Valuing Pilot Project Investments in Incomplete Markets: A Compound Option Approach

Eymen Errais *
Stanford University

Jeffrey Sadowsky †
Stanford University

January 18, 2005

Abstract

We introduce a general framework to value pilot project investments under the presence of both, market and technical uncertainty. The model generalizes different settings introduced previously in the literature. By distinguishing between the pilot and the commercial stages of the project we are able to frame the problem as a compound perpetual Bermudan option. We work on an incomplete market setting where market uncertainty is spanned by tradable assets and technical uncertainty is private to the firm. The value of these investment opportunities as well as the optimal exercise problem are solved by approximate dynamic programming techniques. We prove the convergence of our algorithm and derive a theoretical bound on how the errors compound as the number of stages of the compound option is increased. Furthermore, we show some numerical results and provide an economic interpretation of the model dynamics.

1 Introduction

The value of the investment opportunity to introduce a new product is subject to significant uncertainty. This uncertainty is exacerbated when this product comes together with a technological improvement. Market factors outside the control of a firm, such as product demand, prices of raw materials or labor costs are not the only sources of risk that affect the value of these projects. When new technologies or marginal improvements of existing technologies are involved, there is considerable technical uncertainty with respect to the final implementation and recurrent costs. This may also apply to revenues, as the value that these improvements may have for the final customer and, therefore, his willingness to pay for the product, could be affected by technical factors.

We could think of technical uncertainty as uncertainty with respect to the efficient production frontier of the firm, that is to say the final output/input ratio of the production function. A few examples of improvements in the efficient frontier of the firm are, for instance, a new technology that is able to perform exactly the

*Dept. of Management Science and Engineering, Terman Building , Stanford University, email: eymen@stanford.edu

†Corresponding author, Dept. of Management Science and Engineering, Terman Building , Stanford University, email: sadowsky@stanford.edu. We wish to thank Benjamin Van Roy for his help in the theoretical set up of our dynamic programming algorithm. Helpful discussions with Blake Johnson, Sachin Jain and Pedram Mokrian from Stanford University are also gratefully acknowledged

same functions as the previous one with fewer raw materials, or one that does so with cheaper inputs, or one that requires less time and investment for its implementation to take place. Unlike market uncertainty, which the firm has little control over, technical uncertainty could be greatly reduced by investing in pilot projects that give us considerable information about the performance of a new technology before the huge investment required for its full scale implementation is made.

In general, pilot investments have a very natural option interpretation. In fact, they could be considered as investments that allow the firm to limit its losses in case of negative outcomes, while maintaining the profits resulting from more favorable scenarios. However, unlike a financial option, our underlying, i.e., the value of the commercial stage, is not completely tradable in a market. Moreover, the process of the underlying is affected by the investment decisions taken by the firm. The higher the investment made in the pilot stage, the more technical uncertainty that is resolved. Thus, the value of the project will move more rapidly with increasing investment. These are the two fundamental differences between the real options setting in this paper and standard financial settings, and we will keep them in mind when interpreting our results.

Optimal R&D investment has been the subject of numerous studies in the economic and business literatures, beginning with Lucas (1971). He addressed the problem of optimal allocation of effort throughout the development stage of the project in a general case where effort is controllable and time to completion is random. However, there is no modeling of learning and all uncertainty is private.

The next major work in the area was done by Roberts and Weitzman (1981), who modeled continuous learning through time. Based on a proportionality assumption between cumulative investment and total uncertainty resolved they derive a diffusion process that should be followed by the expected benefits of a project. Since they do not take into account market uncertainty in their model, their results will be only applicable for a small set of projects where market uncertainty is negligible and could be ignored.

Grossman and Shapiro (1986) provide a few interesting models of R&D programs under certainty and uncertainty in progress and time to completion. The market dimension is absent and there is no linkage between the actions taken and the distribution of the time to completion.

McDonald and Siegel (1986) analyze the value of waiting to invest. In their model, the investment opportunity could be effectively translated into an option to exchange one asset for another. Therefore, they are able to adapt existing results in the financial literature (see Margrabe, 1978) to solve for the value of the option to invest. Their setting is useful when time to build is negligible and only market uncertainty is considered. It is also applicable when the development stage of the project has already been completed and a firm is considering whether to launch the commercial stage of a project or wait for more favorable market conditions.

Majid and Pyndick (1987) developed a continuous investment model with time-to-build. In their setting the only role of investment is to bring a project closer to completion and there is no learning involved. Pyndick (1993) is probably the first to take into account market and technical uncertainty into a coherent framework that values investments with uncertain costs. In his model revenues are fixed and costs are driven by market and technical uncertainty. However, unlike in Roberts and Weitzman, the stochastic process for the evolution of costs is not derived from fundamental principles about learning. Moreover, he does not distinguish between development and production stages. This distinction turns out to be relevant as, in many industries, most technical learning takes place in the former.

More applied work has been done recently, most of it focused on specific industries or project characteristics. Messica and David (2000) analyzed the effect of the life cycle of the future project's revenues on the optimal investment allocation in the development stage. Cortazar et al (2001) focused on optimal exploration investments in a mine under price and geological uncertainty. Bach. et Paxson (2001) modeled investment in the drug development process. Schwarz et Soraya (2003), using a model similar to that of Pyndick (1993), analyzed investment in the IT industry both in acquisition and development projects.

In general, the value of investment opportunities in the pilot stage will be driven by both technical and market uncertainty. As it was previously mentioned, market uncertainty will, in most of the cases, be completely exogenous to the firm and correlated to economic fundamentals and tradable assets in the market. Technical uncertainty is private to the firm, and its evolution depends on the level of pilot investment made by the firm. In this paper, by adopting a compound option approach and using dynamic programming techniques, we will value these investment opportunities and find the corresponding optimal investment decisions. In our setting, both market and technical uncertainties are dynamically evolving and affecting the evolution of the value of the option through time. Tradability assumptions will be made on the market uncertainty driving the value of the project and we will value technical uncertainty by specifying a "market price" of technical risk that will play a role of a utility function by characterizing the firm's attitude towards technical risk. This will allow us to work in a unique risk neutral measure and to apply standard financial engineering techniques to obtain the value of our option to invest. Thus, the results obtained shall not be interpreted as strict non-arbitrage prices of financial derivatives, but as plausible economic valuations of these investment opportunities.

Throughout this paper we will use the words pilot and development stage interchangeably to refer to the period of time before a decision was made to launch the full scale project. For the following stage we will use the terms commercial or production stage indistinctly. At some point this may become an artificial distinction; nevertheless, this will allows us to formulate our problem in an optionality setting.

The main contributions of this paper lie under two different areas: Real Options and Dynamic Programming. In the former, we provide a general framework for valuing pilot project investments in an environment, where both market and technical uncertainty affect the value of the underlying asset. We emphasize on the distinction between them, and relate the volatility coefficient driving technical risk to general economic assumptions about the learning process. On the Dynamic Programming side, we successfully apply approximate dynamic programming techniques in an infinite dimensional unbounded state space with independent increments to solve for our value function. We proof the convergence of the proposed algorithm and provide theoretical bounds on how the errors compound as we increase the number of stages of the compound option. The algorithm proposed could also be applied to a variety of financial products, being specially suited to value compound perpetual Bermudan options.

The remainder of this paper is structured as follows. Section 2 sets up the general model explaining our modeling assumptions in detail and defines the general Bellman Equation to solve. It also introduces additional simplifying assumptions that allow us to keep the model computationally tractable without altering the main essence of the problem. Section 3 develops the approximate dynamic programming approach adopted to solve our problem, proves the convergence of the algorithm and derives theoretical bounds on the errors. Section 4 discusses the results of the paper. Section 5 concludes.

2 The Model

2.1 General Structure

In our setting, we consider a firm analyzing the possibility of launching a big commercial project, for example, introducing a new drug, a new model of aircraft, or starting the operation of an oil well. Due to the magnitude and risk level of the project's commercial stage, we assume that it can not be launched before investing in N steps of a pilot project. The goal of the pilot stage, aside from making the launch of the commercial stage feasible, is to resolve most of the technical uncertainty associated with a project of this nature. To consider a pilot step completed the firm has to invest an amount I belonging to the interval $[\underline{I}, \overline{I}]$ that denotes the maximum and minimum possible investment levels per step. Completion of any pilot step takes ΔT units of time regardless of the level of investment. Investment decisions are made at times $t \in \Lambda = \{0 = T_0, T_1, ..., T_i, ...\}$ where $T_i > T_{i-1}$ and $\Delta T = T_i - T_{i-1}$ for any $i \geq 1$. In other words, funding decisions are made at discrete points in time with a periodicity of ΔT. This is a realistic assumption as most firms revise their funding decision periodically. Unlike most compound options in the financial market, it is perfectly possible for the firm to suspend investment on the pilot at a certain time T_i for any $i \geq 1$ if, for instance, market conditions are not favorable, and resume investment at a later point in time T_{i+k}, where

$k \in \mathbb{N}$. Moreover, we will work under the assumption that the progress made in previous steps is not lost if investment is suspended temporarily.

Let S_t be the value of a claim to all revenues of the commercial stage based on the information available at time t. Similarly, K_t, will be the value of a claim to all costs corresponding to the commercial stage, implementation and recurrent costs. We could think of S_t and K_t as the Net Present Value (NPV) of these cash flow streams assessed at time t. In addition, S_t and K_t will be the values of these claims, assuming that the commercial stage could be launched at the current time t. Of course, this is not possible if the pilot stage has not yet been completed, making these claims unfeasible in this case. However, by separating the underlying assets from the option to invest, this abstraction will allow us to formulate our problem consistently and, since once the pilot stage is completed these claims will become feasible, it will not affect the results obtained. More structure to the processes followed by S_t and K_t, as well as the information set available at time t, will be described in greater detail in the next subsection.

We could think of this investment opportunity as a perpetual N-stage Bermudan option. A perpetual Bermudan option is an option that could be exercised at fixed specified dates in the future with no expiration date. In our setting, this set of fixed dates is given by Λ. At this fixed set of dates the firm has the option to acquire an $N-1$ stage perpetual Bermudan option, which will give it the right to obtain, at a fixed set of dates, an $N-2$ stage perpetual Bermudan option and so on. However, the amount of money spent on the option influences the volatility of the underlying assets. This optionality set up will become clearer in the next section once we explain in detail the assumptions regarding the learning process and the processes governing the evolution of revenues and costs.

2.2 Learning and Stochastic Process

In this subsection we will state our assumptions regarding the evolution of the underlying assets S_t and K_t, namely, the value of a claim to revenues and a claim to costs of the commercial stage. The risk free rate in our setting will be denoted by r_f.

Assumption 1: The Process for Revenues

We assume the revenue process to be completely driven by market uncertainty and perfectly correlated with a tradable asset, which, without loss of generality, is taken to be S_t itself. The value of S_t follows:

$$dS_t = \alpha_s S_t dt + \sigma_s S_t dw_t^1 \tag{1}$$

where w_t^1 is a standard Browian motion that represents the market uncertainty driving revenues. Moreover,

S_t accrues dividends at a rate δ_s. In other words, the market price of risk of w_t^1, which we label by λ_1, is uniquely defined by: $\lambda_1 = (\alpha_s + \delta_s - r_f)/\sigma_s$. This market price of risk represents the excess return over the risk free rate per unit of volatility that an investor demands for being exposed to a particular source of risk, in this case, w^1.

The fact that there is no technical uncertainty with respect to revenues is a reasonable assumption for a large number of investment projects of this type, as technological improvements are usually reflected in cost reductions and not revenue increases. Nevertheless, our setting could be easily generalized to account for private uncertainty on the revenues side of a project. The spanning assumption for revenues is a sensible assumption to make in industries where prices and demand are very sensitive to economic fundamentals, especially when the output of a project is a commodity or a marketed asset. These economic fundamentals are in turn correlated with the stock market in general and, in particular, with the stock of the firm. This depends on the characteristics of the industry such as the type of product supplied, the elasticity of demand with respect to price, the competition dynamics, etc. For instance, in the airplane manufacturing industry, it is reasonable to expect the value of a claim to future revenues on a new jet to be developed by *Boeing* or *Airbus* to be highly correlated with the state of the world economy in general and the firms' stock prices in particular. The automobile manufacturing and high-tech industries are other examples where the spanning assumption is also suitable. However, it may be too strong of an assumption to make for industries like the pharmaceutical and food industries that provide basic products characterized by a very inelastic demand with respect to price.

The dividend rate δ_s has a rich economic interpretation. First, it could be thought of as the yearly cash inflow that the project will currently generate if it were properly functioning. In this purely financial view δ_s is completely equivalent to a dividend yield. However, in a competitive market environment, a delay in launching a new product may have more undesirable consequences to a firm. Indeed, market share could be lost to competitors who may launch a close substitute in the market place while the firm is waiting for more favorable conditions. We could think of this second interpretation, as an opportunity cost view on δ_s. In many industries the long term effects of losing market share may be much higher than the yearly revenue forgone, or not be appropriately reflected in the actual dividend rate observed in the market. To account for these effects, the parameter δ_s should be increased in these cases.

Assumption 2: The Process for Costs

The process for the costs K_t is driven by both, market and technical uncertainty. The market uncertainty driving costs, which we will denote by the brownian motion w_t^2, is perfectly correlated with a tradable asset, which we label by c_t. Moreover, this asset accrues dividends at a rate δ_c and follows the following process:

$$dc_t = \alpha_c c_t dt + \sigma_c c_t dw_t^2 \tag{2}$$

We will also assume that w_t^1 and w_t^2, i.e., market uncertainty of revenues and costs, have a correlation of ρ. The previous equation implies that the market price of risk of w_t^2, which we label λ_2 is completely specified by non-arbitrage conditions and equal to $(\alpha_c + \delta_c - r)/\sigma_c$. Total costs are usually influenced by market factors such as prices of materials, capital or labor costs, to cite a few. Thus, the economic justification for the tradability assumption on w^2 is similar to the one given for the market uncertainty of revenues.

If at time $t \in \Lambda$ there are j stages in the pilot project remaining for completion and the firm decides to invest an amount I, the process that K_t will follow between t and $t + \Delta T$ for any $i \geq 0$ is given by:

$$dK_t = \alpha_k(K_t, I, j)dt + \sigma_k K_t dw_t^2 + g(K_t, I, j)dz_t \tag{3}$$

The Brownian term z_t corresponds to technical uncertainty, which is private to the firm and independent of w_t^1 and w_t^2. Unlike the cases of w^1 and w^2, the market price of risk for z, λ_z, is not determined by non-arbitrage considerations, but has to be assessed from the risk preferences of the firm. The growth rate $\alpha_k(K_t, I, j)$ and technical volatility level $g(K_t, I, j)$ are intrinsic to the project and technical characteristics of the investment under consideration and, in general, dependent on the investment level and the number of steps remaining for completion. However, log market volatility, σ_k, is assumed to be constant and independent of the level of investment. More about the structure of the drift and technical diffusion coefficient will be specified below.

Assumption 3: Learning

The diffusion term that drives technical uncertainty is a function of the investment made for the pilot step and the number of steps remaining for completion, $g(K_t, I, j)$. It is here where the main difference between this setting and standard financial options arise, through the pilot investment level we can influence the process of the underlying. We introduce the following general restrictions on the drift and technical volatility terms, $\alpha_k(K_t, I, j)$ and $g(K_t, I, j)$ respectively. As we will see below all of these restrictions correspond to general and reasonable economic assumptions about the learning process. Note that whenever we talk about the "learning process" in this paper, we are referring exclusively to the process driving technical uncertainty.[1]

Assumption 3a : $g(K_t, 0, j) = 0$. If there is no investment, no learning takes place, and thus, no technical uncertainty is resolved.

[1] Our framework could be easily adapted to value investment opportunities, where learning is on the market side, i.e., that help reduce market uncertainty in the commercial stage. Valuing marketing pilot projects is the first example that comes to mind

Assumption 3b : $\delta g(K_t, I, j)/\delta I \geq 0$. The higher the investment made, the higher the resolution of uncertainty.

Assumption 3c : $g(K_t, I, 0) = 0$. There is no technical learning once the pilot stage has been completed.

Assumption 3d : $\Delta_j g(K_t, I, j) = g(K_t, I, j) - g(K_t, I, j-1) > 0$. There is more learning at earlier steps of the pilot stage. We could label this condition decreasing learning with respect to investment. It is consistent with general economic assumptions about decreasing marginal returns of factors of production. In addition, if we consider each step as part of a series of independent experiments that could be placed in a flexible order, the firm will want to reorder this steps such that those that reduce more technical uncertainty are placed first. Marginal learning decreasing in the number of steps completed (i.e increasing in the number of stages remaining for completion) will result as a consequence of this reordering.

Assumption 3e : $\delta g(K_t, I, j)/\delta K_t \geq 0$. This is a logical scaling assumption and says that the higher the total costs the higher the total uncertainty reduced.

The drift term will also be in general dependent on the level of investment and the number of stages remaining for completion. We will assume $\alpha_k(K_t, I, j)$ satisfies the following assumptions

Assumption 3f : $\delta\alpha_k(K_t, I, j)/\delta I \leq 0$. Drift in costs is non increasing with respect to investment. The higher the investment that takes the place the more costs are expected to drop.

Assumption 3g : $\Delta_j \alpha_k(K_t, I, j) = \alpha_k(K_t, I, j) - \alpha_k(K_t, I, j-1) \leq 0$. Expected cost reduction decreases with the number of stages completed. It has a similar economic interpretation as the analogous restriction for technical volatility.

Unlike for the case of technical volatility, when the pilot stage is completed ($j = 0$) or when no investment is taking place, the drift is non-zero, as costs are still expected to change as a result of market uncertainty.

2.3 Change of Measure

In standard financial settings, markets are complete and all different sources of risk are spanned by tradable assets in the economy. Thus, any financial derivative whose payoff is constructed from these tradable assets must have a unique price consistent with non-arbitrage. Moreover, there is a unique risk neutral measure under which the price of any tradable asset discounted by the money market account is a martingale[2], i.e., the price of any tradable asset grows at the risk free rate.

In our setting, the technical uncertainty faced by the projects, unlike market uncertainty driven by the

[2]It is actually a local martingale. However since we are working with "nice" diffusion coefficients through this paper, every local martingale is a martingale.

brownian terms w_t^1 and w_t^2, is specific to the firm and hence not traded in the market. Therefore, as in most real investment opportunities, we are working in an incomplete market setting where there are infinitely many values consistent with the non-arbitrage condition, and infinitely many equivalent risk neutral measures that give us an arbitrage free economic valuation.

As the non-arbitrage bounds could be quite big, and hence of little practical use, additional considerations as how the firm values private risk are needed. This is done by fixing a “market price” for technical uncertainty, which we will call λ_z. Unlike the market prices for w^1 and w^2 which are completely determined by non-arbitrage conditions, λ_z is not a true market price since technical uncertainty is not tradable in the market. However, it will play the role of a utility function and allow us to “complete” the market in order to work in an equivalent risk neutral measure parameterized by this λ_z. By changing its value we will allow for different risk specifications with regard to technical uncertainty, obtaining different equivalent risk neutral measures, thus creating different plausible economic valuations for our investment opportunity.

Following standard financial mathematics argurments, we deduce from assumptions 1 and 2 that the market price of risk for the market brownian of revenues (w^1) and costs (w^2) are given by $(\alpha_s + \delta_s - r_f)/\sigma_s$ and $(\alpha_c + \delta_c - r_f)/\sigma_c$ respectively.

Now, fix λ_z and denote the unique risk neutral measure defined from it as $Q(\lambda_z)$. For ease of notation we will omit the explicit dependence hereafter. Since we want to work with a standard three-dimensional Brownian Motion, we first decompose w_t^2 in w_t^1 and a component orthogonal to both w_t^1 and to z_t, which we will call $w_t^{2'}$ (the market component of costs that is orthogonal to that of revenues). By Cholesky decomposition, we can then rewrite the process for c_t as:

$$dc_t = \alpha_c c_t dt + \rho\sigma_c c_t dw_t^1 + \sqrt{1-\rho^2}\sigma_c c_t dw_t^{2'}$$

we can now redefine the new risk premium for $w^{2'}$ which we denote by $\lambda_{2'}$ as the solution to:

$$\alpha_c + \delta_c - r_f = \lambda_1 \rho\sigma_c + \lambda_{2'}\sqrt{1-\rho^2}\sigma_c$$

This simplifies to :

$$\lambda_{2'} = \frac{\alpha_c + \delta_c - r_f - \rho\sigma_c(\alpha_s + \delta_s - r_f)/\sigma_s}{\sigma_c\sqrt{1-\rho^2}}$$

Following standard arguments in the finance literature (Duffie (2001)), we construct the equivalent measure as follows:

Let $\eta = (\lambda_1, \lambda_{2'}, \lambda_z)$ be the vector of market prices of risk for the three orthogonal Brownian Motions. We define a three-dimensional Brownian Motion, B_t, as $B_t = (w_t^1, w_t^{2'}, z_t)$.

Let ξ_t be a density process defined as:

$$\xi_t = exp\left(-\int_0^t \eta dB_s - \frac{1}{2}\int_0^t \eta.\eta ds\right)$$

By Ito's Lemma it is easy to verify that $d\xi_t = -\xi_t \eta dB_t$. Moreover, since the market price of risk η is bounded, ξ_t is a martingale[3] with finite variance and, therefore, the density process of an equivalent probability measure Q, which we defined by $\frac{dQ}{dP} = \xi_t$.

Then by Girsanov's Theorem, a standard Brownian Motion under Q, B_t^Q, is defined by $B_t^Q = (w_t^{1Q}, w_t^{2'Q}, z_t^Q) = B_t + \eta t = (w_t^1 + \lambda_1 t, w_t^{2'} + \lambda_{2'} t, z_t + \lambda_z t)$. Let w_t^{2Q} be a Brownian motion independent of z^Q but that has a correlation of ρ with w^{1Q},

Using Girsanov's theorem and the values for market prices of risk we obtain the process followed by the underlying claims under the risk neutral measure.

When switching from the real to the risk neutral measure an important modification needs to be made when working with the cost process K_t. If the firm is risk averse, standard economic theory suggests that it asks for a lower reduction in expected costs to undertake technological investments that are more volatile. [4] In other words, as a risk averse firm will be more conservative in its assessments of costs, it will use a lower rate to discount them. The effect of risk aversion on the discount factor of costs is therefore opposite to the effect on the discount factor of revenues. Thus, we add instead of substracting the term $\lambda\sigma$ from the drift rate in the risk neutral measure. Equivalently we can work with the process of $-K_t$, since, consistent with our definition, a claim on costs forces the firm to pay out all the cash disbursements required by the implementation of the projects. With these consideration, the process for S_t and K_t under Q are given by:

$$dS_t = (r_f - \delta_s) S_t dt + \sigma_s S_t dw_t^{1Q} \tag{4}$$

$$dK_t = (r_f K_t - \delta_k^{\lambda_z}(K_t, I, j))dt + \sigma_k K_t dw_t^{2Q} + g(K_t, I, j) dz^Q \tag{5}$$

where $\delta_k^{\lambda_z}(K_t, I, j)$ is given by:

[3] Actually, Novikov's condition, i.e., $E\left[exp(\frac{1}{2}\int_0^T \eta_s.\eta_s ds)\right] < \infty$ is a much weaker requirement for ξ_t to be a martingale and accommodates stochastic or time dependent market prices of risk. See Duffie (2001) or Musiela and Rutkowski (1997) for a more rigorous treatment of the subject. This condition is automatically satisfied in our setting since we will work with constant market prices of risk.

[4] A risk averse agent that follows orthodox economic principles is risk averse for losses as well as for revenues. However, there is a growing literature in Behavioral Finance and Economics that focuses on non standard utility functions. In some cases, due to framing effects an agent may be risk averse for wins and risk seeking on loses. However, for our setting, this framing effect makes little sense as costs are not evaluated separately as pure losses.

$$\delta_k^{\lambda_z}(K_t, I, j) = [r_f K_t - \lambda_2 \sigma_k K_t - \lambda_z g(K_t, I, j)] - \alpha_k(K_t, I, j) \tag{6}$$

$$\delta_k^{\lambda_z}(K_t, I, j) = [r_f - \frac{\sigma_k}{\sigma_c}(\alpha_c + \delta_c - r_f)]K_t - \lambda_z g(K_t, I, j) - \alpha_k(K_t, I, j)$$

where the superscript λ_z indicates that the function δ_k is λ_z dependent.

The term $\delta_k^{\lambda_z}$ is analogous to a dividend rate of a financial asset, but it has a more appealing economic intuition. From expression (6) we could see that it represents the difference between the rate at which firms should discount costs according to their risk preferences (the term inside the brackets) and the actual growth rate of these costs. It is through this term that changes in risk preferences with regard to private uncertainty modifies the economic valuations obtained using risk neutral pricing.

Note that, since we will work exclusively in the risk neutral measure Q, we will omit the superscripts Q from the Brownian terms from here on.

2.4 Bellman Equation

We would like to determine the value of the investment opportunity $V(S_t, K_t, j)$ at each $t \in \Lambda$ as a function of the underlying assets and the number of pilot steps remaining for completion. We would also want to find the optimal feasible investment policy $I(S_t, K_t, j)$ to follow at each revision time $t \in \Lambda$. To be feasible $I(S_t, K_{t,}j)$ has to be either 0 (no pilot investment is realized) or take any level value in the interval $[\underline{I}, \overline{I}]$, i.e., we complete one more step in the pilot project and have $j-1$ steps remaining for completion at time $t + \Delta T$.

The value of the investment opportunity $V(S_0, K_0, j)$ is given by the solution to the following Bellman Equation:

$$V(S_t, K_t, j) = \max_{I \in [\underline{I}, \overline{I}] \cup \{0\}} \left\{ -I 1_{\{I>0\}} + \alpha E^Q [V(S_{t+\Delta T}, K_{t+\Delta T}^{j,I}, j - 1_{\{I>0\}})] \right\} \tag{7}$$

where $\alpha := e^{-r_f \Delta T}$ is the discount factor, $t \in \Lambda$ and S_t and $K_t^{j,I}$ follow (4) and (5) respectively with starting points S_t and K_t. The expectation is taken under the measure Q and, as previously mentioned we omit the superscript hereafter.

In the case case where there is no possibility of waiting to launch the commercial stage after the pilot project is completed, the boundary condition of the problem is given by:

$$V(S_T, K_T, 0) = [S_T - K_T]^+$$

where T corresponds to the time of the completion of the last pilot step. However, if the firm has the flexibility to wait before deciding to launch the project and, moreover, could do so at any point in time after the pilot stage is completed, the boundary condition becomes the value of a perpetual Option to exchange one asset for another. When S_t and K_t follow log normal process this option has a closed form solution. We refer the reader to Merton (1973) and Margrabe (1978) for additional details.

The model presented above sets up the most general framework to work with. Due to its flexibility, it can be used to model many investment situations seen in the real world. However, when it comes to valuing the pilot project investments, we face a computational challenge mainly due to what is known in the dynamic programming literature as the *curse of dimensionality.* Indeed, we have to keep track of the value of S, K and j at each time step. In addition, the space of decisions is infinite, and this makes finding the optimal strategy cumbersome. Therefore we introduce the following simplifying structure to come up with a computationally tractable model without altering the main essence of the problem.

1. Revenues will be considered fixed at a level $\bar{S}$
2. We set only one level of investment $\overline{I} = \underline{I} = I$. Hence the decision becomes a binary go-no go decision.
3. We work with $g(K_t, I, j) = \sigma_z K_t I j^{\gamma}$ where σ_z is a proportionality constant and $\gamma > 0$. This form satisfies assumptions 3a through 3e with $\gamma > 0$ guaranteeing that learning is higher in earlier steps than in later steps.
4. For the drift, we will use for simplicity: $\alpha_k(K_t, I, j) = (\alpha_k - \beta I j^{\gamma})K_t$ where α_k and $\beta > 0$ are constant terms. This form is consistent with assumptions 3f and 3g. We could think of α_k as the log normal drift driven by market factors since it is independent of the level of investment. The second term is the expected reduction in drift following technical investment and has a similar form as the one assumed for $g(K_t, I, j)$. Thus, investment is expected to reduce costs more drastically at earlier steps of the pilot. From (6), $\delta_k^{\lambda_z}$ could now be simplified to:

$$\delta_k^{\lambda_z}(K_t, I, j) = [r_f - \alpha_k - \frac{\sigma_k}{\sigma_c}(\alpha_c + \delta_c - r_f) - (\lambda_z \sigma_z + \beta) I j^{\gamma}]K_t$$

   Note that we obtain a log normal expression. To take advantage of this simplified form we will work with the dividend lognormal coefficient, which we will label $\hat{\delta}_k^{\lambda_z}(I, j)$ which is given by:

$$\hat{\delta}_k^{\lambda_z}(I, j) = r_f - \alpha_k - \frac{\sigma_k}{\sigma_c}(\alpha_c + \delta_c - r_f) - (\lambda_z \sigma_z + \beta) I j^{\gamma} \quad (8)$$

The dividend coefficient above below has an interesting economic interpretation. The first three terms inside the brackets are drift effects caused by market factors regardless of the level of investment. The last term only affects the drift when technical investment, and thus, learning, takes place. This investment effect on the drift could, in turn, be decomposed in two separate ones. The β term represents the intrinsic effect on the drift that depends on the technological characteristics of the investment under consideration. The λ_z term is the effect on the present value of costs caused by an increase in the discount factor, as the future costs outflows are subject to less uncertainty when learning is taking place.

5. After the pilot stage is completed the firm has the ability to defer the launch of the commercial stage to any point in time. When revenues are fixed at $\bar{S}$ this boundary condition amounts to nothing more than the value of a perpetual American put, which is given by:

$$V(K_T, 0) = 1_{\left\{K_T \geq -\frac{\rho \bar{S}}{1-\rho}\right\}} \frac{(-\rho)^{-\rho} \bar{S}^{1-\rho} K_T^{\rho}}{(1-\rho)^{1-\rho}} + 1_{\left\{K_T \leq -\frac{\rho \bar{S}}{1-\rho}\right\}} [\bar{S} - K_T]^+ \tag{9}$$

with ρ and δ defined as:

$$\rho = 1/2 - (r_f - \delta)/\sigma_k^2 - \sqrt{2r_f/\sigma_k^2 + \left[(r_f - \delta)/\sigma_k^2 - 1/2\right]^2}$$

$$\delta = r_f - \alpha_k - \frac{\sigma_k}{\sigma_c}(\alpha_c + \delta_c - r_f)$$

These additional simplifications make the problem tractable while maintaining the essence of the paper, which is to capture the value of technical learning under the presence of both, market and technical uncertainty and the tradeoffs between the value of investing in the pilot stage and its costs.

Consequently, the Bellman equation that we need to solve is a simplification of (7):

$$V(K_t, j) = \max\{-I + \alpha E[V(K_{t+\Delta T}^{j,I}, j-1)], \alpha E[V(K_{t+\Delta T}^{j,0}, j)]\}$$

where between times t and $t + \Delta T$, $K_t^{j,0}$ and $K_t^{j,I}$ follow :

$$dK_t^{j,0} = \left(r_f - \hat{\delta}_k^{\lambda_z}(I, j)\right) K_t^{j,0} dt + \sigma_k K_t^{j,0} dw_t^2 \tag{10}$$

$$dK_t^{j,I} = \left(r_f - \hat{\delta}_k^{\lambda_z}(I, j)\right) K_t^{j,I} dt + \sigma_k K_t^{j,I} dw_t^2 + \sigma_z K_t I j^{\gamma} dz_t \tag{11}$$

with $\delta_k^{\lambda_z}(I, j)$ given by (8) and the boundary condition given by (9).

3 Solution Approach: Approximate Dynamic Programming

We would like to find the value function $V(K_t, j)$ and optimal investment policy $I(K_t, j)$ at any step $j = 0, 1, ...n$. The value function and optimal investment policy at step 0 are given by the boundary condition.

The value function at any particular stage $j > 0$, $V(K_t, j)$ can be determined by an iterative process if we know the value function and optimal investment process at the the previous step, $V(K_t, j-1)$ and $I(K_t, j-1)$. In order to do this, we define an iterative operator which has the contraction properties needed to reach convergence. The first subsection lays down the theoretical foundations under which our algorithm will be based, followed in the next subsection by a description of the algorithmic procedure used.

3.1 Theoretical Foundation

First, we start by carrying out a change of variable to make the implementation easier. Indeed, since K_t is lognormal and in order to take advantage of the independence of increments of the normal distribution, we work in the log space. Letting $k_t = \ln K_t$, by Ito's Lemma

$$dk_t^{j,0} = \left(r_f - \hat{\delta}_k^{\lambda_z}(0, j) - \frac{\sigma_k^2}{2} \right) dt + \sigma_k dw_t^2 \tag{12}$$

$$dk_t^{j,I} = \left(r_f \hat{\delta}_k^{\lambda_z}(0, j) - \hat{\delta}_k^{\lambda_z}(I, j) - \frac{\sigma_k^2 + \sigma_z^2 I^2 j^{2\gamma}}{2} \right) dt + \sigma_k dw_t^2 + \sigma_z I j^\gamma dz_t \tag{13}$$

It is clear from equations (12) and (13) that the increments of $k_t^{j,0}$ and $k_t^{j,I}$ are independent. This independente increment characteristic of the state variable will be of major importance in the theoretical analysis that follows. For notational simplicity we use the same letters V and I used before to denote the value function $V(k_t, j)$ and optimal investment decision $I(k_t, j)$:

$$V(k_t, j) = \max\{-I + \alpha E[V(k_{t+\Delta T}^{j,I}, j-1)], \alpha E[V(k_{t+\Delta T}^{j,0}, j)]\} \tag{14}$$

The boundary condition given in (9) is also adjusted accordingly.

We will work in the measure space $L_2(\Re^d, B(\Re^d), \lambda)$, where $B(\Re^d)$ and λ denote the Borel σ-field and the Lebesgue measure respectively. We construct a Hilbert Space on L_2 by defining the following inner product : $< J, \widetilde{J} >= \int J(x)\widetilde{J}(x)dx$, where $J, \tilde{J} \in L_2$. The norm induced by this inner product is defined by $\|J\|^2 = \int J^2(x)dx$.

In order to solve for the value function at a given stage j, $V(k_t, j)$, let us define for any function $J \in$

$L_2(\Re^d, B(\Re^d), \lambda)$, a family of dynamic programming operators $\left(T^j\right)$ for $j = 1, ...n$ by:

$$T^j J(k_t, j) = \max\left\{-I + \alpha E[V(k^{j,I}_{t+\Delta T}, j-1)], \alpha E[J(k^{j,0}_{t+\Delta T}, j)]\right\} \tag{15}$$

where the diffusions of $k^{j,I}_{t+\Delta T}$ and $k^{j,0}_{t+\Delta T}$ are given by equations (12) and (13).More explicitly,

$$T^j J(x, j) = \max\left\{-I + \alpha \int V(x + z_1, j-1) f(z_1) dz_1, \alpha \int J(x + z_2, j) f(z_2) dz_2\right\} \tag{16}$$

where $f(.)$ is the normal distribution density and z_1 and z_2 are normal random variables distributed as follows:

$$z_1 \sim N\left(\left(r_f - \hat{\delta}^{\lambda_z}_k(I, j) - \frac{(\sigma_k^2 + \sigma_z^2 I^2 j^{2\gamma})}{2}\right)\Delta T, (\sigma_k^2 + \sigma_z^2 I^2 j^{2\gamma})\Delta T\right)$$

$$z_2 \sim N\left(\left(r_f - \hat{\delta}^{\lambda_z}_k(0, j) - \frac{\sigma_k^2}{2}\right)\Delta T, \sigma_k^2 \Delta T\right)$$

These values are consistent with equations (12) and (13). Note also that the integral is taken over the entire real line, so we are working in an infinite dimensional space.

We will show that the operators T^j are contractions in the $L_2(\Re^d, B(\Re^d), \lambda)$ under the norm previously defined.

Proposition 1 *Given $J(k_t, j-1)$, the operator T^j is a contraction under the introduced λ-norm, i.e $\|T^j J - T^j \widetilde{J}\| \leq \alpha \|J - \widetilde{J}\|$ for any J and $\widetilde{J}$ in $L_2(\Re^d, B(\Re^d), \lambda)$ and $\alpha \in (0, 1)$. Moreover, for any function $\widetilde{J}$ the sequence $T^j_m(\widetilde{J})$ converges to the value function $V(k_t, j)$ as $m \to +\infty$.*

Proof. First note that once we know $V(k_t, j-1)$ we could think as the entire first term on the max of (15) as a deterministic function of k_t and j, which we will label $g(k_t, j)$. In other words $g(., j) = -I + \alpha E[V(., j-1)]$. Thus, (15) and (16) and simplify to:

$$\begin{aligned} T^j J(k_t, j) &= \max\left\{g(k_t, j), \alpha E[J(k^{j,0}_{t+\Delta T}, j)\right\} &(17)\\ T^j J(x, j) &= \max\left\{g(x, j), \alpha \int J(x + z_2, j) f(z_2) dz_2\right\} &(18)\end{aligned}$$

Let $PJ = \int J(x + z_2, j) f(z_2)$ be the expected value of the function after one period. Let $z_2 = z$ for notational simplicity. We can establish that $\|PJ\|^2 \leq \|J\|^2$ since:

$$\|PJ\|^2 = \int (PJ)^2(x) dx = \int \left[\int J(x + z, j) f(z) dz\right]^2 dx$$

But applying Jensen's Inequality and Fubini's Theorem respectively

$$\begin{aligned} \int \left[\int J(x + z, j) f(z) dz\right]^2 dx &\leq \int \int J^2(x + z, j) f(z) dz dx \\ &= \int \left[\int J^2(x + z, j) dx\right] f(z) dz \end{aligned}$$

Now conditioning on z, $\int J^2(x+z,j)dx = \int J^2(x,j)dx$ and hence

$$\begin{aligned}\int \left[\int J(x+z,j)f(z)dz\right]^2 dx &\leq \int \|J\|^2 f(z)dz \\ &= \|J\|^2\end{aligned}$$

Thus,

$$\|PJ\|^2 \leq \|J\|^2.$$

Now, we are ready to show that our operator T^j is a contraction. Take any two arbitrary value functions J and $\widetilde{J}$. We have that:

$$\|T^j J - T^j \widetilde{J}\|^2 = \int \left[\max\left\{g(x,j), \alpha \int J(x+z,j)f(z)dz\right\} - \max\left\{g(x,j), \alpha \int \widetilde{J}(x+z,j)f(z)dz\right\}\right]^2 dx$$

But note that $|\max\{a,b\} - \max\{a,c\}| \leq |b-c|$ so that the previous expression becomes:

$$\|T^j J - T^j \widetilde{J}\|^2 \leq \int \left[\alpha \int \left(J(x+z,j) - \widetilde{J}(x+z,j)\right) f(z)dz\right]^2 dx = \alpha^2 \|P(J-\widetilde{J})\|^2$$

Now, since we have just shown that $\|PJ\|^2 \leq \|J\|^2$ for any J this implies that: $\|T^j J - T^j \widetilde{J}\| \leq \alpha \|J - \widetilde{J}\|$ completing our proof. ■

Since T^j is a contraction then the sequence $T^j_m J$ converges to a unique fixed point as m goes to infinity. This unique fixed point is the solution to Bellman equation: $J = T^j J = \max\{g(.,j), \alpha PJ\}$, which is nothing more than equation (17), and it is equal to the value function $V(k_t, j)$ at step j. However, since the space L_2 is infinite dimensional we need to resort to an approximation of the value function. In this paper, we parameterize the value function in the following form :

$$\hat{V}(x,j,r) = \sum_{k=1}^{q} r_{j,k}\phi_k(x)$$

where $\Phi = (\phi_1,\ \phi_2,...,\ \phi_q) \in L_2(\Re^d, B(\Re^d), \lambda)$ is the vector of "basis functions" and $r_j = (r_{j,1},...,r_{j,q})$ is a vector of scalar weights. We perform the iteration over the subspace of L_2 generated by the linearly independent basis.

We now define Π to be the operator that projects any function in $L_2(\Re^d, B(\Re^d), \lambda)$ onto the subspace spanned by Φ. For each value function J, we define $\Pi J = \hat{r}.\Phi$ to be the projection of J over the space generated by the basis functions Φ, where $\hat{r}$ is the q -dimensional vector of coefficients of the q linearly independent basis functions that solves:

$$\hat{r} = arg \min_{r \in \Re^q} \|J - r.\Phi\|^2$$

Let Ψ^j be the approximate operator for the value function with j stages to go. It is defined recursively as :

$$\Psi^1 J = \Pi T^1 J \text{ for } j = 1 \tag{19}$$

$$\Psi^j J = \Pi \hat{T}^j J \text{ for } j > 1 \tag{20}$$

with $\hat{T}^j$ for $j > 1$ given by

$$\hat{T}^j J(k_t, j) = \max\left\{-I + \alpha E[\hat{V}(k^{j,I}_{t+\Delta T}, j-1)], \alpha E[J(k^{j,0}_{t+\Delta T}, j)]\right\}$$

Note that in defining $\hat{T}^j$ for $j > 1$ we are not using the previous true value function, but its approximation value, which we denote by $\hat{V}(., j-1)$. Since the value function is known at the expiration date, we use the "true" T^1 for the definition of the approximate operator at the first stage. We are now ready to prove the following proposition.

Proposition 2 *Given $J(k_t, j-1)$, Ψ^j is a contraction under the λ -norm, i.e., $\|\Psi^j J - \Psi^j \widetilde{J}\| \leq \alpha \|J - \widetilde{J}\|$ for any J and $\widetilde{J} \in L_2$. Moreover, starting from any basis approximation $r^0_j \Phi$ the sequence $r^m_j \Phi$ generated by using the approximate iterator converges to a fixed point $\hat{V}(., j) = r^\star_j.\Phi$ as $m \to +\infty$.*

Proof. First note that we can mimic the proof of Proposition 1 to show that the operator $\hat{T}^j$ is a contraction under the introduced λ-norm. All the steps given there will go through using a different deterministic function $\hat{g}(k_t, j)$ that will now be defined by the approximate value function $\hat{V}(k_t, j-1)$ instead of the true value $V(k_t, j-1)$. Thus, for any J and $\widetilde{J} \in L_2$, $\|\hat{T}^j J - \hat{T}^j \widetilde{J}\| \leq \alpha \|J - \widetilde{J}\|$ with $\alpha \in (0,1)$. On the Hilbert Space defined on L_2, we know that for any $J \in L_2$, ΠJ and $J - \Pi J$ are orthogonal. By the Pythagorean theorem:

$$\|J\|^2 = \|\Pi J\|^2 + \|J - \Pi J\|^2$$

so that $\|\Pi J\|^2 \leq \|J\|^2$. For any J, $\widetilde{J} \in L_2$ and $j > 1$ we have:

$$\|\Psi^j J - \Psi^j \widetilde{J}\|^2 = \|\Pi(\hat{T}^j J - \hat{T}^j \widetilde{J})\|^2 \leq \|\hat{T}^j J - \hat{T}^j \widetilde{J}\|^2 \leq \alpha^2 \|J - \widetilde{J}\|^2$$

where the last inequality is the result of $\hat{T}^j$ being also a contraction. Since $\Psi^j_m = r^m_j \Phi$ then the sequence $r^m_j \Phi$ converges to a unique fixed point as m goes to infinity. ■

We have shown above that the approximate value iteration will converge to a fixed point, which will give us the approximate value function. However, it is important to provide bounds on the errors obtained and see how these errors compound as we increase the number of stages. This is provided by following proposition.

Proposition 3 *Let $\hat{V}(k_t, j) = r^\star_j.\Phi(k_t)$ be the approximate value function obtained at each given stage j. Then, the approximation error at any stage j is bounded by:*

$$\|V^j - \hat{V}^j\|^2 \leq \sum_{i=0}^{j-1} \frac{\alpha^{2i}}{(1-\alpha^2)^{i+1}} \|V^{j-i} - \Pi V^{j-i}\|^2 \tag{21}$$

Proof. : First for notational simplicity let $V(.,j) = V^j$ and $\hat{V}(.,j) = \hat{V}^j$. Decomposing the error difference in two components, one parallel and one orthogonal to Φ, we get by the Pythagorean Theorem:

$$\|V^j - \hat{V}^j\|^2 = \|V^j - \Pi T^j V^j\|^2 + \|\Pi T^j V^j - \hat{V}^j\|^2 \tag{22}$$

Adding and substracting $\Pi \hat{T}^j V^j$, and using the fact that the approximated value function is a fixed point of the approximate operator Ψ^j, i.e., $\hat{V}^j = \Psi^j \hat{V}^j = \Pi \hat{T}^j \hat{V}^j$, we can, by means of the triangle inequality, bound the second sum of the previous expression:

$$\|\Pi T^j V^j - \hat{V}^j\|^2 \leq \|\Pi T^j V^j - \Pi \hat{T}^j V^j\|^2 + \|\Pi \hat{T}^j V^j - \Pi \hat{T}^j \hat{V}^j\|^2 \tag{23}$$

Since $\hat{T}^j$ is a contraction and Π is non-expanding the second term is trivially bounded by $\alpha^2 \|V^j - \hat{V}^j\|^2$. The proof is complete if we can provide a bound on the first term. Going back to our definitions we have :

$$\|TV^j(x) - \hat{T}^j V^j(x)\|^2 = \int \left[\max\left\{ g(x,j), \alpha \int V(x+z,j) f(z) dz \right\} - \max\left\{ \hat{g}(x,j), \alpha \int V(x+z,j) f(z) dz \right\} \right]^2 dx$$

Note again that the only difference lies on the first term of the max operator. For the true T operator we use a deterministic $g(k_t, j)$ based on the correct value function at the previous stage V^{j-1} whereas for $\hat{T}$ we used a different function $g(k_t, j)$ based on its approximate value function V^{j-1}. In other words: $g(k_t, j) = -I + \alpha E[V(k_{t+\Delta T})]$ and $\hat{g}(k_t, j) = -I + \alpha E[\hat{V}(k_{t+\Delta T})]$. Using the fact that $|\max\{a,b\} - \max\{a,c\}| \leq |b - c|$, as in the proof for Proposition 1, we can simplify the following expression to:

$$\begin{aligned}
\|T^j V^j(x) - \hat{T}^j V^j(x)\|^2 &\leq \|g(x,j) - \hat{g}(x,j)\|^2 \leq \int \left[\alpha \left(\int V^{j-1}(x+z) - \hat{V}^{j-1}(x+z) \right) f(z) dz \right]^2 dx \\
&\leq \int \left[\alpha \left(\int V^{j-1}(x+z) - \hat{V}^{j-1}(x+z) \right)^2 f(z) dz \right] dx \\
&\leq \int \alpha^2 \left[\int \left(V^{j-1}(x+z) - \hat{V}^{j-1}(x+z) \right)^2 dx \right] f(z) dz \\
&\leq \alpha^2 \int \|V^{j-1}(x) - \hat{V}^{j-1}(x)\|^2 f(z) dz \leq \alpha^2 \|V^{j-1}(x) - \hat{V}^{j-1}(x)\|^2
\end{aligned}$$

where we have used the definitions of $g(.,j)$ and $\hat{g}(.,j)$, Jensen's Inequality and conditioning over a fix z respectively. Thus, as Π is non-expansive this implies $\|\Pi T^j V^j - \Pi \hat{T}^j V^j\|^2 \leq \alpha^2 \|V^{j-1} - \hat{V}^{j-1}\|^2$. Replacing this result together with (23) in (22) we get:

$$\|V^j - \hat{V}^j\|^2 \leq \|V^j - \Pi V^j\|^2 + \alpha^2 \|V^{j-1} - \hat{V}^{j-1}\|^2 + \alpha^2 \|V^j - \hat{V}^j\|^2$$

Rearranging the terms we can get a bound for the approximation error at any stage j at a function of the error at the previous stage.

$$\|V^j - \hat{V}^j\|^2 \leq \frac{1}{1-\alpha^2} \|V^j - \Pi V^j\|^2 + \frac{\alpha^2}{1-\alpha^2} \|V^{j-1} - \hat{V}^{j-1}\|^2$$

However, we know the correct value function when $j = 0$ and hence we can expand the previous recursive inequality until we reach the stage 0 and obtain expression (21). ■

3.2 The Algorithm

Recall that the Bellman equation we would like to solve is given by :

$$V(k_t, j) = \max\left\{-I + \alpha E[V(k^{j,I}_{t+\Delta T}, j-1)], \alpha E[V(k^{j,0}_{t+\Delta T}, j)]\right\}$$

To ease the notation it is useful to introduce two Q-functions, the exercise value $Q_e(k_t, j)$ and the continuation value $Q_c(k_t, j)$ for any $j, k_t \geq 0$. They are defined respectively as:

$$\begin{aligned} Q_e(k_t, j) &= -I + \alpha E[V(k^{j,I}_{t+\Delta T}, j-1)] \\ Q_c(k_t, j) &= \alpha E[V(k^{j,0}_{t+\Delta T}, j)] \end{aligned}$$

Notice that in implementing the algorithm, we immediately face with the problem of the unboundedness of the state space. In order to work in a bounded space while maintaining the theoretical properties shown above, we must work with basis functions that vanish outside the bounded interval $[-D, D]$ for $D > 0$ large enough that guarantees a good approximation for the value function in our range of interest. In general, the initial choice of D should depend on the parameters k_0, the number of pilot steps and the diffusion and drift terms for the process of k_t in the risk neutral measure.

To solve this Bellman equation, we start by computing the value function at the end of the pilot stage (when $j = 0$) and then compute the current value function by backward induction as described below.

- At the end of the pilot project period, depending on the value of k_T (recall that T is the arbitrary time required to complete the N steps of the pilot stage), the company needs to decide whether to launch the commercial project immediately or wait for better market conditions, i.e., lower implementation costs. Therefore, the value function at the end of the pilot stage, $V(k_t, 0)$, is found by replacing K_T for $\exp(k_t)$ in (9).

  We generate an n-dimensional uniform vector $k_T = \left(k^1_T, ..., k^n_T\right)$ with each component uniformly distributed on the grid $[-D, +D]$. The higher the number of simulated points the best the value function approximation obtained. Now, recall that in the previous section, we assumed the following parametrization of the value function

  $$\hat{V}(k_t, j) = \sum_{i=1}^{q} r_{j,i}\phi_i\left(k_t\right)$$

  To find the vector of scalar weights $r_0 = (r_{0,1}, ..., r_{0,q})$, we compute the vector $V\left(k_T, 0\right)$ using equation (9) and then regress this vector on the basis functions $\Phi = (\phi_1,\ \phi_2, ...,\ \phi_q)$.

- After obtaining the vector of scalar weights r_0, the exercise value function with one step to go is trivially computed as

$$Q_e\left(k_{T-\Delta T},1\right)=-I+\alpha\sum_{k=1}^{q}r_{0,i}E\left[\phi_i\left(k_T\right)\right]$$

  Note that the value functions vanish at any sample path in which k_T falls outside the range $[-D,+D]$. As in step (1), the vector $k_{T-\Delta T}=\left(k_{T-\Delta T}^1,...,k_{T-\Delta T}^n\right)$ is obtained by simulating an n-dimensional vector, with each component uniformly distributed along the grid $[-D,+D]$. Also, for the continuation value $Q_c\left(k_{T-\Delta T},1\right)$, we assume that it is parameterized as :

$$Q_c(k_{T-\Delta T},1)=\sum_{i=1}^{q}\alpha r_{1,i}E_{T-\Delta T}\left[\phi_i\left(k_T\right)\right]$$

  where $E_{T-\Delta T}$ is the risk neutral expected value conditional on the information available at $T-\Delta T$, which, since our process is Markovian, is equivalent to the expectation conditional on knowing $k_{T-\Delta T}$.

  To compute the value of the vector of scalar weights r_1, we start by assuming that it is equal to some arbitrary vector $r_1^0=\left(r_{1,1}^0,...,r_{1,q}^0\right)$ and then compute $Q_c(k_{T-\Delta T},1,r_1^0)=\alpha\sum_{k=1}^{q}r_{1,i}^0E_{T-\Delta T}\left[\phi_i\left(k_T\right)\right]$. Note that a third argument in the Q-function continuation value is used to stress that it is dependent on the choice of vector r_1^0. Hence the value function is now given by :

$$\hat{V}(k_{T-\Delta T},1,r_1^0)=\max\left\{Q_e\left(k_{T-\Delta T},0\right),Q_c(k_{T-\Delta T},1,r_1^0)\right\}$$

  Now that we have the value function, the vector coefficient $r_1^1=\left(r_{1,1}^1,...,r_{1,q}^1\right)$ is found by regressing $\hat{V}(k_{T-\Delta T},1,r_1^0)$ on the vector $\Phi\left(k_{T-\Delta T}\right)=\left(\phi_1\left(k_{T-\Delta T}\right),\phi_2\left(k_{T-\Delta T}\right),...,\phi_q\left(k_{T-\Delta T}\right)\right)$. We then replace these coefficients in the continuation value to obtain $Q_c(k_{T-\Delta T},1,r_1^1)=\sum_{k=1}^{q}\alpha r_{1,i}^1E\left[\phi_i\left(k_T\right)\right]$ and find the vector coefficient r_1^2 through a further regression. The procedure is repeated until $\|r_1^{l+1}-r_1^l\|<\varepsilon$ for $l\geq 0$ where ε is a criteria of maximal error.

- In the general case for any $j>0$ investment opportunities, the two steps are repeated as follows. At this point of the algorithm we have already computed the vector coefficients r_{j-1}, therefore the exercise value is given by :

$$Q_e\left(k_t,j\right)=-I+\alpha\sum_{i=1}^{q}r_{j-1,i}E\left[\phi_i\left(k_{t+\Delta T}\right)\right]$$

  Again $k_t=\left(k_t^1,...,k_t^n\right)$ is obtained by generating an n-dimensional uniform vector on $[-D,+D]$. Also to compute the continuation value $Q_c\left(k_t,j\right)$, we assume that it is parameterized as following

$$Q_c(k_t,j)=\sum_{i=1}^{q}\alpha r_{j,i}E\left[\phi_i\left(k_{t+\Delta T}\right)\right]$$

  In a similar manner to the second step of the algorithm, the coefficients vector r_j are found by iteration as follows: We start with an initial arbitrary value r_j^0 and then compute the approximated continuation value:

$$Q_c(k_t, j, r_j^0) = \sum_{k=1}^{q} \alpha r_{j,i}^0 E_t \left[\phi_i \left(k_{t+\Delta T}\right)\right]$$

We then compute the approximate value function:

$$\hat{V}(k_t, j, r_j^0) = \max \left\{ Q_e \left(k_t, j-1\right), Q_c(k_t, j, r_j^0) \right\}$$

and find the vector of coefficients r_j^1 by regressing $\hat{V}(k_t, j, r_j^0)$ on the vector $\Phi\left(k_t\right) = (\phi_1\left(k_t\right), \phi_2\left(k_t\right), ..., \phi_q\left(k_t\right))$. Finally, we compute $Q_c(k_t, j, r_j^1)$ and $\hat{V}(k_t, j, r_j^1)$ by replacing the vector r_j^0 by the vector r_j^1 in the expressions above and regress again to find the next iterate r_j^2. We continue this iterative procedure until $\|r_j^{l+1} - r_j^l\| < \varepsilon$ for $l \geq 0$ where ε is a criteria of maximal error.

- The above procedure is repeated until $j = N$. At this point we have found all $(N+1)$ q-dimensional vector of coefficients $r_j = (r_{j,1}, r_{j,2}, ..., r_{j,q})$ for $j = 0, .., N$. This provides an approximate value function $\hat{V}(k_t, j)$ at any given number of steps for completion and for any given value of the cost process.

4 Numerical Solution and Economic Analysis

4.1 Approximate Value Function and Investment Thresholds

We present in this section a simple example and give some interpretations of the obtained results. In this example we set the parameters to the following values : $N = 5$, $\bar{S} = 100$, $\Delta T = .5$, $\sigma_M = 0.2$, $r_f = 0.05$, $I = 5$, $\gamma = 1$, $\sigma_z = 0.008$, $\beta = 0.009$, $\lambda_z = 0.1$. More explicitly, the firm makes semiannual funding decision and the pilot stage requires 5 steps for completion, each of them requiring an investment of 5 units of money (UM). After the pilot stage is completed a fixed payment of 100 UM could be received. By setting $\gamma = 1$ the learning coefficient is taken to be greater than 0 implying higher learning at earlier steps of the pilot stage. Since the total technical uncertainty is given by $Ij^{\gamma}\sigma_z$ we make choose σ_z such that the overall technical volatility at the first step, i.e., $j = 5$ is around 20%.

Table (3) shows the values of the investment opportunity at each step j and for a range of costs going from 40 to 110. The value functions as a function of K_t are plotted in figure (1) for each given j.

We could immediately verify that the investment opportunity is decreasing in K_t at all stages. This effect does not deserve further explanation. Another fairly intuitive effect is that for fixed costs the value of the investment opportunity is decreasing in the number of stages remaining for completion, since there is an extra fixed payment required to receive the final payoff. Moreover, we can also verify that for fixed K_t the difference of the values of the option at two consecutive steps, i.e., $V(K_t, j-1) - V(K_t, j)$, is decreasing in j. This could be explained by the fact that the extra payment of I that has to made in $V(K_t, j)$ is discounted

Table 1: Value of investment opportunity $V(K_t, j)$ as a function of expected costs and number of steps remaining for completion

K — j	**0**	**1**	**2**	**3**	**4**	**5**
40	60	52.7591	46.0756	39.67684	33.79725	28.44393
45	55	47.64556	40.95205	34.61529	28.90278	23.80467
50	50	42.53202	35.8285	29.55373	24.0083	19.16541
55	45	37.41848	30.70494	24.49216	19.11381	14.52614
60	40	32.30493	25.58137	19.43059	14.21932	9.886861
65	35.00013	28.99098	22.80878	17.13183	12.37619	8.665299
70	30.51426	25.19435	19.81973	14.88551	10.7538	7.52747
75	26.85638	22.0985	17.38238	13.05381	9.430863	6.599658
80	23.83266	19.53937	15.36759	11.53967	8.33728	5.832697
85	21.30313	17.3985	13.6821	10.27299	7.422431	5.191088
90	19.16463	15.58858	12.25716	9.202125	6.649005	4.648662
95	17.33972	14.04406	11.04117	8.288292	5.988994	4.185778
100	15.76933	12.71496	9.994777	7.501913	5.421036	3.787453
105	14.40776	11.5626	9.087526	6.8201	4.9286	3.442094
110	13.21918	10.55664	8.295541	6.224911	4.498728	3.140612

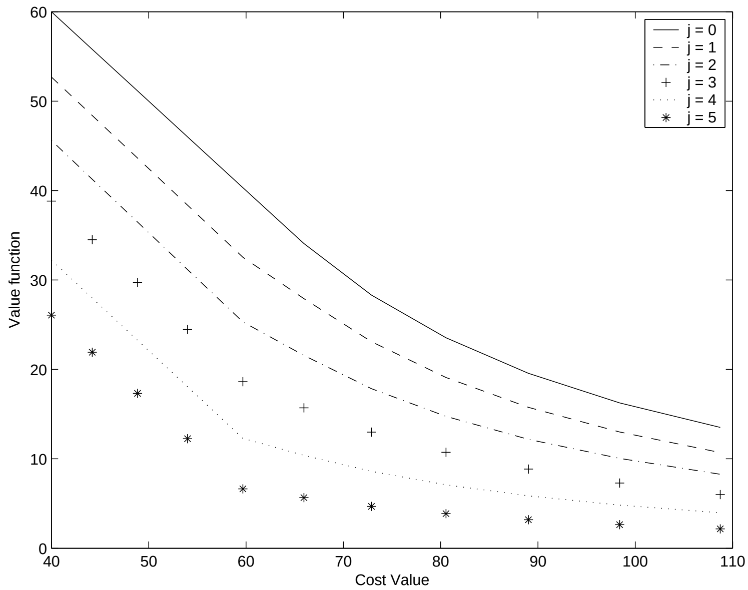

Figure 1: Plot of Value function with respect to number of investments to go

with a larger time horizon as we increase the number of steps (move away from completion). However, for an unusually high level of technical volatility, strongly decreasing learning with respect to investment across pilot steps and low I, this property may break down, as the difference between technical learning at two consecutive steps may counteract the previous effect.

Let us denote by $\bar{K}^j$ the optimal exercise threshold with j stages to go, i.e., the value for expected costs such that if $K_t > \bar{K}^j$ the firm will decide not to fund the pilot project and if $K_t < \bar{K}^j$ the firm will proceed with the required investment and advance to the following stage. This threshold value is found by solving :

$$-I + \alpha E[\hat{V}(K_{t+\Delta T}, j-1) = \alpha E[\hat{V}(K_{t+\Delta T}, j)] \quad (24)$$

Table (2) shows the threshold below which investment is optimal as a function of the number of stages remaining for completion.

Table 2: Thresholds for each step

j	**1**	**2**	**3**	**4**	**5**
K^j	74.32293	73.30621	71.14852	66.7037	62.25501

We can verify that, as our intuition suggests, the exercise threshold decreases with the number of stages for completion. This has an appealing explanation that we can related to the fact that $V(K_t, j-1) - V(K_t, j)$ is decreasing in j. Investing at step $j-1$ one will acquire the option $V(., j-2)$ in ΔT years from now. Doing so at step j one will acquire $V(., j-1)$. The option to invest is decreasing in j, implying a higher incentive to invest at $j-1$. However, by not investing at $j-1$ one keeps the option $V(., j-1)$ at $t+\Delta T$ and, by doing so at $j-2$ one keeps $V(., j-2)$. Thus, one has also more incentive not to invest in j with respect to $j-1$. Nevertheless, if the difference of the option value at two consecutive stages is decreasing in j, the higher incentive to invest at step $j-1$ with respect to j will be relatively more important than the higher incentive to continue. Thus, at step $j-1$, the firm can afford to have a higher indifference threshold than at a previous step.

4.2 Comparative Statics

In this section, we provide additional economic insights, by studying how changes in the main parameters driving our model affect the value of the investment opportunity and the decision to invest.

As the distinction between market and technical uncertainty is one of the key features of our setting it will be useful to compare how they affect the solution to our model.

Figure (2) shows the joint impact of market and technical risk by plotting together for the case $j = 1$ four different combinations corresponding to scenarios of high and low market uncertainty, as well as high and low technical uncertainty. Table (3) shows the investment threshold as a function of j for each of these combinations. The market uncertainty is taken to be 10% and 25% for the low and high scenarios respectively. σ_z is taken so that the technical uncertainty resolved at the last stage ($j = 5$) is around 25% and 40% for the low and high scenarios respectively.

From the figure, we can see that the value of the option to invest increases in both market and technical uncertainty. This a basic property of any option, consequence of the convexity of its payoff function. More

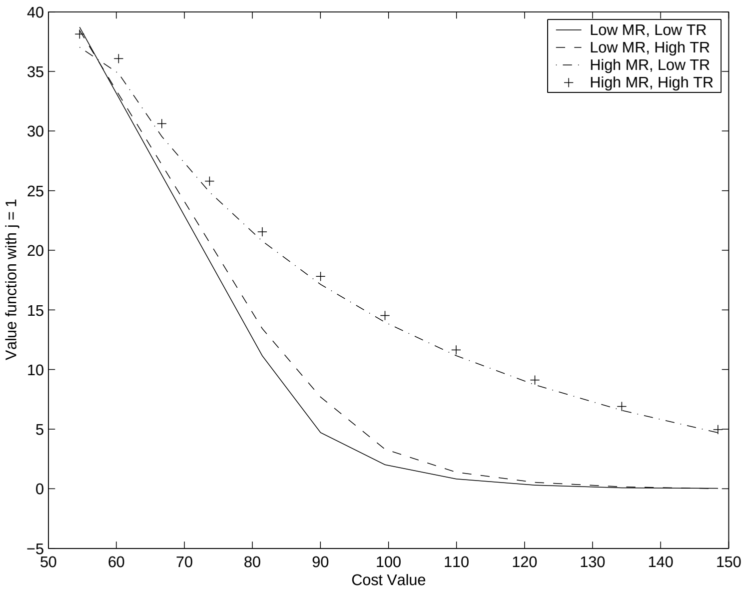

Figure 2: Sensitivity of The Value function with One investment to go with respect to Market Risk (MR) and Technical Risk (TR).

Table 3: Thresholds for each step and different market and technical risks

MR-TR — j	**1**	**2**	**3**	**4**	**5**
Low-Low	89.88019	87.71959	87.0693	87.16034	86.80213
Low-High	90.21455	89.31893	89.88441	90.74021	89.76489
High-Low	77.90904	76.62582	71.46393	64.3008	58.39973
High-High	78.82331	77.71367	73.43698	69.00896	67.78572

interestingly, this effect is much stronger for market than for technical uncertainty, since market uncertainty will always be present, while technical uncertainty requires an investment to be resolved. The table of thresholds complements our intuition. At any given stage j we can verify that the investment threshold is decreasing in market uncertainty and increasing in technical uncertainty. For a risk averse firm, market uncertainty makes a firm more reluctant to undertake irreversible investment. Thus, it requires a higher incentive to do so. This results in a lower indifference threshold. As for technical uncertainty, the history is dramatically different. Technical uncertainty benefits the firm only if it can invest to take advantage of it. Thus, higher technical uncertainty results in a higher incentive to invest and a higher indifference threshold. This results are consistent with those obtained by Pindyck (1993) in a slightly different setting that did not differentiate between a pilot and a commercial stage.

Figure (3) shows the effect of the intrinsic reduction in expected costs due to technical investment (the parameter β) in the value of the option and optimal investment threshold. A higher β implies that each additional monetary unit invested in a given pilot step will result in a greater reduction in expected cost.

Thus, the value of the investment opportunity and the threshold value at which is optimal to start investment, are both increasing in β.

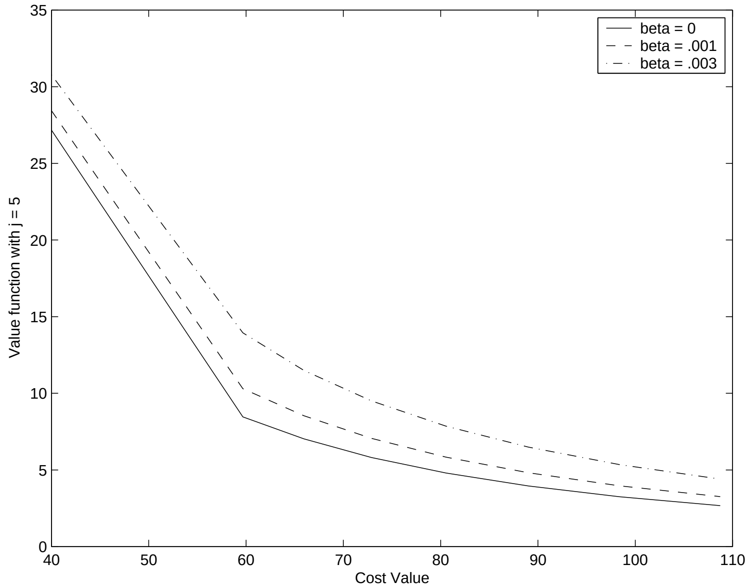

Figure 3: Sensitivity of The Value function with One investment to go with respect to Market Risk (MR) and Technical Risk (TR).

Finally, we analyze the effects of changing the time between funding decisions, ΔT in our result. Figures (4) and (5) draw the optimal solution as a function of Δ_T for one and three steps to go respectively.

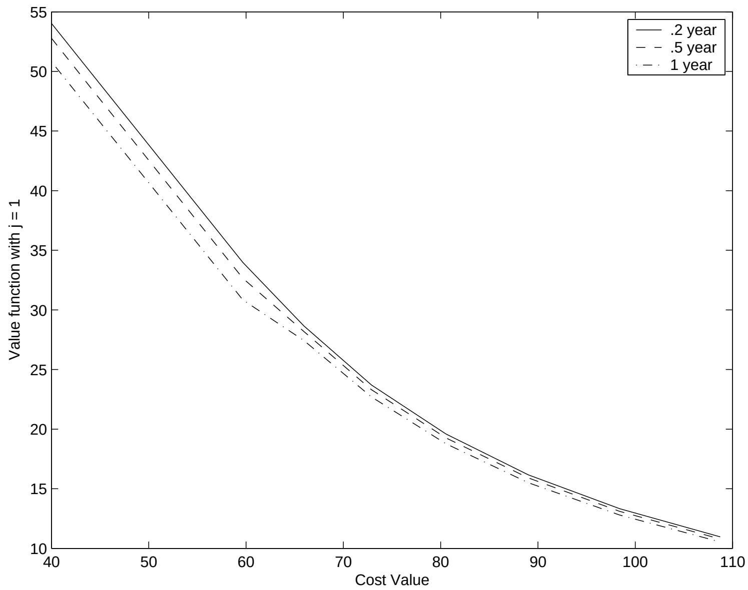

Figure 4: Sensitivity of The Value function with One investment to go with respect to change of time step

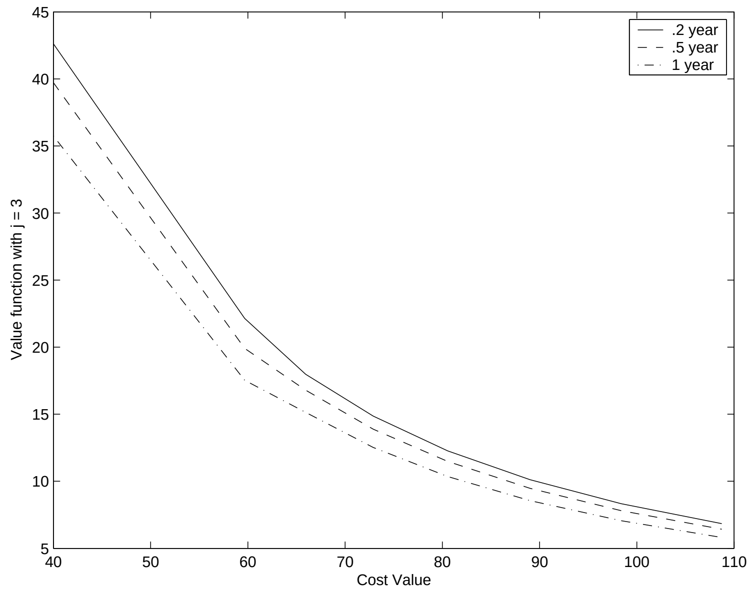

Figure 5: Sensitivity of The Value function with One investment to go with respect to change of time step

Speeding up time to completion, by reducing up the time between pilot increases the value of the option, as it puts the final payoff closer in time. This increase is more dramatic when K_t is low and the commercial stage has a high probability of being undertaken. Even when costs are high and investment is not currently optimal a decrease in ΔT increases the value of the investment opportunity by allowing the firm to monitor market information more frequently. Of course, decreasing this time interval may not be feasible for the firm beyond certain point, as technical considerations will put a limit on the minimum time that a pilot step could take to be completed.

5 Conclusions

We have proposed a general modeling approach to value sequential investments where the underlying assets are subject to two types of risk: market and technical uncertainty. Each affect the value of our underlying in different ways. Market uncertainty is generally related to economic fundamentals and always driving the value of a project. Technical uncertainty is generally private to the firm and only affects the value of a project when the firm invests in activities whose purpose is to reduce this technical risk, i.e., pilot projects. We frame the problem as a perpetual compound option. Spanning assumptions are made on market uncertainty and technical uncertainty is parameterized with a ”market price” of technical risk that allows for different risk specifications to value the non tradable payoffs. In our model technical uncertainty only affects costs, while market risk influences both, the revenues and costs of the underlying commercial stage. The technical volatility and drift terms of the stochastic process for costs are assumed to satisfy some restrictions that we relate with economic assumptions governing the learning process.

We work in an equivalent risk neutral measure that is uniquely characterized by the market price of technical risk, and thus, are able to draw on financial engineering techniques to solve our problem. Using Approximate Dynamic Programming we solve the Bellman Equation resulting from a simplified version of our general setting. To solve for the value function at each state we define a fix point operator, which in turn depends on the approximate value function at the previous stage. By adapting a set up suggested by Van Roy (1998) to handle problems in an infinitely dimensional state space with independent increments, we prove that our operators are contractions. This guarantees the convergence of our algorithm. We provide bounds on the maximum permissible errors at a given stage, showing how the maximum approximation errors compound as we add stages in the compound option. This last result might be new to the literature.

Interesting economic insights are obtained from our model. Among the most important ones is that market and technical have different effects over our results. Although both increase the value of the option to invest (the former more significantly than the latter), they have opposite effects on the investment decisions: higher technical uncertainty results in an incentive to invest, while the contrary occurs for higher market uncertainty. This is consistent with results previously obtained in the literature.

Solving for the general setup with a variable investment level at each stage, may yield additional economic insights about the tradeoff between the benefits of investing to reduce technical uncertainty and its cost. Allowing for a more general range of learning diffusion coefficients may also do so. However, tractability may be greatly sacrificed on the way. Focusing on learning in the commercial stage, i.e., ”learning by doing”, might be a natural complement of the work presented here. All of the above are left for further research.

Stanford, California.

January 2005

References

[1] Bertsekas, D.P., Dynamic Programming and Optimal Control, Athena Scientific, 1995.

[2] Bertsekas, D.P. and Tsitsiklis, J.N., Neuro Dynamic Programming, Athena Scientific, 1996

[3] Brach M. and Paxson D., A Gene to Drug Venture: Poisson Options Analysis, *R&D Management*, vol 31 No 2 203-214, 2001.

[4] Cortazar G., et. al, Optimal Exploration Investments under Price and Geological-Technical Uncertainty: a real options model, *R&D Management*, v31 No 2 181-189, 2001.

[5] Duffie D., Dynamic Asset Pricing Theory, *Princeton University Press*, 2001.

[6] Dixit, A and Pindyck, R., Investment Under Uncertainty, Princeton University Press, 1994.

[7] Grossman G. and Shapito C., Optimal Dynamic R&D Programs, *Rand Journal of Economics,* Vol 17 581-593, 1986.

[8] Lucas R. E. Jr., Optimal Management of a Research and Development Project, *Management Science*, Vol 17 No 11 679-697, 1971.

[9] Margrabe W., The Value of an Option to Exchange One Asset for Another, *Journal of Finance*, 33, 177-86 (1978)

[10] Majid S. and Pyndick R. S., Time to Build, Option Value and Investment Decisions, *Journal of Financial Economics* 18, 7-27, 1987.

[11] Merton R.C.,Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141-183, 1973.

[12] Messica A. and David I., Optimal expenditure patterns for risky R&D projects with time dependent returns, *R&D Management*, vol 30 No. 3 247-253, 2000.

[13] McDonald R. and Siegel D., The Value of Waiting to Invest, *Quarterly Journal of Economics,* 101 707-728, 1986.

[14] Pyndick R. S., Investments of Uncertain Cost, *Journal of Financial Economics* 34, 53-76, 1993.

[15] Roberts K. and Weitzman M. L. , Funding Criteria for Research, Development and Exploration Project, *Econometrica* 49 p. 1261-1288, 1981.

[16] Schwartz E. and Zoraya C., Investment under uncertainty in Information Technology, Acquisition and Development Projects, Management Science, vol 49 57-70, 2003.

[17] Van Roy B., Learning and Value Function Approximation in Complex Decision Processes, PhD Thesis, *Massachusetts Institute of Technology*, May 1998.