. The value of the model we use is to identify performance objective through expert opinion, capture this information in a robust decision framework, and consider technical, market, and cost factors all together. In addition, by considering a multidimensional approach, decision makers can more easily identify future scenarios as a function of the management actions to be implemented in each controllable dimension, during the project review stages (fact that is not possible to achieve by taking into account more traditional single dimensional techniques such as DCF, ROI, and others).

Although the model considered in this paper is a powerful approach, which represents resolution of uncertainty over the new product development process, generalization to the full R&D life cycle still needs to be investigated. Specifically, the case study we consider deals with later stages of the full R&D process (i.e., we do not focus on the front end stages such as preliminary investigation, concept generation, and others). To deal with front end stages, one still needs to come up with efficient ways to assess the range of controllable variables used in the model (i.e., the project technical performance). We further debate this idea in the last section of this paper.

There are still a number of issues that need to be further investigated in order to successfully use quantitative models to value real options. These problems generally fall into three main categories: finding the right model (the one whose assumptions match those of the project being analyzed), determining the model inputs, and being able to mathematically solve the option pricing algorithm (for more details see, e.g., [13] and [14]).

One alternative to deal with the dilemma of whether or not using the model to the full R&D life cycle is to consider a series of (qualitative) statements that can yield insight into factors that enhance the attractiveness of a project (see, e.g., [12]), instead of trying to quantify performance levels of the project and evaluate the overall value of it. On the other hand, one of the negative aspects of the qualitative assessment of a project, as pointed out by [15], is that it cannot capture interactions among factors, nor multi periods effects that need also to be taken into account during early stages of the research and development process.

Another issue that still needs to be further discussed is whether or not to extend the strict logic of (financial) option theory to the entire organization, as the current state of the art remains best suited for project level financial evaluation [16]. In Sections VI and VII, we provide some discussion in line with these implementation issues.

One cannot conclude that the approach and results presented in this case study will be exactly the same as the ones obtained in other applications of the same decision support methodology since this is just one application of it. Nevertheless, we believe the discussion we provide can be beneficial to academics, as well as practitioners. We also hope this report will encourage future applications of the method and generate future research to improve it. The rest of the paper is organized as follows.

Next, in Section II, we provide some background of the development process. Then, we define the model in Section III, and describe the modeling approach in Section IV. The evaluation results are stated in Section V and discussed in Section VI. Finally, we conclude in Section VII.

II. Background

The ophthalmoscope is a telescope for viewing the inside of the eye [17]. Since its invention by Von Helmoholz in the 1850s [18], improvements in this device have provided a more precise image of the eye. Some examples of these improvements are: 1) the capture of a primitive-quality fundus photograph of the optic nerve by Jackson and Websterwith (1885) and 2) the invention of the reflex ophthalmoscope by Thorner (1899), which was later perfected by Gullstrand (1910) [19].

Ophthalmoscopes can image an object either directly or indirectly. A direct ophthalmoscope is a hand-held instrument used in general medicine that sheds light on the patient's retina. It is made in such a way that the use of the device can place his or her retina in optical conjugation with the patient's. An indirect ophthalmoscope uses a set of lenses to transfer the real image of the eye to a plane, where it can be inspected with a magnifying ocular [17].

A major advance in ophthalmic retinal photography was made with the invention of the scanning laser ophthalmoscope (SLO) by Webb *et al.* [20], [21]. The SLO uses a different approach from the traditional ophthalmoscopes, by allowing a beam of light to sweep over the object, delivering all its energy to a very small spot during a very short time. The energy returned from the spot is detected and synchronously decoded to form an image on an electronic display medium (see, e.g., [17]).

Currently, the AOSLO does not exist in the market nor does a substitute capable of providing a resolution comparable to it. Although this revolutionary product is initially being targeted for research applications (studying mice eyes), it could also be used to image a human eye. The goal for AOSLO is to visualize, initially in research applications, the ora serrata (back of the retina) with such a high resolution that clinicians and researchers can study new vessel formation and cellular-scale physiology. By using this instrument, researchers and physicians will be able to study diseases such as macular degeneration, diabetic retinopathy, and glaucoma at an earlier stage and, consequently, test new techniques that can identify and treat them before they reach an advanced stage (e.g., blindness caused by diabetic retinopathy). Thus, it is expected that this product can further the research and treatment of eye diseases, and offer society a high benefit.

In order to develop an AOSLO, it is necessary to merge another technology into the traditional SLO. Adaptive optics technology uses a microelectromechanical deformable mirror [22] to correct high order aberrations of the retinal image and can potentially provide better resolution.

Recently, some researchers have started to investigate the use of microelectromechanical mirror adaptive optics in the human eye [23], [24]. However, no product that merges these two technologies is currently available on the market for human applications or for research purposes. The design process does not merely consist of merging two technologies, namely, the scanning laser ophthalmoscope and the microelectromechanical mirror. On the contrary, multiple design issues can arise when designing an SLO that uses adaptive optics. For instance, issues such as how to capture and improve the imaging process and how to make a compact/commercial system are not straightforward applications of the current knowhow.

III. Decision Model

The purpose of this section is to present a model, derived from the model introduced in [1] and further developed in [2], and emphasize the specific assumptions that were made to analyze the AOSLO product development environment. The main difference between the model presented in this section and the two previously described is that we consider a multidimensional state variable to characterize the management process, instead of considering scalar variables. Moreover, we extend the notion of expected improvement benefit to a multidimensional vector which does not necessarily have only positive components (we further explain this idea in the development dynamics subsection).

Assume the project is initiated at stage $t = 0$, reviewed at stages $t = 1, \ldots, T-1$, and is launched to the market at stage T.

A. Performance State and Management Options

The state of the project at the review point t is characterized by a value representing an assessment (at review point t) of the product's performance when it is launched into the market. This performance state, when observed/realized, is a deterministic vector. However, before its realization, it involves some uncertainty and is modeled by a random vector $\underline{X}_t$ whose elements represent assessment dimensions of the product development process.

We denote management decision options at stage t by u_t. These may represent continuation, improvement (usually with more then one option of improvement), or abandonment of the project. These decisions are made at instances $t = 0, \ldots, T$.

B. Development Dynamics and Uncertainty

Given an action on the part of management at stage t, the next performance state of the project cannot be determined with complete certainty due to uncertainty in the development process. In addition, there may be external sources of uncertainty. In other words, the performance state of the project at stage $t+1$ depends on the performance state at stage t, $\underline{X}_t$, management decision at stage t, u_t, and development uncertainty during stage t which we will denote by $\underline{\omega}_t$, i.e.,

$$\underline{X}_{t+1} = \phi_t(\underline{X}_t, u_t, \underline{\omega}_t).$$

We assume that the model is additive in the following sense:

$$\underline{X}_{t+1} = \underline{X}_t + \underline{k}_t(\underline{X}_t, u_t) + \underline{\omega}_t.$$

$\underline{k}_t(\underline{X}_t, u_t)$ represents an expected improvement vector that can be achieved if decision u_t is selected at stage t. As can be seen, with the above additivity assumption, uncertainty and expected improvement are added to the state at the current stage to obtain the state of the project at the next stage.

The expected improvement vector represents managerial options considered at each review stage

$$\underline{X}_{t+1} = \begin{cases} \underline{X}_t + \underline{k}(u_t) + \underline{\omega}_t, & \text{if } u_t = \text{Continue or Improve} \\ \text{Stopped}, & \text{if } u_t = \text{Abandon} \end{cases}$$

where $\underline{k}(\text{Continue}) = 0$ and $\underline{k}(\text{Improve})$ corresponds to expected impact on each dimension of the performance vector $\underline{X}$. That is, if the project is funded at the "continue" level, the performance at the next stage is expected to be the same as that at the current stage plus some uncertainty. On the other hand, if the project is funded at the "improve" level, there is some improvement added to the vector $\underline{X}$ plus the uncertainty effect.

It should be noted that the expected improvement in a dimension i of vector $\underline{X}$ could impact the other dimensions j $(j \neq i)$ in three different ways: 1) it might deteriorate j and, in this case, some scalar would be subtracted from it; 2) it might not affect j, which is modeled by adding the scalar 0 (zero) to it; or 3) it might improve j which, in this case, would be accomplished by adding some positive scalar to it.

Stopped represents the stopped state for the project. As it is common in modeling practice, we assume that once the project is stopped, it remains in the stopped state in the following stages.

The effect of development uncertainty is assumed to be as follows: We consider a general distribution for the development uncertainty, i.e., $\underline{\omega}_t$. Generally, we assume that the random vector $\underline{\omega}_t$ represents the effect of "pure" uncertainty during stage t. Therefore, we assume it has mean zero, i.e., $E[\underline{\omega}_t] = 0$. In some cases, we also assume that the development uncertainty is symmetric, in the sense that $\underline{\omega}_t =^d -\underline{\omega}_t$ for $t = 0, \cdots, T-1$, where $=^d$ denotes equality in distribution. This assumption is not very restrictive.

C. Development Cost and Market Payoff

The development cost at each stage depends on management decisions. If the abandon option is selected, the project is immediately terminated, and consequently no further development cost will be incurred. Under other options, the cost incurred generally depends on: 1) the state of the project at review point t; 2) the action taken by management at stage t; and 3) the development stage t. We denote this cost by $c_t(\underline{x}_t, u_t)$. An initial investment I is needed to begin the project at stage $t = 0$. Upon completion of the project, the product is launched to the market at stage T. If the performance state of the product is $\underline{x}$, it acquires a random market payoff, denoted by $\Pi(\underline{x})$.

The development costs are assumed to be as follows. Under the continue option, development proceeds to the next stage and a continuation cost $c(t)$ is incurred. If management decides to improve the project, there will be an improvement cost of $\alpha(t)$ in addition to the continuation cost $c(t)$. Also, we assume there are multiple possibilities of improving the project. In summary, the cost at each stage is $c_t(\underline{x}_t, u_t) = c(u_t)$

$$c(u_t)=\begin{cases} 0, & \text{if } u_t = \text{abandon} \\ c(t), & \text{if } u_t = \text{continue} \\ c(t)+\alpha(t, \text{option } h), & \text{if } u_t = \text{improve option } h \end{cases}.$$

To model the market payoff, it is assumed that the market requires a certain level of performance denoted by $\underline{D}$. If the product meets or exceeds customer requirements, the market will yield a high payoff M. In this case, the product will have a competitive advantage when compared with those of the competitors. On the other hand, if customer requirements are not met, the company will receive a lower payoff, m. The required level of product performance $\underline{D}$ is not known before marketing the product. The company's information about $\underline{D}$ is modeled as a distribution with mean $\underline{\mu}$ and covariance matrix Σ. Under these assumptions the market payoff is given by

$$\Pi(\underline{x}) = \begin{cases} M, & \text{with prob. } F(\underline{x}) \\ m, & \text{with prob. } 1-F(\underline{x}) \end{cases}$$

where $F(\underline{x}) = P(\underline{D} \leq \underline{x})$.

D. Basic Tradeoff at a Review Point

The essential tradeoff that is the basis for a management decision at a review point is captured by the dynamic programming equation

$$V_t(\underline{x})=\max_{u_t}\left\{-c_t(u_t)+\frac{1}{1+r}E\left[V_{t+1}\left(\underline{X}_{t+1}(\underline{x},u_t,\underline{\omega}_t)\right)\right]\right\}. \quad (1)$$

In other words, the cost of taking an action during stage t is compared with the discounted expected payoff to be "obtained" at the completion of stage t, and the action that provides the best tradeoff is selected. At the terminal stage $t = T$

$$V_T(\underline{x}) = E\left[\Pi(\underline{x})\right].$$

$V_t(\underline{x})$ represents the expected discounted value of the project (evaluated at the beginning of state t assuming that the performance state of the project is $\underline{x}$). V_t is called the value function of the project.

IV. AOSLO Product Development Management

In this section, we show how the decision model described in the previous section is used to assess the development process of the AOSLO. However, we do not provide extensive details regarding each step of the development process, since our main goal is to illustrate the application of the decision model. Considering the fact that this is a novel product, the company did not have enough information to both assess the market and the product development process.

The first step taken in the development process was to gather information from potential users of the product. To this end, we carried on in-depth interviews with venture capitalists specialized in the ophthalmic industry, ophthalmologists, and leading researchers (senior scientists and/or principal investigators of research grants). The objective of the interviews was to understand: 1) the (ophthalmic) industry; 2) the research environment; and 3) the researchers needs, hopes, and expectations for a product like the AOSLO.

Through the interviews with venture capitalists, we were able to better understand the market potential for a scanning laser ophthalmoscope that targets markets for both applications: for the human eye and research-oriented projects. The interviews with ophthalmologists made some product requirements more explicit, and also improved our understanding of the eye and the environment we were dealing with.

To capture the market needs, one on one interviews were conducted with nine researchers from the following institutions: Harvard Medical School, Massachusetts General Hospital, Wellman Laboratories of Photomedicine, Schepens Eye Research Institute, Joslin Diabetes Center, Wilmer Ophthalmological Institute (Johns Hopkins Hospital), New York Eye and Ear Infirmary, and Ohio State University. Given the fact that an *in vivo* scanning laser ophthalmoscope with the ability to see inside a living mouse's eye is a very specific product, which targets specific users, we consider these interviewees a representative group of the potential market/users.

By means of these interviews, we were able to collect information regarding product features (e.g., resolution, contrast, zoom, portability, conditions of use, among others) and the importance of each of them. In addition to the researchers' needs, we collected data concerning technical specifications (e.g., lateral and axial resolution, zoom, and weight) and information regarding price range for this potential product.

Next, we describe how the expected payoff, the controllable variables of the project, the possible managerial actions were determined, and the tradeoff that the decision makers would be facing at each review stage was captured.

A. Market Potential

1) Pricing Range: When developing a revolutionary product, it is critical for the product development team to have a broad view of the pricing possibilities [25]. One way is to analyze the new product supply chain and obtain a general sense of the price possibilities (see, e.g., [26] and [27]). Another alternative is to establish a ceiling (maximum) and a floor (minimum) price for the future product (see, e.g., [25] and [28]).

Price Ceiling: Decision makers can estimate this maximum price by considering the product's benefits and ensuring that each and every potential price point is brought up for discussion.

By means of in-depth interviews, it was possible to estimate reasonable and maximum prices that a scientist would pay, using research grants funds, for this product. Six scientists estimated what they thought would be a reasonable price for a device such as the one being designed, the other three did not have a price in mind. Five of the interviewees mentioned that the reasonable price should be \$50 000. Moreover, based on the interviews, we realized that the maximum price a researcher would pay for a device as the one being designed would be \$75 000. Therefore, the price ceiling was defined as \$75 000.

Price Floor: The price floor is determined by the development cost and is evaluated automatically by selecting the abandon option. In other words, the product design process should be abandoned if the benefit of developing it is not greater than its cost. When computing the project value using the dynamic programming recursive equation, we can establish the price floor.

2) Market Size: We were able to assess the market potential through intensive discussions with venture capitalists specialized in product development for the ophthalmic industry. For a product like the AOSLO, the venture capitalists expect a small market on the order of 50 products to be sold. If the company eventually decides to focus on human applications, the market would be much bigger (the total expected market for human applications would be around 25 000 units in the U.S. during the course of three years). This estimation is based on the fact that there are 16 000 ophthalmologists in the U.S. (three times this number in the whole world) and 60 000 optometrists.

B. Development Horizon and Review Stages

The product development process started in the summer of 2002 and, since then, scientists and engineers have been working together to develop the product. The first critical review stage using the decision model described in Section III took place in July 2003, during the prototyping phase and before the first tests. For this reason, we are not considering the initial investment and tradeoff decision regarding whether or not to start the development process. The remaining review stages for this project are: 1) test, in August 2003; 2) analyze, in October 2003; 3) design improvement, in January 2004; 4) manufacturing a compact system, in March 2004; and 5) market launch, in September 2004. We are not providing more details regarding the tasks performed at each project phase because we would be deviating from our major goal in this case study, which is to discuss the application of the decision support methodology. Fig. 1 illustrates the development phases and review stages.

C. Product Features—State Characterization

In this section, we describe how the controllable variables and the stage transition dynamics were modeled in order to evaluate the technological success of the design process. These controllable variables would work as a means for the decision makers to assess the product development performance at each review stage.

1) Controllable Variables: In order to accomplish the needs of future users (researchers), we assessed the development performance using the product features highlighted in the interviews with experts and scientists. Examples of these features are: lateral and axial resolution, contrast, adjustable field of view (zoom), ease of use, portability, among others.

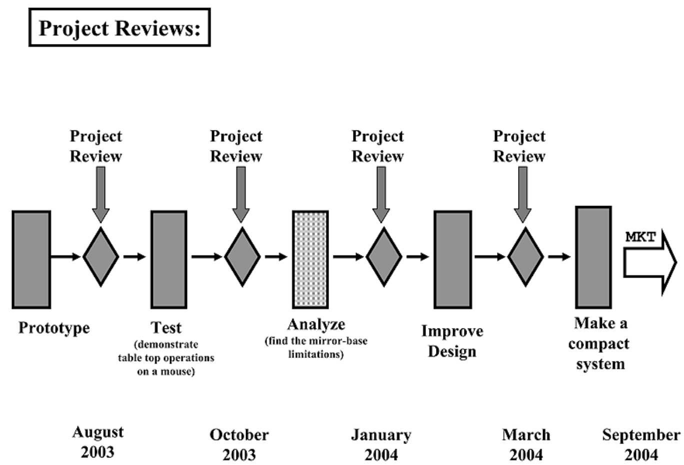

Fig. 1. Activities of the product development process.

Through the interviews, we learned that both lateral and axial dimensions of resolution of the device are critical factors in its success. The better the resolution, the better the product acceptance would be, and the stronger the argument to sell it. Although the other features are also important, most researchers expect the product to meet specific required levels and not continuously improve them. For example, many researchers selected a discrete required level for zoom (not continuously improve it) and, for the development team, this was seen as a "must achieve" feature specification. The same applies to requirements such as compactness and weight.

Thus, we decided to assess the development performance using two controllable variables. We called the first one "percentage of improvement in resolution," which would capture both dimensions of resolution improvement. Since lateral and axial resolutions are physically coupled (i.e., if one is improved, the other is also improved), the decision makers decided to follow up the development process by assessing only the lateral resolution performance, which would be easier to measure and follow up its performance at each project review.

Furthermore, the product development team decided that it would be better to measure the percentage of resolution improvement, instead of considering the resolution value at each stage. The baseline value for measuring the improvement is the current resolution of a scanning laser ophthalmoscope that do not use adaptive optics. These products are available on the market and have been used by both ophthalmologists for human treatment and scientists for research purposes.

The decision makers foresee the range of possible performance outcomes for the improvement in resolution as a nonlinear measure. In other words, considering only the noise effect between two review stages ("pure uncertainty" affecting the AOSLO resolution), it is easier to worsen the current resolution level than to improve it. To deal with this development assumption, the decision makers decided to follow up the logarithm of the percentage of improvement in resolution and assumed discrete intervals to measure each possible state of the development process. As a consequence, the logarithmic value of resolution improvement would be assessed from one review phase to the subsequent one on a linear scale representing the range of its possible outcomes.

At launch stage, we consider the worst possible development outcome in terms of resolution improvement and also the best case scenario (which corresponds to the diffraction limit). By considering these two states, we were able to create a lattice (discrete states) that would represent the performance measured during the development review phases given the uncertainty encountered at each phase transition. The lattice could only be determined after modeling the transition dynamics of resolution performance, which represents the performance evaluation at each review stage. The transition dynamics are described in the next subsection.

The second controllable variable is usability or fitness for use. Although the percentage of resolution improvement could be seen as the stronger argument for selling the AOSLO, the product should also meet other needs (market requirements) in order to be accepted by the market. The requirement in usability dimension could also be considered in the development of many other products. For example, despite the fact that nowadays much faster computer processors for laptops can be built, which could (for some market segments) become the most highly demanded feature by customers, some constraints (e.g., weight and cooling system) could block the viability or commercialization of this product. As a consequence, the market would not accept the product and the company cannot rely only on the processor performance to assess the technological success of a product.

Similarly, we decided to examine the dimension "usability." This dimension would incorporate assessment of features such as ease of use, portability, adjustable field of view (zoom), and external appearance. One important characteristic of the usability dimension is that it is independently perceived by the market when compared with the percentage of improvement in resolution dimension. At each review stage, we assign a value for this dimension based upon the probability of not meeting this requirement, which was determined during the interviews with scientists. We also consider a lattice for this dimension and, in the next subsection, we explain how it was modeled.

2) Development Uncertainty and Dynamics: The transition performance from one review stage to another is not deterministic, for either performance dimensions considered in this project. To model performance evolution, we consider that from one stage to another the product's expected "percentage in improvement in resolution" and "usability" (at launch) would only be affected by the level of investment selected at each review stage plus some "noise" effects. Furthermore, we consider uncertainty in one dimension independent from that in the other. The uncertainty assumptions were used to predict the possible range of performance levels for each attribute at each transition stage (between review stages).

To determine the percentage of improvement in resolution, we consider that the range of reachable performance levels from the prototyping stage to the testing stage is the highest one. In other words, the uncertainty regarding the possible outcome of the prototyping stage is the highest one when compared with the uncertainty of the performance outcome between the other review stages. Then, from the testing stage to the analysis phase there is also a higher uncertainty, which nevertheless is smaller than the first transition stage. To determine the subsequent

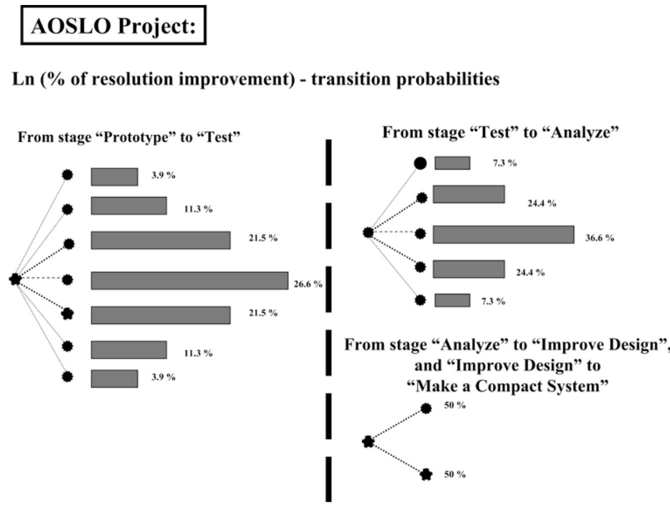

Fig. 2. Transition probabilities for resolution improvement.

stages, namely, "improve the design" and "make a compact system," the transition uncertainty is considered to be smaller than the previous ones ("prototyping" and "testing"). This assumption is quite intuitive and reflects the acceptable belief that the product development process has a higher performance variability in the earlier stages when compared with the expected performance at review stages closer to market launch.

We assume the noise effect in the transition outcome to be normally distributed, and then approximate the transition probabilities at each review stage to a discrete case that arbitrarily considers 2% of error (i.e., we approximate the normally distributed transition probabilities to a discrete case given some error). This arbitrary error assumption was made in order to highlight for the decision makers the fact that the model cannot capture all development outcomes (one can disregard this error and assume a triangular distribution to capture uncertainty in the next stage). At each review stage, we consider discrete intervals of the logarithmic value of the percentage in resolution improvement that would capture the product development performance in this dimension (in this case, we consider an interval equal to 0.10 units of performance measure). Fig. 2 illustrates our assumptions for the transition probabilities of the logarithmic value of the percentage of resolution improvement.

Given the transition probabilities for the logarithmic value of the percentage of improvement in resolution, at each development stage, we will have a range of measurable states, as shown in Fig. 3 (please note that the state space is not completely illustrated in the figure, that is one still needs to consider the states that can be reached through the improvement options).

A note on the reason for considering a lattice state space is in order at this point. Although we could have considered a general state space for both dimensions, we decided to use a finite-discrete state space. The main reason for doing so is to reduce the computational effort and the curse of dimensionality (for more details, see, e.g., [29] or [30]). For example, at launch state, we evaluated the project value at 279 points (31 possibilities in the percentage in resolution improvement times nine possibilities in the usability dimension). However, if we had considered a general state space, we would have had to evaluate the project

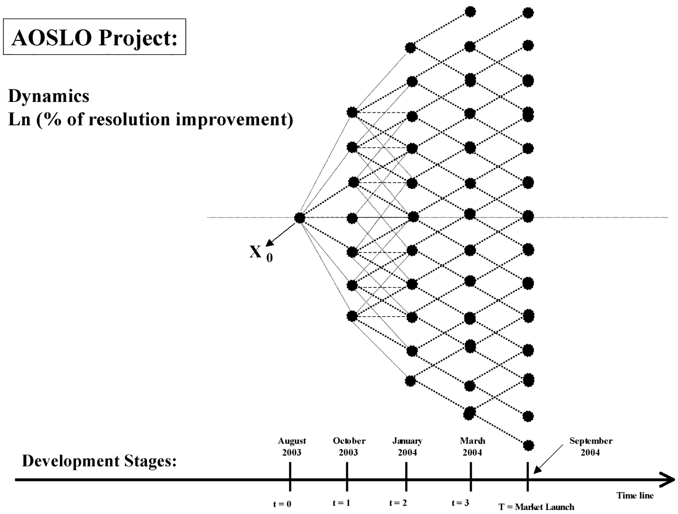

Fig. 3. Transition dynamics for resolution improvement.

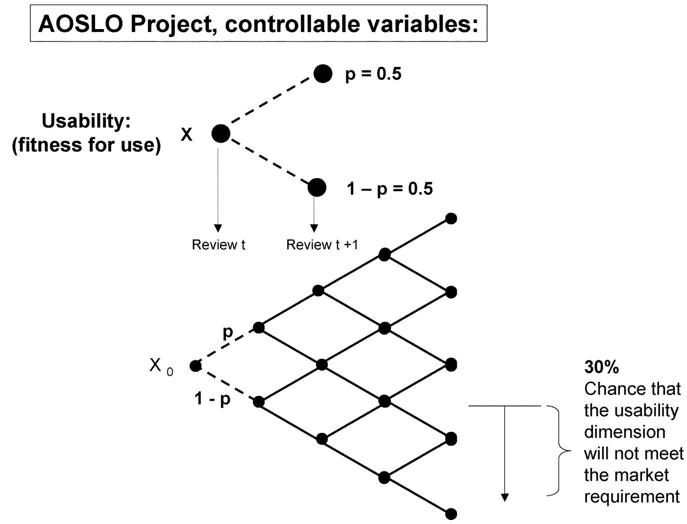

Fig. 4. Transition probabilities and dynamics for usability.

at 544 320 points (2240 possibilities in the percentage in resolution improvement times 243 possibilities in the usability dimension). Moreover, the further discretization of the state space would not lead to a concrete benefit for the project management activity since it would not match the project environment, and the market would not be able to differentiate among so many levels of resolution improvement/states.

To monitor the usability dimension, we consider a binomial lattice and assume the project is equally likely to increase or decrease its usability performance during the transition from one review stage to the next one (given the current product development performance, we consider that the next stage could be either 0.5 units of measure better or worse than the current usability value). Consequently, at launch stage, we can evaluate the probability distribution that the project will reach each performance state and, to set a value to the coordinate at the initial review stage, we consider the chance that we would not satisfy the market needs (probability of failure). In this project, the decision makers agreed that there is a 30% probability that the project will not meet the market requirement in the usability dimension (see Fig. 4. Also, as in Fig. 3, the state space is not completely illustrated). In the next section, we provide more details on how to define the initial value for this dimension, given the expected payoff for the project.

D. Expected Payoff

We evaluated the expected payoff as follows. We assume both percentage in resolution improvement and fitness for use as being independently perceived by the market and/or researchers. As a consequence, the expected payoff will be given by the probability that the design meets the required level in both dimensions.

To model the usability dimension, we took into account the information acquired in the interviews and consider that the market perceives it almost as a binary function. Therefore, we assume an approximation of a step function and position it at the point that matches the probability of failure, described in the previous section (30% for the prototyping stage).

To model the percentage in resolution improvement, we took into account the information acquired in the interviews for both lateral and axial resolution specifications. We approximate the required level for the lateral resolution to a normal distribution with standard deviation $\sigma = 0.49$ (units of the logarithmic value of the percentage of resolution improvement) and mean $\mu =$ 3.5 μm (micrometers) or 0.95 of percentage in resolution improvement. As discussed before, axial resolution performance would be monitored through lateral resolution performance.

We can evaluate AOSLOs expected payoff as described in (1), whose parameters are: 1) the probability of meeting the market requirements for both performance attributes (we consider the requirements independent of each other); 2) the ceiling price, which corresponds to "M" in our model; 3) the minimum payoff price "m" as being zero, since the development team does not consider any salvage value for this new product development project; and 4) the final states coordinates as described in the previous section. Fig. 5 illustrates the expected project payoff.

E. Managerial Actions & Development Cost

At each review stage, the management team assesses the performance of the development process and decides about future investments. The decision upon future investments is a function of the expected reward when selecting a given managerial action and the cost incurred to carry on the project. The managerial options at each stage are: 1) continue the project as planned by investing a baseline cost to carry on the development at that stage; 2) improve the project by spending extra resources on it, by which the team expects to improve the performance of the project; or 3) abandon the project, thus avoiding future expenses.

Initially, the management team decided that it would be possible to improve the usability dimension by only one state between March 2004 ($t = 3$) and September 2004 ($t = 4$), during the development phase "make a compact system." However, it is expected to improve the resolution dimension by 0.10 units of the logarithmic of resolution improvement (percentage) during the phases that start in August 2003 ($t = 0$), October 2003

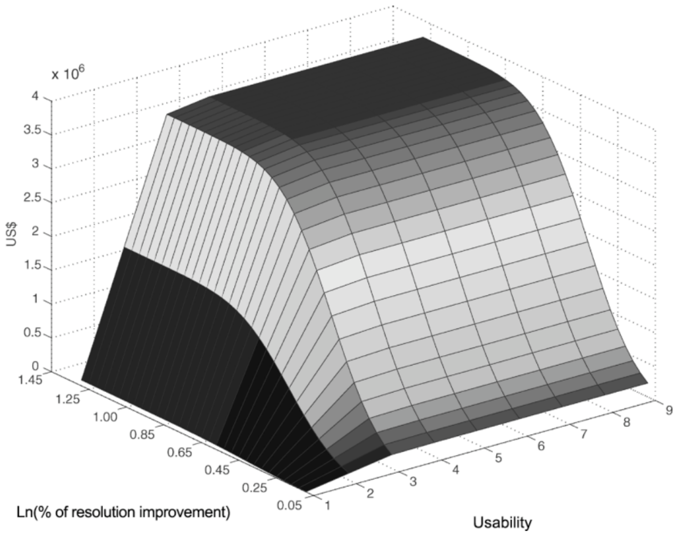

Fig. 5. AOSLO: expected payoff versus development performance.

$(t = 1)$, and in January 2004 $(t = 2)$. The cost at each stage (t) are as follows.

$(t = 0)$ $10 000 to continue the project as initially planned, and an additional $20 000 if the development team decides to add an extra engineer to the project.

$(t = 1)$ $20 000 to continue the project as initially planned, and an additional $40 000 if the development team decides to add an extra engineer to the project.

$(t = 2)$ $20 000 to continue the project as initially planned, and an additional $150 000 to review and improve the resolution of the prototype.

$(t = 3)$ $300 000 to continue the project as initially planned, and an additional if the company decides to improve the usability dimension.

$(t = 4)$ $1 750 000 corresponding to the production cost.

F. Basic Tradeoff at a Review Stage

The essential tradeoff that is the basis for management decision at a review point is captured by the dynamic programming recursive equation (1).

It is essential to reiterate that the project is reviewed four times (t =August 2003, October 2003, January 2004, and March 2004). In addition to these review periods, at the launch stage (September 2004), the management team still has the option to produce more units (production cost) and, therefore, there is one tradeoff regarding whether or not to produce them (either abandon the production or invest $1 750 000 to produce the AOSLO).

The performance vector $\underline{x}$ represents the dimensions: improvement in resolution (x_1) and usability (x_2). At each review stage, the management actions (control options) with their respective costs are the ones listed before. At the terminal stage, $t = T$ ($T =$ September 2004), the expected payoff is given by $V_T(\underline{x}) = E[\Pi(\underline{x})]$, as shown in Fig. 5.

The results, reported in the next section, were obtained by using an arbitrary discount rate of 1% per month. In addition, we also evaluated the project value using other discount rates, namely, risk-free and no discount rates, in order to assess the evaluation sensitivity as a function of this rate (please see discussion regarding different discount rates in the next section).

V. Results

The project value in August 2003 was $54 250, according to our decision support method. In addition to the AOSLOs expected value, the estimated product development performance was 4.3 μm for the resolution dimension, and we considered a baseline value for the usability dimension which is relative to its market requirement (assumed to be 0 in August 2003). The management team was advised to invest more money of the project budget in order to improve the resolution dimension. The expected improvement benefit is 0.10 logarithmic units, which corresponds to an expected resolution equal to 3.9 μm. The usability dimension in the next stage is expected to remain at the same level as in August 2003.

The positive value of the project, calculated in August 2003, indicates a profitable development. Although, at this point, the potential immediate profit is considered small, the decision makers strategically envision this project as a means to reach other markets and ramp up the company's development/commercialization capacity. Another expected benefit to be drawn from this development process, by the decision makers and the product design team, is that it could also be seen as a learning experience which will facilitate the development of future projects.

This decision support methodology also permits the characterization of future scenarios at subsequent review phases and highlights actions that can impact the development process. For instance, the company can predict the expected project value which corresponds to possible development outcomes. As can be seen in Fig. 6, for each possible pair of project performance values (usability and logarithmic value of the percentage in resolution improvement), there is an expected project value. The project value at each development stage is obtained through the tradeoff: managerial action (option to fund the development process) versus expected payoff given the selected action.

Similarly, one can map the possible managerial actions at each stage, as indicated in the Fig. 7. For each project performance value, the company needs to select actions for each of its dimensions. In other words, the development team can choose to act according to the managerial option of funding the usability dimension or that of funding the dimension "logarithmic value of the percentage in resolution improvement."

More specifically, Fig. 7 is displayed as follows: in the usability axis, the odd rows' elements represent recommended action regarding the "percentage in resolution improvement" dimension, and the even rows' elements represent actions recommended to be taken in the usability dimension. The scale of the axis "percentage of resolution improvement" is as described previously (i.e., there is no change in the displayed value). In order to better understand this, one could take the following example in which one assumes a certain value of project performance (i, j), where i represents the percentage of resolution improvement and j represents the usability dimension. In Fig. 7, the action recommended on the percentage of resolution improvement dimension is represented by the

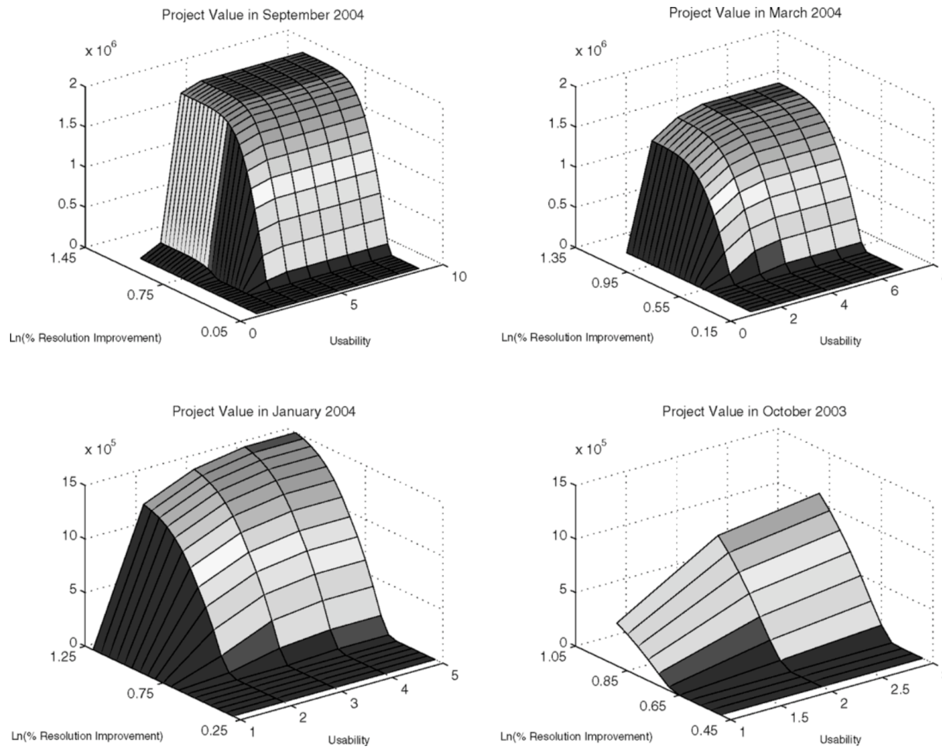

Fig. 6. Expected project value for the next development stages.

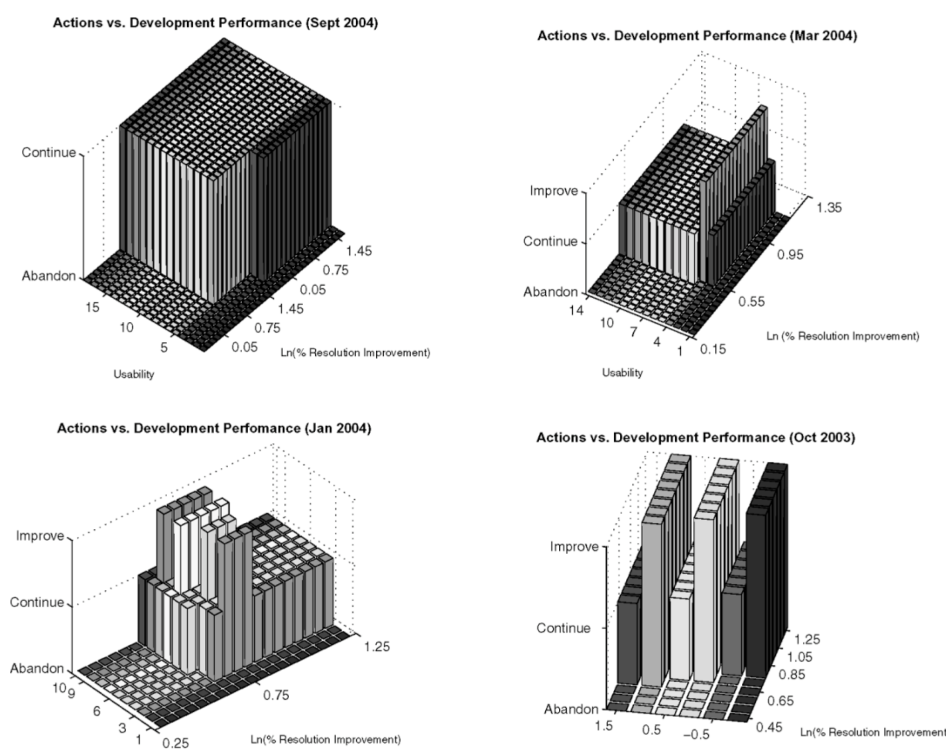

Fig. 7. Expected managerial actions for the next development stages.

element $(i, k-1)$, and the action recommended on the usability dimension is represented by (i, k), where $k = 2 * j$.

Although only the results concerning the subjective discount rate are reported, the company also evaluated the project value considering the risk-free interest rate and considering no discount factor $(r = 0)$. For all scenarios, the recommended action at the first review stage was the same, namely, invest more in the percentage in resolution improvement dimension. Furthermore, in the case in which no discount factor was used, the recommended actions for almost all states were the same as the ones when a discount rate equal to 1% was used (just 3 out of 304 states had different optimal actions evaluated). This fact illustrates that, for this particular project, the development costs are relatively small when compared with the final payoff. Moreover, one can see that for this particular new product design, uncertainty, and not the development costs, plays a core role.

The future outcomes prediction and characterization (project value, recommended managerial actions, and project performance at the next review stages) allow a company to establish contingent plans to deal with adverse scenarios, for example, if the project performance turns out to be worse than expected. At each review stage, the product development team can also identify and characterize (represented in each subfigure) the areas where more money should be invested in the project to improve its performance.

VI. Discussion of the Results

As argued by several researchers since the 1970s ([31]–[33]), one of the most important benefits of evaluating a product development process through a structured methodology is the discussion it triggers among the decision makers. In the management process of the AOSLO project, we also identified the importance of sharing information among members of the product development team, and in many instances, the decision model worked as a means to ensure a uniform knowledge among them.

As previous experiences using quantitative methods [34], this project evaluation is limited by the value of the input information. However, although this quantitative model is working with imprecise information and cannot be seen as a way to determine the exact price of the device, the product development team recognized the importance of the model as a way to assess a revolutionary product with no similar competitor in the market. This experience would be especially useful in cases where decision makers do not have enough information about past experiences, but need to understand inherent risks of the project and manage in uncertainty.

An important lesson drawn from this experience is the importance of considering a multidimensional vector to assess project performance. For instance, if we had considered only the percentage of resolution improvement to manage the project, we would be assuming that the usability dimension is accomplished with probability one, as shown in Fig. 8.

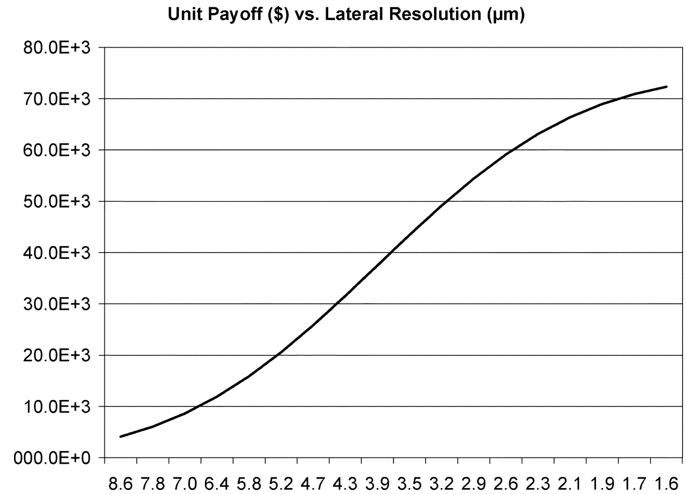

Fig. 8. Expected payoff versus improvement in resolution.

Although Fig. 8 lends itself to an intuitive interpretation, the AOSLO value that considers only one dimension (percentage in resolution improvement) would, in August 2003, be \$177 350 which is different from the one described in the previous section (\$54 250). One of the reasons for this discrepancy is that many undesirable scenarios were neglected in order to obtain the project payoff, as shown in Fig. 8. To list a few causes for this higher value, the company did not consider areas where the project should be improved in the usability dimension or moments when the project should be abandoned no matter which resolution value the product development team achieved, as shown in Fig. 6.

On the other hand, if the unidimensional payoff (considering only the dimension percentage in resolution improvement) had been assumed as 30% smaller, since there is 30% of probability of abandoning the project, the optimal action in August 2003 would be to abandon the development process. In other words, the project value would be smaller than the one evaluated for a two-dimensional case because there would be no flexibility to improve the usability dimension. As a clear consequence, the product development team would be taking undermined decisions and evaluating the project incorrectly. In summary, this example justifies the need to consider a multidimensional vector to assess the project characteristics and obtain a more realistic value, as well as recommended actions for it.

A second lesson/benefit acquired by using the model is to incorporate future generations of the current development project into the analysis (for example, low-cost derivative products that would lead to a higher payoff). Specifically, the future benefits can be taken into account when evaluating the minimum and maximum market payoff prices. In this case, the expected project value would have been evaluated higher than the one initially obtained (the company opted not to consider this possibility in this case study) and the recommended decisions would have been based on future deployments of the current development project (as if a platform project and the future derivatives of it were being considered all together). Similar to incorporating future generations into the analysis, product life cycle can be taken into account to try to assume more accurate market potential (see, for example, [35]).

One technique that can be used to enhance the discussion among the team members and the knowledge about the project is to perform a model sensitivity analyses. Some examples of analysis that can be done are: what is the impact in the project value and in the recommended managerial actions if: 1) the discount factor is increased or decreased?; 2) different subjective probabilities for transitions between review stages are considered; and 3) different expected improvement levels are considered; among others. This technique can explicitly highlight the models' potentialities and limitations.

As a fourth benefit, in the early stages of the development process, the project value obtained by means of the decision model could be used as a baseline for discussions with venture capitalists. Although the market price will probably not be exactly as stated, since several intrinsic uncertainties were taken into account when calculating it, the company managers and venture capitalists can use the methodology to assess an estimated project value and base their analysis for future investments on it. Another direct application of the methodology is to guide decision makers at each review stage on the project resource allocation by examining the tradeoff: development cost versus potential benefit.

It is critical to highlight the importance behind assumptions in this model. In other words, slightly different assumptions for the model input parameters can easily affect the final evaluated result. This fact makes the in depth discussion among decision makers crucial. For instance, different values for the parameters such as price ceiling (M), salvage value (m), transition probabilities, and expected improvement level can change the final result. As a first example, assume decision makers had considered, as the price ceiling, the reasonable price (\$50 000) instead of the maximum expected price (\$75 000), as a consequence, the recommended managerial action for this project would have been to abandon it in August 2003. As a second example, assume the company had considered some salvage value for this development, say $m =$\$200 000 (say, for example, some prototyping platforms and knowledge acquainted can be used in an other development), in this case, the project value would have been evaluated as \$70 516 and the future expected scenarios, in which the project would have been abandoned, would have been reduced (recall that the project value according to the original assumptions is 54 250). As a third example, if the company expected, by investing more money in the resolution dimension at each review stage, to achieve twice the benefit (i.e., 0.20 logarithmic units instead of 0.10 for the expected improvement, $\underline{k}$(Improve), in the logarithmic value of percentage in resolution improvement), the project value evaluated in August 2003 would have been \$329 235. These examples illustrate the importance behind the assumptions when evaluating the project value and optimal managerial actions at each review stage, and how these quantities can be easily affected by the model input parameters. Nonetheless, at the same time, they also unleash discussions among decision makers which can make them more confident about the model potentialities and limitations.

VII. CONCLUSION

In this case study, we presented one of the first applications of the decision support method introduced in [1] and further developed in [2], and illustrated the importance of considering a multidimensional state space to evaluate new product development performance.

Although it is only one application of the methodology, we believe this case study illustrates its viability. Furthermore, we hope this paper will encourage future applications of the model and generate future improvements of it. Topics such as transition probabilities, discount rate, and market definition should be discussed and improved further in future applications and/or research.

As future research, some practical implementations issues need to be discussed further. For instance, one topic that deserves more investigation could be to study when the decision model should be taken into account during the development process; that is, at which stage of the new product development process the company would be better off by using this decision support method. For example, if a company decides to use the model to evaluate the process at the very beginning of the development process (to decide whether or not to initiate the development of a project), one question that needs further investigation is how to assess the product development performance. It is clear that the product development performance could be easily assessed during the detailed engineering design, through physical prototyping or computer simulation. On the other hand, assessing product development performance in the concept generation phase(s) is not so straightforward. Therefore, a more detailed discussion regarding when the decision support method would be more efficient in evaluating research and development projects would be of great value to practitioners. At the heart of the discussion is the tradeoff between the quality of information versus the need to manage development uncertainty, where theoretically the model would be more appropriate.

Another topic to be further investigated is to identify suitable values for transition probability at each phase and for different applications. For example, transition probabilities values that better fit: 1) certain industry applications, such as pharmaceutical, electronics/high-technological products, and mass consumption; or 2) types of development projects such as platform, derivatives, and breakthrough products.

The type of market input information considered can also be enhanced. As an example, in future applications, more interviews should be performed in order to increase the accuracy of the information to be collected. Given that this product has such a specific market, namely, principal investigators of research grants who are studying mice's eyes, the interviewees were still considered a representative group. Another example would be to consider the product life cycle to evaluate the potential return (see, e.g., [35] for a discussion considering option evaluation). Another possibility is to further investigate the definition of the market potential for a revolutionary product, specifically how to assess the expected market payoff. For example, one can consider the "market payoff" as a "strategy payoff," in a broader definition, which allows for interactions with other products that might use similar technology, market positioning, etc. A similar study can also be carried out to investigate suitable discount rates (i.e., whether or not to consider a risk-free discount rate), which could also focus on industry applications and/or type of development projects.

ACKNOWLEDGMENT

The authors would like to thank P. Bierden, Boston Micromachines Corporation, for our meetings and discussions about development projects. They also would like to show appreciation to Schepens Eye Research Institute, in particular, Dr. S. Burns and Dr. R. Webb, and to the Massachusetts General Hospital's Wellman Laboratories of Photomedicine, in particular, Dr. C. Lin. Last, but not least, they would like to thank the Department Editor and two anonymous referees for their many valuable comments on an earlier draft of this paper.

REFERENCES

[1] A. Huchzermeier and C. H. Loch, "Project management under risk: using the real options approach to evaluate flexibility in R&D," *Manage. Sci.*, vol. 47, pp. 85–101, Jan. 2001.

[2] L. P. Santiago and P. Vakili, "On the value of flexibility in R&D projects," *Manage. Sci.*, 2005, to be published.

[3] A. K. Dixit and R. S. Pindyck, *Investment Under Uncertainty*. Princeton, NJ.: Princeton Univ. Press, 1994.

[4] L. Trigeorgis, *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. Cambridge, MA: MIT Press, 1996.

[5] M. Amram and N. Kulatika, *Real Options: Managing Strategic Investment in an Uncertain World*. Boston, MA.: Harvard Business School, 1999.

[6] T. W. Faulkner, "Applying "options thinking" to R&D valuation," *Res. Technol. Manage.*, vol. 39, pp. 50–56, May–Jun. 1996.

[7] T. A. Luehrman, "What's worth? A general manager's guide to valuation," *Harvard Bus. Rev.*, pp. 132–142, May–Jun. 1997.

[8] C. Y. Baldwin and K. B. Clark, "Design rules: The power of modularity," in *The Power of Modularity*. Cambridge, MA: MIT Press, 2000, vol. 1.

[9] P. D. Childs and A. J. Triantis, "Dynamic R&D investment policies," *Manage. Sci.*, vol. 45, pp. 1359–1377, Oct. 1999.

[10] M. J. Brennan and L. Trigeorgis, *Project Flexibility, Agency, and Competition: New Developments in the Theory and Application of Real Options*. New York: Oxford Univ. Press, 2000.

[11] G. Cortazar, E. S. Schwartz, and M. Salinas, "Evaluating environmental investments: A real options approach," *Manage. Sci.*, vol. 44, pp. 1059–1070, Aug. 1998.

[12] R. G. McGrath and I. C. MacMillan, "Assessing technology projects using real options reasoning," *Res. Technol. Manage.*, vol. 43, pp. 35–49, Jul./Aug. 2000.

[13] E. H. Bowman and G. T. Moskowitz, "Real options analysis and strategic decision making," *Organ. Sci.*, vol. 12, pp. 772–777, Nov.–Dec. 2001.

[14] D. M. Lander and G. E. Pinches, "Challenges to practical implementation of modeling and valuing real options," *Quart. Rev. Econ. Finance, Special Issue*, vol. 38, pp. 537–567, 1998.

[15] C. H. Loch and S. K. Kavadias, "Dynamic portfolio selection of NPD programs using marginal returns," *Manage. Sci.*, vol. 48, pp. 1227–1241, Oct. 2002.

[16] M. L. Barnett, "Falling off the fence? A realistic appraisal of a real options approach to corporate strategy," *J. Manage. Inquiry*, vol. 12, pp. 185–196, Jun. 2003.

[17] R. H. Webb, *Noninvasive Diagnostic Techniques in Ophthalmology*. New York: Springer-Verlag, 1990, ch. Scanning Laser Ophthalmoscope, pp. 438–450.

[18] R. W. Rodieck, *The First Steps in Seeing*. Sunderland, MA: Sinauer, 1998.

[19] M. Nagpal and S. Banait, "Scanning laser ophthalmoscopy in macular degeneration," *Bombay Hosp. J.*, vol. 44, Jul. 2002.

[20] R. H. Webb, G. W. Hughes, and O. Pomerantzeff, "Flying spot TV ophthalmoscope," *Appl. Optics*, pp. 2991–2997, 1980.

[21] R. H. Webb and G. W. Hughes, "Scanning laser ophthalmoscope," *IEEE Trans. Biomed. Eng.*, vol. 28, pp. 488–492, Jul. 1981.

[22] T. G. Bifano, J. Perreault, R. K. Mali, and M. N. Horenstein, "Microelectromechanical deformable mirrors," *IEEE J. Sel. Topics Quantum Electron.*, vol. 5, pp. 83–89, Jan./Feb. 1999.

[23] N. Doble, G. Yoon, L. Chen, P. Bierden, B. Singer, S. Oliver, and D. R. Willians, "Use of a microelectromechanical mirror for adaptive optics in the human eye," *Optics Lett.*, vol. 27, pp. 1537–1539, Sep. 2002.

[24] A. Roorda, F. Romero-Borja, W. J. Donnelly, III, and H. Queener, "Adaptive optics scanning laser ophthalmoscopy," *Optics Express*, vol. 10, pp. 405–412, May 2002.

[25] M. V. Marn, E. V. Roegner, and C. C. Zawada, "Pricing new products," *McKinsey Quart.*, no. 3, pp. 40–49, 2003.

[26] ——, "The power of pricing," *McKinsey Quart.*, no. 1, pp. 27–39, 2003.

[27] J. M. Abele, B. R. Elliott, A. A. O'Hara, and E. V. Roegner, "Fighting for your price," *McKinsey Quart.*, no. 4, pp. 116–125, 2002.

[28] E. V. Roegner, T. Seifert, and D. D. Swinford, "Putting a price on solution," *McKinsey Quart.*, no. 3, pp. 94–97, 2001.

[29] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Belmont, MA: Athena Scientific, 2000, vol. 1.

[30] D. P. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*. Belmont, MA: Athena Scientific, 1996.

[31] A. F. Helin and W. E. Souder, "Experimental test of a q-sort procedure for prioritizing R&D projects," *IEEE Trans. Eng. Manage.*, vol. EM-21, pp. 159–164, Nov. 1974.

[32] W. E. Souder, "A system for using R&D project evaluation methods," *Res. Manage.*, vol. 21, pp. 29–37, Sep. 1978.

[33] K. B. Clark and S. C. Wheelwright, *Managing New Product and Process Development: Text and Cases*. Boston, MA: Harvard Business School Press, 1993.

[34] M. J. Cetron, J. Martino, and L. Roepcke, "The selection of R&D program content—survey of quantitative methods," *IEEE Trans. Eng. Manage.*, vol. EM-14, pp. 4–13, Mar. 1967.

[35] N. P. B. Bollen, "Real options and product life cycles," *Manage. Sci.*, vol. 45, pp. 670–684, May 1999.

Leonardo P. Santiago received the B.E. degree in mechanical engineering and the M.S. degree in production engineering from Federal University of Minas Gerais, Minas Gerais, Brazil, in 1997 and 1999, respectively, and the Ph.D. degree in systems engineering from Boston University, Boston, MA, in 2004.

He is an Assistant Professor with the Production Engineering Department, Federal University of Minas Gerais, where he joined in 2004. He is a founding Member of the Brazilian Institute for Product Development Management. He has five years experience in the automotive industry, first as an engineer, and subsequently, as a consultant and researcher. His research focuses on new product development process, product portfolio definition and management, and Markov decision processes.

Thomas G. Bifano received the B.S. and M.S. degrees in mechanical engineering from Duke University, Durham, NC, and the Ph.D. degree in mechanical engineering from North Carolina State University, Raleigh, in 1988.

He is a Professor and Chair of the Manufacturing Engineering Department, Boston University, Boston, MA, where he joined in 1988. He is a founding Faculty Member of the Photonics Center, Boston University, and a Faculty Member of the Fraunhofer Center for Manufacturing Innovation. He is also Chief Technology Officer of Boston Micromachines Corporation. His research focuses on development of an advanced class of microelectromechanical systems (MEMS) for optical and biomedical applications.