

Learning Options and Binomial Trees

Graeme Guthrie*

Victoria University of Wellington

April 6, 2010

Abstract

This paper modifies the standard binomial option pricing approach to real options analysis so that it can incorporate learning options. These options allow a manager to gather information about a potential investment payoff prior to investment occurring. The project's overall volatility will vary in the run-up to investment, being higher when the manager gathers more information about the eventual investment payoff. This paper shows how to construct a recombining tree for the project's anticipated value by making the time steps shorter during periods of high volatility. It describes a simple scheme for calculating the lengths of these steps and the risk-neutral probabilities that are needed to calculate arbitrage-free asset prices.

JEL Classification codes: G31, D81, C61

Keywords: real options; learning options; information gathering; binomial option pricing model

1 Introduction

Managers considering investing in a new project should compare the anticipated present value of the completed project with the opportunity cost of undertaking the investment. This opportunity cost includes the value of the funds required to complete the project, as well as the value of any real options that are destroyed by the act of investment. In particular, managers must take account of the real option to wait and possibly invest in the project at a later date. In typical real options models, the value of waiting derives from the possibility of investing in more favorable market conditions in the future (as well as not investing if conditions are unfavorable). However, an important source of value for many real-world projects, especially those involving new technologies and markets, is the ability to learn more about the eventual investment payoff by delaying investment. For example, a firm can delay launching a new product while it conducts extensive market research. If the results of this market research indicate the new product is unpopular, then the firm can avoid incurring unnecessary investment expenditure. In contrast, if the product is likely to be popular, then the firm can still invest, albeit at a slightly later date.

This paper describes a straightforward approach to incorporating learning options into real options analysis that uses a simple modification of the binomial option pricing model. While this approach will not be suitable for every situation, it has the important advantage of being

*School of Economics and Finance, PO Box 600, Victoria University of Wellington, Wellington, New Zealand. Ph: 64-4-4635763. Fax: 64-4-4635014. Email: graeme.guthrie@vuw.ac.nz

straightforward to implement.¹ It has the potential to provide an approximate indication of the importance of learning options for a wide variety of investment problems faced by corporate finance practitioners.

When pricing financial derivatives, more sophisticated techniques than the binomial option pricing model are typically used, and with good reason. These techniques are also available for analyzing real options. However, while many corporate finance practitioners are comfortable with the binomial approach to valuing financial derivatives, the number that can carry out valuation analysis with more sophisticated techniques is much smaller. Moreover, the specific models that are needed for a real-options analysis of a problem often vary from one problem to the next, so that models need to be tailor-made. For smaller projects, especially, it may be too costly to justify the investment required to build one of the more sophisticated models available. It follows that the binomial option pricing model is an important tool for real options analysis in practice. Guthrie (2009) provides a textbook treatment of this approach to real options analysis.

Most of the practitioner literature focuses on the market-climate motivation for delaying investment.² When uncertainty about the potential investment payoff is considered, it is usually only resolved at fixed points in a sequence of required investments. For example, a pharmaceutical firm undertaking research and development will typically be assumed to learn more about a new product only at each predetermined stage of the R&D process. The learning and investment activities of the firm are thus inextricably linked. In contrast, the learning options considered in this paper separate the information-gathering activities of a firm from its investment activities: firms can gather information without investing, and vice versa.

The underlying model of information gathering in this paper is relatively simple. The state variable that determines the investment payoff is assumed to be a function of two terms. The first is an observable, market-wide variable, such as an output or commodity price. The second is a project-specific constant that can be observed only once investment is complete. This constant reflects aspects of the underlying project that are initially uncertain and about which information can be gathered prior to investment. For example, the inherent demand for a new product will be uncertain, but a campaign of market research has the potential to reduce this uncertainty. Learning is modeled by allowing the decision maker to make a series of noisy observations of this constant prior to investment taking place. As more observations are accumulated, uncertainty surrounding the constant falls. Each observation induces the decision maker to update the forecast of the constant's value, injecting volatility into the anticipated investment payoff. This combines with volatility in the market-wide variable to determine the overall volatility of the investment payoff.

High current uncertainty about the project's true value and the ability to rapidly reduce this uncertainty each make the option to wait and gather more information especially valuable. In these circumstances, the minimum anticipated present value of the completed project that is required for immediate investment to be optimal has to be relatively high to compensate the firm's owner for the loss of the option. Since uncertainty falls over time (as a result of the ongoing information-gathering), the investment threshold also falls over time. If managers adopt investment policies that do not incorporate such behavior, even though they learn more about the project while they wait to invest, then they will not be maximizing their firms' market values.

Although consideration of learning options has not made much headway in the practitioner

¹A spreadsheet implementation of the technique described in this paper is available at <http://contemporaryfinance.com/forpractitioners.html>.

²See, for example, the books by Copeland and Antikarov (2003) and Shockley (2006). An exception is Guthrie (2009), who shows how learning options can be incorporated in a quite general investment problem. However, the precise implementation is specific to each particular investment problem, with the form of information gathering varying from one problem to another. Rees (2008) also discusses learning options.

literature, several authors have described such models in the academic literature. For example, Epstein et al. (1999) use the model of information gathering that is described in this paper. However, their model does not include the market-wide state variable, so that the true value of the project does not vary over time. Childs et al. (2001, 2002a,b) also omit a market-wide state variable, but this is less important in their model since they allow the project-specific variable to evolve over time. As in this paper, their decision maker can observe a noisy signal of the project's true value, but their noise terms are serially correlated. Guthrie (2007) supposes that changes in the underlying project's value are at least partially unobservable, and that the owner of the project rights must incur a cost in order to learn the exact change. All of these authors develop their models in continuous time, which, because of the complexities introduced by learning options, requires solution of partial differential equations using complicated numerical methods.

An important contribution of this paper is to show how a model of learning options can be formulated in discrete time and then implemented using a simple modification of the standard binomial option pricing model. The modification is necessary because new information has a bigger impact on predicted project value when the uncertainty surrounding this value is greater—and this occurs early in the lifetime of the investment option, before much information has been gathered. The volatility of the anticipated market value of the underlying project will therefore decline in the run up to investment. Using a standard binomial tree to model the behavior of the anticipated market value is impossible since the declining volatility means that the tree would not recombine. For example, an up move (when volatility is high) followed by a down move (when it is low) would result in a higher market value than a down move followed by an up move: the move during the high-volatility period will dominate the move during the low-volatility period. Tree-based real options analysis is impracticable when the tree for the underlying state variable does not recombine as the number of nodes in the tree grows exponentially, rather than linearly as in the case of a recombining tree.

This paper ensures that the binomial tree recombines by keeping the sizes of up and down moves constant throughout the tree, while dividing the lifetime of the option into periods of unequal length. Initially, when volatility is high, the time periods are short, so that the state variable changes frequently by the standardized amount. Subsequently, when volatility is lower, the time periods are long, so that changes in the state variable are less frequent. We derive a closed form solution for the period length and step size that are required to match the dynamics of the underlying state variable. Implementation of the method is no more difficult than the standard binomial option pricing model. Therefore, corporate finance practitioners should now find it straightforward to include a simple learning option when they carry out real options analysis of investment problems.

The next section describes the investment-timing problem that is used to illustrate the general modeling approach and recaps the standard approach to solving this problem when there is no embedded learning option.³ Section 3 introduces the model of information gathering that is the basis of the paper, shows how to build the recombining tree for the state variable, and discusses calculation of the risk-neutral probabilities that are needed to implement the option analysis. An extended numerical example appears in Section 4, which demonstrates implementation of the techniques introduced in the paper and uses them to demonstrate the importance of information gathering in real options analysis. Section 5 concludes the paper.

³This particular application is used here solely for illustrative purposes: the modeling technique described in this paper can be applied to a much wider variety of problems. For example, it can be applied to analyze sequential investment problems, or other situations involving compound timing options.

2 The baseline model

A firm has the right to invest in a particular project at any time up to and including date T . Investment is instantaneous, requires lump sum expenditure of I , and is irreversible. The project will be worth $P_t \exp(X)$ immediately after construction is completed, where t is the date on which investment occurs, P_t is the level of an exogenous state variable at this date, and X is a project-specific random variable whose value is revealed immediately after investment occurs. P_t reflects readily observable market-wide factors that will evolve independently of when the project is undertaken. In contrast, X reflects those characteristics of the project itself that do not change over time and are not directly observable until after the project has been completed. For example, for a firm that is considering when to develop a gas field, P_t might reflect the market value at date t of each unit of proven developed gas reserves, while X represents the quantity of this particular gas reserve. For a firm that is considering launching a new product in an overseas market, P_t might reflect the exchange rate at date t and X some measure of the inherent demand for the new product.

Because X reflects project-specific characteristics, the forecast error, $X - E_t[X]$, is assumed to be uncorrelated with the returns on all other assets. This means that the risk associated with the project turning out to be more or less valuable than was expected immediately before investment can be diversified away. Consequently, investors do not receive any reward for bearing this risk. The rights to the completed project are therefore worth

$$V_t \equiv E_t[P_t \exp(X)] = P_t E_t[\exp(X)]$$

immediately before investment occurs. The anticipated payoff from investing at date t is therefore $V_t - I$. The optimal investment policy is to wait as long as this investment payoff is less than the value of waiting. This “delay-option” value will depend on the current level of V_t , as well as the parameters that determine how it evolves over time. This paper shows how to modify the familiar binomial-tree approach in order to estimate the value of waiting when delay allows the manager to learn more about X .

We make the assumption, which is common in real options analysis, that the exogenous state variable, P_t , follows a geometric Brownian motion. This implies that $\log P_t$ evolves according to the random walk

$$\log P_{t+\Delta t} - \log P_t = \left(\mu - \frac{1}{2}\phi^2 \right) \Delta t + Z, \quad Z \sim N(0, \phi^2 \Delta t),$$

where Δt is the length of time over which the change in $\log P$ is measured, and the constants μ and ϕ are the drift and volatility of the geometric Brownian motion.

Before we analyze the learning model itself, we will consider what happens if the firm’s manager cannot learn any more about X than is known at date 0. In this case, the underlying project is worth $V_t = P_t E_0[\exp(X)]$ at date t , where $E_0[\exp(X)]$ is the expected value calculated using all information available to the manager at date 0. Since the manager’s forecast of $\exp X$ will not change over time in this case, the only fluctuations in V_t come from changes in the market-wide variable P_t . In fact, $\log V_t$ evolves according to the random walk

$$\log V_{t+\Delta t} - \log V_t = \left(\mu - \frac{1}{2}\phi^2 \right) \Delta t + Z, \quad Z \sim N(0, \phi^2 \Delta t). \quad (1)$$

The procedure for constructing a binomial tree that represents the behavior of V_t as described in (1) is well known.⁴ We construct a binomial tree with the desired number of time steps (N),

⁴See, for example, Chapter 12 of Guthrie (2009). Alternative specifications for the key parameters (the size and probability of up moves) have been suggested by various authors. The differences disappear in the limit as the time steps become very small, but some combinations of parameters will lead to faster convergence than others.

each of a constant length $\Delta t = T/N$. The binomial tree starts with the variable taking the value $V_0 = P_0 E_0[\exp(X)]$. Each period, it changes by a factor of either $U = \exp(\phi\sqrt{\Delta t})$ or $D = 1/U$. Since $U > 1$ and $D < 1$, we call these “up” and “down” moves respectively. An up move occurs in any given period with probability

$$\frac{1}{2} + \frac{(\mu - \frac{1}{2}\phi^2)\sqrt{\Delta t}}{2\phi}.$$

This ensures that the change in $\log V$ during each period has an expected value that equals $(\mu - \frac{1}{2}\phi^2)\Delta t$ and a standard deviation that is approximately equal to $\phi\sqrt{\Delta t}$, consistent with the random walk in (1). Once this tree has been built, we can use it to estimate the value of the option to delay investment. This is demonstrated in Section 4.

We now turn to the main focus of this paper, the case in which the firm’s manager can learn more about X by delaying investment.

3 Modeling the learning process

We model the manager’s uncertainty about the investment payoff by assuming that, based on all the information that is available to him at date t , X is normally distributed with mean $X_t \equiv E_t[X]$ and variance $a_t^2 \equiv \text{Var}_t[X]$. The completed project is therefore worth

$$V_t = P_t E_t[\exp(X)] = P_t \exp\left(E_t[X] + \frac{1}{2}\text{Var}_t[X]\right) = P_t \exp\left(X_t + \frac{1}{2}a_t^2\right) \quad (2)$$

immediately before investment occurs. In order to derive an optimal investment policy for the firm, we must be able to identify the process that describes how V_t evolves over time. Taking logarithms of equation (2) shows that the log project value at date t is

$$\log V_t = \log P_t + X_t + \frac{1}{2}a_t^2. \quad (3)$$

The discussion in Section 2 shows how $\log P_t$ behaves over time. Therefore, once we know the behavior of X_t and a_t^2 we will be able to calculate the process that determines how the log project value evolves over time. Section 3.1 derives this behavior, starting from a particularly simple model of information gathering.

However, before we proceed it is worthwhile thinking about what the assumption above means for the distribution of possible project values. Let $V_t' = P_t \exp(X)$ denote the market value of the completed project immediately after investment occurs, assuming it occurs at date t . Then the assumption above implies that (when viewed immediately prior to investment) $\log V_t'$ is normally distributed with mean

$$\log P_t + X_t = \log V_t - \frac{1}{2}a_t^2.$$

and variance a_t^2 . Equivalently,

$$\frac{\log V_t' - \log V_t + \frac{1}{2}a_t^2}{a_t}$$

is drawn from the standard normal distribution. Consequently, if $x\%$ of the standard normal distribution lies between $-q$ and q , then the $x\%$ confidence interval for the post-investment market value is described by

$$-q \leq \frac{\log V_t' - \log V_t + \frac{1}{2}a_t^2}{a_t} \leq q,$$

which can be rewritten as

$$V_t \exp\left(-qa_t - \frac{1}{2}a_t^2\right) \leq V_t' \leq V_t \exp\left(qa_t - \frac{1}{2}a_t^2\right). \quad (4)$$

Knowledge of this result can be useful in interpreting our starting value for a_t . It can even provide a useful means of calibrating the model: it may be difficult, if not impossible, to estimate a_0 accurately, but if managers have an intuitive estimate of the confidence interval for the project under consideration, we can infer an appropriate value of a_0 using (4).

3.1 A simple model of information gathering

The information-gathering process is modeled by assuming that the firm's manager can obtain noisy observations of X at a sequence of dates t_0, t_1, \dots, t_{N-1} , where $t_0 = 0$ and $t_{N-1} < T$. This occurs until the firm invests in the project or the investment option expires, whichever occurs sooner. For example, the manager of a firm considering developing a gas field can drill test wells to learn more about the quantity of gas contained in the reserve; the manager of a firm considering expanding into a new overseas market can carry out small test marketing campaigns to assess demand. The noise in the observations means that the manager does not learn the exact level of X . However, if the manager observes a series of high values, this gives a strong indication that the project-specific characteristic is also high, which might be sufficient to make investment optimal.

Each new piece of information gathered potentially changes the forecast of X and reduces the standard deviation of the prediction error. Because the precise content of each piece of information is unpredictable, information-gathering causes the forecast X_t to fluctuate randomly over time. However, as more information is gathered, the firm's manager will obtain an increasingly accurate estimate of the value of X . Eventually, there is little information left to gather, and the volatility of X_t will die out.

We assume that if the manager of the firm gathers information at date t_n , then he immediately observes $Y_n = X + \varepsilon_n$, where $\varepsilon_n \sim N(0, \theta^2/(t_{n+1} - t_n))$ and θ is a constant. All noise terms are independently distributed and are uncorrelated with the returns on all assets in the economy.⁵ Notice that if t_{n+1} is a relatively long time after t_n then the noise component of Y_n has a relatively small variance. That is, when the manager spends a considerable amount of time gathering information (equal to the time elapsed between dates t_n and t_{n+1}), he gets a relatively precise observation of X . The particular functional form of the variance is chosen because it ensures that the amount of information gathered over any interval of time depends on the length of the interval and not how it is split up into subintervals. For example, we will see that gathering information between dates t_0 and t_2 has the same effect as gathering information between dates t_0 and t_1 and then between dates t_1 and t_2 . This result ensures that the information-gathering process is comparable regardless of the number of periods into which the investment option's lifetime is divided. This is important because when we implement real options analysis in practice we use a large number of small time steps—and we want our analysis to reflect the same underlying reality regardless of the number of time steps we use.

It can be shown that by the time n observations have been made, the manager's information about X has evolved to the point that he believes X is normally distributed with variance

$$a_{t_n}^2 = \frac{a_0^2}{1 + \left(\frac{a_0}{\theta}\right)^2 t_n} \quad (5)$$

⁵We could allow the noise terms to be serially correlated, as in Childs et al. (2001, 2002a,b). However, this significantly complicates the behavior of V_t . The procedure for constructing a recombining tree that is developed in this paper requires the assumption that the noise terms are uncorrelated.

and mean⁶

$$X_{t_{n+1}} = \left(\frac{\frac{\theta^2}{t_{n+1}-t_n}}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}} \right) X_{t_n} + \left(\frac{a_{t_n}^2}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}} \right) Y_n. \quad (6)$$

Equation (5) shows that the uncertainty concerning the actual value of X falls as more information is gathered (that is, as t_n increases). It starts, at date $t_0 = 0$, at a_0^2 , and falls to $a_0^2/(1 + (a_0/\theta)^2 T)$ by the time the firm's investment option expires; if information could be gathered indefinitely, then the uncertainty would eventually disappear altogether.⁷ The constant θ determines how rapidly the variance falls. For instance, information needs to be gathered for

$$h = 3 \left(\frac{\theta}{a_0} \right)^2 \quad (7)$$

years in order to halve the standard deviation.⁸ A large value of θ (which means that the manager's observations of X are very noisy) makes h large, indicating that learning about the project's underlying characteristics is a slow process. The parameter h defined in equation (7) is a natural way to measure the rate of learning, and will provide a convenient means of interpreting the results in Section 4.⁹

Equation (6) shows that the new estimate of X is a weighted average of the current estimate (X_{t_n}) and the latest observation (Y_n). At first, the uncertainty concerning the actual value of X (which is measured by $a_{t_n}^2$) will be relatively large, so that the manager will place a relatively high weight on the new information Y_n . However, as more information is gathered and uncertainty falls, $a_{t_n}^2$ will become relatively small, so that a relatively low weight will be attached to the new information. As a result, the manager's point estimate of the project-specific characteristic (X_{t_n}) will initially be quite volatile as the new information is processed. Eventually, however, this volatility will fall as the new information begins to have only a small effect on the point estimate. For example, the first few test wells drilled in an undeveloped gas field will have a much greater effect on the forecast of the gas reserve's size than later tests (by which time much of the uncertainty will already have been resolved).

It can be shown that when viewed from date t_n , the manager believes that the revised forecast of X (that is, $X_{t_{n+1}}$) is normally distributed with mean X_{t_n} and variance

$$\text{Var}_{t_n}[X_{t_{n+1}}] = \frac{\theta^2(t_{n+1} - t_n)}{\left(t_n + \left(\frac{\theta}{a_0} \right)^2 \right) \left(t_{n+1} + \left(\frac{\theta}{a_0} \right)^2 \right)}. \quad (8)$$

Therefore the forecast of X evolves according to

$$X_{t_{n+1}} - X_{t_n} = W_{n+1}, \quad W_{n+1} \sim N \left(0, \frac{\theta^2(t_{n+1} - t_n)}{\left(t_n + \left(\frac{\theta}{a_0} \right)^2 \right) \left(t_{n+1} + \left(\frac{\theta}{a_0} \right)^2 \right)} \right).$$

There are three fundamentally different standard deviations involved in this model. First, at any given date t_n , there is the standard deviation of the forecast error, which is denoted

⁶A heuristic derivation, along the lines of the one in Epstein et al. (1999), is included in Appendix A.1. A formal derivation of the corresponding result in continuous time, known as the Kalman-Bucy filter, can be found in Øksendal (1998).

⁷Note that the uncertainty remaining at date t_n , as measured by the variance in equation (5), depends on the total time elapsed (t_n) and not on the number or timing of the individual noisy observations of X .

⁸To see why, substitute $t_n = h$ into (5). After some manipulation, the variance simplifies to $a_{t_n}^2 = a_0^2/4$.

⁹In the numerical example in Section 4, the model is actually specified in terms of h and the implied value of θ is calculated using $\theta = a_0\sqrt{h/3}$.

Table 1: Behavior of the three standard deviations

n	$\theta = 0.25$				$\theta = 0.50$			
	Date	Uncertainty	Noise	Volatility	Date	Uncertainty	Noise	Volatility
0	0.000	1.000	0.559	0.873	0.000	1.000	1.118	0.667
1	0.200	0.488	0.559	0.321	0.200	0.745	1.118	0.413
2	0.400	0.368	0.559	0.202	0.400	0.620	1.118	0.301
3	0.600	0.307	0.559	0.148	0.600	0.542	1.118	0.237
4	0.800	0.269	0.559	0.117	0.800	0.488	1.118	0.195
5	1.000	0.243	n/a	n/a	1.000	0.447	n/a	n/a

n	$\theta = 0.25$				$\theta = 0.50$			
	Date	Uncertainty	Noise	Volatility	Date	Uncertainty	Noise	Volatility
0	0.000	1.000	0.559	0.434	0.000	1.000	2.291	0.400
1	0.014	0.901	0.403	0.434	0.048	0.917	1.889	0.400
2	0.038	0.790	0.333	0.434	0.118	0.825	1.487	0.400
3	0.081	0.660	0.295	0.434	0.231	0.721	1.082	0.400
4	0.190	0.497	0.278	0.434	0.444	0.600	0.671	0.400
5	1.000	0.243	n/a	n/a	1.000	0.447	n/a	n/a

a_{t_n} . It measures the uncertainty currently surrounding the manager’s forecast of the project-specific characteristic X . Second, there is the standard deviation of the noise term in the noisy observation of the true value of X , which equals $\theta/\sqrt{t_{n+1} - t_n}$. This measures the amount of noise in the new information gathered by the manager. Third, there is the standard deviation of the change in the estimated value of X , which equals the square root of the variance in equation (8). This measures the volatility of the forecast of the project-specific characteristic, and directly determines the value of the option to delay investment in order to gather more information.

Table 1 uses a simple numerical example to demonstrate how these three standard deviations evolve over time. In the top panel, an interval spanning one year has been divided into five periods of equal length. This panel is split into two parts, with the left-hand side showing how the three standard deviations evolve when there is relatively little noise contained in the manager’s new information; there is much more noise in the case shown in the right-hand side of the panel. The columns labeled “Uncertainty” show that the forecast error’s standard deviation gradually falls from its initial value of $a_0 = 1$ to the much lower values of $a_1 = 0.243$ and $a_1 = 0.447$, depending on whether the manager can gather low-noise or high-noise information, respectively. Notice that uncertainty falls rapidly at first, and more slowly as time progresses. Because the time periods are all of equal length, the standard deviation of the noise term (shown in the columns labeled “Noise”) is constant over time. However, the standard deviation of the change in the forecast (shown in the columns labeled “Volatility”) is falling over time, reflecting the fact that much more information is acquired in the first period than in later ones of the same length.

The bottom panel of Table 1 has the same format as the top panel, but now the year is broken into five periods of unequal length. (The method for calculating the length of each period will be described shortly.) As in the top panel, the uncertainty in X starts at $a_0 = 1$ and falls to either $a_1 = 0.243$ and $a_1 = 0.447$ by the end of the year. However, for the dates used in the table, the standard deviation of the change in the forecast of X is the same for each period, which we will see allows us to model the evolution of the project value using a recombining tree. This is achieved by making the periods relatively short at first (when volatility is high) and then increasing over time as the amount of new information gathered falls.

3.2 Building a tree for the predicted market value of the completed project

It is straightforward to combine the process for the predicted value of the project-characteristic X with the one for the observable state variable P_t to obtain a simple process that describes how the pre-investment market value of the completed project evolves over time. The information-gathering process is independent of the economy as a whole. Therefore, the new information revealed about the project-specific characteristic is uncorrelated with shocks to the market-wide variable P_t . As shown in Appendix A.2, the logarithm of the pre-investment market value of the completed project evolves according to

$$\log V_{t_{n+1}} - \log V_{t_n} = \left(\mu - \frac{1}{2} \sigma_n^2 \right) (t_{n+1} - t_n) + Z'_{n+1}, \quad Z'_{n+1} \sim N(0, \sigma_n^2 (t_{n+1} - t_n)), \quad (9)$$

where

$$\sigma_n^2 = \phi^2 + \frac{\theta^2}{\left(t_{n+1} + \left(\frac{\theta}{a_0} \right)^2 \right) \left(t_n + \left(\frac{\theta}{a_0} \right)^2 \right)}. \quad (10)$$

Equations (9) and (10) are key, since they completely determine the behavior of the firm's anticipated investment payoff. The first equation implies that the pre-investment market value of the completed project is expected to grow at rate μ . That is, the market value grows at the same rate as the market-wide state variable on average. Its volatility, σ_n , is a function of three parameters—the volatility of the market-wide variable (ϕ), the initial uncertainty surrounding the unobservable project-specific characteristic (a_0), and the amount of noise in the information gathered by the firm's manager (θ)—and time. The volatility in equation (10) has two distinct components. The first one results from ongoing fluctuations in the market-wide variable. The assumption that the state variable follows geometric Brownian motion means that this volatility is constant over time. The second component results from the information-gathering process. It is large when the initial uncertainty around the project-specific characteristic is large, and diminishes over time as the incremental information flow falls. As t grows very large (so that information gathering has ceased to inject any significant volatility into the process for V_t), the volatility converges to ϕ , which reflects solely fluctuations in the market-wide variable.

The time-varying nature of volatility poses problems when we come to build a binomial tree for V_t . In the standard procedure described in Section 2, $\log U$ equals the standard deviation of the change in $\log V_t$ over each time step. One way to capture time-varying volatility might therefore be to allow the size of each up move to vary over time. For example, if volatility is falling over time we might use a relatively large value, U_0 , when calculating the size of the change in V_t over the first period and a smaller value, U_1 , over the second period. Unfortunately, such a choice means that the tree does not recombine: that is, an up move followed by a down move results in a different value for V from a down move followed by an up move. Suppose, for example, that $V_0 = 100$, $U_0 = 1.25$, and $U_1 = 1.10$. Since the sizes of the corresponding down moves are $D_0 = 1/U_0 = 0.80$ and $D_1 = 1/U_1 = 0.91$, V_2 equals

$$V_0 U_0 D_1 = 100 \times 1.25 \times 0.91 = 113.64$$

in the first case and

$$V_0 D_0 U_1 = 100 \times 0.80 \times 1.10 = 88.00$$

in the second case. If the binomial tree is to recombine then the sizes of the up moves, U_0 and U_1 , must satisfy $U_0 U_1^{-1} = U_0^{-1} U_1$, which requires that $U_0 = U_1$. That is, each up move must have the same size.

At first glance, this is difficult to reconcile with time-varying volatility: if volatility changes over time then how can the up moves possibly have the same size over time? However, this

problem can be overcome by allowing different time steps to represent different lengths of time. During periods when volatility is relatively high it is appropriate to have a small time step, so that V changes by the standardized amount in a short space of time. In contrast, when volatility is relatively low the time steps should be quite long, so that V takes longer to change by the standardized amount.

The key to implementing this approach is the calculation of the length of the various time steps. Suppose we have divided time up into a sequence of periods, defined by dates $t_0 < t_1 < \dots < t_{N-1} < t_N$, where $t_0 = 0$ and $t_N = T$. Extending the standard approach described in Section 2 to this situation suggests that an up move should occur between dates t_n and t_{n+1} with probability

$$\frac{1}{2} + \frac{(\mu - \frac{1}{2}\sigma_n^2)\sqrt{t_{n+1} - t_n}}{2\sigma_n}$$

and that the size of this move should be

$$\exp\left(\sigma_n\sqrt{t_{n+1} - t_n}\right).$$

This follows from using the appropriate values of the drift and volatility of V_t , and replacing the constant length of each period, Δt , with the particular length, $t_{n+1} - t_n$. The size of up moves will be constant throughout the tree as long as

$$\exp\left(\sigma_n\sqrt{t_{n+1} - t_n}\right) = U \quad (11)$$

for some *constant* U , for all values of n . During periods when V is especially volatile, σ_n will be large, so that the distance between t_n and t_{n+1} will be small. That is, high volatility corresponds to short time steps, consistent with the discussion above.

Appendix A.3 shows that setting the (log) size of an up move equal to

$$\log U = \left(\frac{T}{N} \left(\phi^2 + \frac{a_0^2}{T + \left(\frac{\theta}{a_0}\right)^2} \right) \right)^{1/2} \quad (12)$$

allows us to divide an interval of T years into N periods in such a way that the tree recombines. Each date t_n can be calculated using

$$t_n = - \left(\frac{a_0^2 + \left(\frac{\phi\theta}{a_0}\right)^2 - n(\log U)^2}{2\phi^2} \right) + \sqrt{ n \left(\frac{\theta \log U}{a_0\phi} \right)^2 + \left(\frac{a_0^2 + \left(\frac{\phi\theta}{a_0}\right)^2 - n(\log U)^2}{2\phi^2} \right)^2 }. \quad (13)$$

We are now able to build a recombining tree for the anticipated market value of the completed project. However, before we can use this tree to value the right to undertake the project, we must consider how we will calculate the risk-neutral probabilities.

3.3 Risk-neutral probabilities

The no-arbitrage approach to valuation compensates for risk by adjusting the probabilities that are used to calculate the expected values of future cash flows. Consider a cash flow that will be received one period from now and which equals Y_u if the anticipated value of the project considered in this paper experiences an up move, and Y_d if a down move occurs instead. The current market value of this cash flow equals

$$\frac{\pi_u Y_u + (1 - \pi_u) Y_d}{R},$$

where R is the total return on a one-period risk-free bond and

$$\pi_u = \frac{\frac{ZR}{V} - D}{U - D}. \quad (14)$$

In the expression for the risk-neutral probability, π_u , V is the current anticipated value of the completed project, U is the size of an up move, and $D = 1/U$ is the size of a down move. Z is the current price of the so-called “spanning asset”, which pays UV after one period if an up move occurs and DV if a down move occurs instead.¹⁰ The key to valuation is calculating Z .

Holders of the spanning asset are exposed to three different types of risk: systematic shocks to the market-wide state variable, unsystematic shocks to this variable, and new information gathered about the project-specific characteristic. Of these, all but the first one can be diversified away. Therefore, holders are rewarded only for bearing the systematic component of P -risk. We will suppose that the continuously-compounded risk-adjusted discount rate for P -risk equals $r + \lambda$, where the constants r and λ are the risk-free interest rate and risk-premium for bearing P -risk, respectively.

Suppose an investor buys one unit of the spanning asset at date t_n . This asset generates a single payoff, equal to $V_{t_{n+1}}$, that is paid out at date t_{n+1} . Since V is expected to grow at rate μ , the expected payoff of this asset, when viewed from the perspective of date t_n , is

$$E_{t_n}[V_{t_{n+1}}] = e^{\mu(t_{n+1}-t_n)}V_{t_n}.$$

That is, we take the level of V at date t_n and compound it forward at rate μ until date t_{n+1} . To calculate the market value of this asset at date t_n , we take this expected payoff and discount it back to date t_n using the risk-adjusted discount rate $r + \lambda$. The market value of the spanning asset at date t_n is therefore

$$Z_{t_n} = \frac{e^{\mu(t_{n+1}-t_n)}V_{t_n}}{e^{(r+\lambda)(t_{n+1}-t_n)}} = e^{(\mu-r-\lambda)(t_{n+1}-t_n)}V_{t_n}.$$

The price at date t_n of a risk-free discount bond that pays one dollar at date t_{n+1} is

$$R_{t_n} = e^{-r(t_{n+1}-t_n)}.$$

Equation (14) therefore implies that the risk-neutral probability of an up move occurring at date t_n equals

$$\pi_{u,t_n} = \frac{e^{(\mu-\lambda)(t_{n+1}-t_n)} - U^{-1}}{U - U^{-1}}. \quad (15)$$

In the special case that the firm’s manager cannot learn anything more about the project-specific characteristic by delaying investment, we saw in Section 2 that the binomial tree is made up of periods each of constant length Δt . In this case the risk-neutral probability of an up move,

$$\pi_u = \frac{e^{(\mu-\lambda)\Delta t} - U^{-1}}{U - U^{-1}},$$

takes the same value throughout the tree. However, the periods are of unequal length in the general case, with the dates being close together early in the life of the firm’s investment option and further apart later on. In this case, the risk-neutral probability of an up move is no longer constant through the tree. However, it can easily be calculated using equation (15).

¹⁰See, for example, Chapter 3 of Guthrie (2009).

3.4 Summary

We conclude this section by briefly summarizing the steps that need to be undertaken in order to build the binomial tree for the pre-investment market value of the completed project.

1. Choose the number of steps (N) that the lifetime of the option (T) will be divided into.
2. Calculate the (log) size of each up move using equation (12).
3. Calculate the dates using equation (13).
4. Fill in the entries in the binomial tree for V by setting

$$V(i, n) = V_0 D^i U^{n-i} = V_0 U^{n-2i},$$

for all values of $n = 0, 1, \dots, N$ and $i = 0, 1, \dots, n$, where i is the number of down moves and n is the number of periods.

5. Calculate the risk-neutral probability of an up move at each node using equation (15).

This procedure will be demonstrated in Section 4.

4 Numerical example

This section features a simple numerical example using the techniques developed in Section 3. This example demonstrates that the calculations needed to incorporate learning are no more difficult than the standard binomial option pricing techniques familiar to corporate finance practitioners. In addition, this section illustrates conceptual material such as the role that information gathering plays in determining optimal investment policies.

The firm in this example has the right to invest in a project at any time within the next $T = 5$ years. Investment requires lump sum expenditure of $I = 9$. Based on all information available at date 0, the market value of the project would be $V_0 = 10$ if investment occurred immediately. However, this estimate is imprecise, and the manager believes that the standard deviation of the prediction error is $a_0 = 0.5$. For example, if investment occurred immediately, the actual value of the completed project will be found to be less than 6.30 with probability 0.25 and greater than 12.36 with probability 0.25.¹¹ In the absence of any learning by the firm's manager, the anticipated market value of the completed project evolves according to a geometric Brownian motion with drift $\mu = 0.02$ and volatility $\phi = 0.1$. The risk-free interest rate is $r = 0.05$ and the risk-premium that is applied to the market value of the completed project is $\lambda = 0.04$.

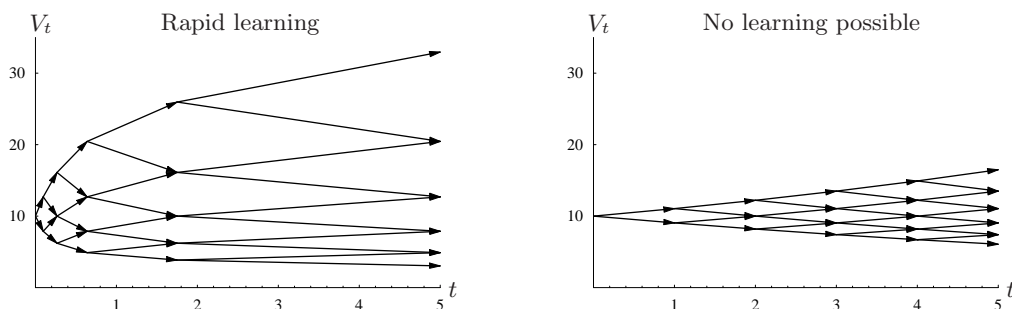
We consider two different information-gathering scenarios. In the first, the manager can halve the prediction error with just one year of information gathering (that is, $h = 1$). Given the choice of h , the parameter θ is determined by equation (7).¹² In the second scenario, which corresponds to standard investment timing models, delaying investment does not provide the manager with any ability to acquire additional information regarding the project-specific characteristic. This is approximated by setting $h = 10^6$, so that it would take one million years of investigation to halve the standard deviation of the project's prediction error.

Once the desired number of time steps (N) has been chosen, the size of each up move is calculated using equation (12). The precise value of each date in the tree is calculated using equation (13). For example, if $h = 1$ and the lifetime of the investment option is to be split

¹¹This follows from noting that the upper quartile of the standard normal distribution is $q = 0.6745$ and evaluating the confidence interval in (4).

¹²The solution for θ is given in footnote 9.

Figure 1: Binomial trees for two different learning scenarios



into $N = 5$ periods, then equation (12) implies that each up move has size $U = 1.2693$, while equation (13) shows that the binomial tree uses dates

$$t_0 = 0.0000, \quad t_1 = 0.0960, \quad t_2 = 0.2666, \quad t_3 = 0.6384, \quad t_4 = 1.7500, \quad t_5 = 5.0000.$$

If it is to be split into $N = 10$ periods instead, each up move has size $U = 1.1370$ and the dates are

$$t_0 = 0.0000, \quad t_1 = 0.0421, \quad t_2 = 0.0960, \quad t_3 = 0.1676, \quad t_4 = 0.2666, \quad t_5 = 0.4112, \\ t_6 = 0.6384, \quad t_7 = 1.0279, \quad t_8 = 1.7500, \quad t_9 = 3.0539, \quad t_{10} = 5.0000.$$

Notice that in both cases $t_N = T$, so that the final date in the tree coincides exactly with the expiry date of the option.¹³ The case with $h = 1$ and $N = 5$ results in the binomial tree shown in the left-hand graph of Figure 1. Using $h = 10^6$ and $N = 5$ leads to the right-hand graph. The graphs show that, early in the life of the option, volatility is much greater when the manager is able to gather information about the project than when learning is impossible.

The top panel of Table 2 gives the details of the binomial trees in Figure 1. In each table, the columns correspond to the period n , while the rows correspond to the number of down moves i . The left-hand table shows the case when $h = 1$ (that is when one year of information gathering will halve the standard deviation of the prediction error in the pre-investment market value of the completed project), while the right-hand table shows the case when h is so large that the manager learns nothing about the project while he delays investment. Consistent with the graphs in Figure 1, the potential for information gathering leads to a much wider range of possible values for the future anticipated market value of a completed project.

The risk-neutral probabilities of up moves occurring are shown in the middle panel of Table 2. Consistent with the formula in equation (15), the risk-neutral probability depends on the date n but not on the number of down moves i . The right-hand table shows that in the special case where no learning is possible the risk-neutral probability takes the same value at every node.

Once the trees for the state variable and risk-neutral probability have been built, construction of the tree for the market value of the project rights is straightforward. At the option's expiry date, the project rights are worth the greater of the anticipated investment payoff and zero:

$$F(i, N) = \max \{V(i, N) - I, 0\}.$$

¹³When $h = 10^6$, so that there is almost no ability to learn about the project, the time steps are equally spaced. For example, when $N = 5$ the binomial tree has $t_0 = 0$, $t_1 = 1$, $t_2 = 2$, $t_3 = 3$, $t_4 = 4$, and $t_5 = 5$.

Table 2: Binomial trees for two different learning scenarios

Rapid learning							No learning possible						
$V(i, n)$	0	1	2	3	4	5	$V(i, n)$	0	1	2	3	4	5
0	10.00	12.69	16.11	20.45	25.96	32.95	0	10.00	11.05	12.21	13.50	14.92	16.49
1		7.88	10.00	12.69	16.11	20.45	1		9.05	10.00	11.05	12.21	13.50
2			6.21	7.88	10.00	12.69	2			8.19	9.05	10.00	11.05
3				4.89	6.21	7.88	3				7.41	8.19	9.05
4					3.85	4.89	4					6.70	7.41
5						3.03	5						6.07

$\pi_u(i, n)$	0	1	2	3	4	5	$\pi_u(i, n)$	0	1	2	3	4	5
0	0.437	0.434	0.425	0.395	0.310	n/a	0	0.376	0.376	0.376	0.376	0.376	n/a
1		0.434	0.425	0.395	0.310	n/a	1		0.376	0.376	0.376	0.376	n/a
2			0.425	0.395	0.310	n/a	2			0.376	0.376	0.376	n/a
3				0.395	0.310	n/a	3				0.376	0.376	n/a
4					0.310	n/a	4					0.376	n/a
5						n/a	5						n/a

$F(i, n)$	0	1	2	3	4	5	$F(i, n)$	0	1	2	3	4	5
0	2.23	4.04	7.11	11.45	16.96	23.95	0	1.00	2.05	3.21	4.50	5.92	7.49
1		0.84	1.75	3.69	7.11	11.45	1		0.44	1.00	2.05	3.21	4.50
2			0.16	0.37	1.00	3.69	2			0.14	0.37	1.00	2.05
3				0.00	0.00	0.00	3				0.01	0.02	0.05
4					0.00	0.00	4					0.00	0.00
5						0.00	5						0.00

The only change from the standard approach arises because the steps in the tree represent differing lengths of time. As a result, the market value of the project rights at node (i, n) equals

$$F(i, n) = \max \left\{ V(i, n) - I, \frac{\pi_u(i, n)F(i, n+1) + (1 - \pi_u(i, n))F(i+1, n+1)}{e^{r(t_{n+1}-t_n)}} \right\}.$$

The first expression inside the brackets is the payoff from investing immediately, while the second one is the payoff from waiting until date t_{n+1} and then reevaluating the investment decision. This payoff is the expected value of the project rights at date t_{n+1} (calculated using the appropriate risk-neutral probabilities of up and down moves) discounted back to date t_n using the risk-free interest rate.

The bottom panel of Table 2 shows the market value of the project rights throughout the lifetime of the investment option for the two situations under consideration. The shaded cells indicate nodes of the respective binomial trees where it is optimal to invest in the project. For the parameters chosen here, when delay allows the manager to gather information about the project it is optimal to initially delay investment in the project. However, the right-hand table shows that it is optimal to invest immediately if the manager is unable to learn about the project's underlying characteristics by delaying. Finally, comparison of the two tables in the bottom panel shows that the possibility of learning more about the project makes the project rights substantially more valuable: the initial value of 1.00 without information-gathering potential rises to 2.23 when information gathering is possible.

The trees used so far have been kept small for illustrative purposes. However, a larger number of time steps is required if the results are to be reliable. Therefore, for the remainder of this section the five-year lifetime of the investment option is divided into 1000 steps.

Figure 2 illustrates the effect of learning options on the optimal investment threshold. Each curve plots the threshold V_t^* above which investment is optimal at date t for three different

Figure 2: Optimal investment policies for three different learning scenarios




Table 3: Effect of information gathering potential on the market value of the project rights

a_0	h													
	1	2	3	4	5	6	7	8	9	10	100	1000	∞	
0.0	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05
0.1	1.15	1.13	1.12	1.11	1.10	1.10	1.09	1.09	1.09	1.09	1.06	1.05	1.05	1.05
0.2	1.36	1.32	1.29	1.27	1.25	1.23	1.22	1.21	1.20	1.19	1.07	1.05	1.05	1.05
0.3	1.63	1.56	1.52	1.48	1.45	1.42	1.40	1.38	1.36	1.34	1.10	1.06	1.05	1.05
0.4	1.93	1.84	1.77	1.72	1.68	1.64	1.61	1.58	1.55	1.53	1.14	1.06	1.05	1.05
0.5	2.23	2.12	2.04	1.98	1.92	1.88	1.84	1.80	1.77	1.74	1.20	1.06	1.05	1.05
0.6	2.54	2.42	2.32	2.24	2.18	2.12	2.07	2.03	1.99	1.95	1.26	1.07	1.05	1.05
0.7	2.86	2.71	2.60	2.51	2.44	2.37	2.32	2.26	2.22	2.17	1.33	1.08	1.05	1.05
0.8	3.17	3.01	2.88	2.78	2.70	2.63	2.56	2.50	2.45	2.40	1.41	1.09	1.05	1.05
0.9	3.48	3.30	3.16	3.05	2.96	2.88	2.81	2.74	2.68	2.63	1.49	1.10	1.05	1.05
1.0	3.78	3.59	3.44	3.32	3.22	3.13	3.05	2.98	2.91	2.85	1.58	1.11	1.05	1.05

scenarios. The dotted curve corresponds to the case where no learning is possible.¹⁴ When h (and hence the standard deviation θ of the noise term associated with each episode of information gathering) is very large, the initial level of uncertainty (as measured by a_0) has no effect on the optimal investment threshold, so that this curve applies to all values of a_0 . The two solid curves correspond to scenarios where $h = 1$. The middle curve has $a_0 = 0.5$, while the top one has $a_0 = 1$. Comparison of these curves with the dashed one shows that the threshold is higher when learning is possible. The optimal threshold falls rapidly towards the no-learning case as the extent of the remaining uncertainty falls. Finally, for a given rate of learning, greater initial uncertainty leads to a higher development threshold.

Table 3 shows the effect of information-gathering potential on the market value of the project rights at date 0. The parameters that reflect the potential to gather information about the project vary as indicated in the table, with the rows corresponding to different values of a_0 and the columns to different values of h .¹⁵ All other parameters take the same values as in the rest of this section. For each combination of a_0 and h , the construction follows the same procedure as in the examples at the start of this section—calculating U , then the 999 separate dates t_1, t_2, \dots, t_{999} , followed by trees for the anticipated market value of the completed project, the risk-neutral probability of an up move, and the market value of the project rights. The entries in the table give the market value of the project rights at node $(0, 0)$.

If there is no uncertainty about the current value of the project (the top row, where $a_0 = 0$),

¹⁴It is actually calculated using $h = 10^6$.

¹⁵The row labeled $a_0 = 0$ actually uses $a_0 = 10^{-6}$, while the column labeled $h = \infty$ actually uses $h = 10^6$. These approximate values are used to avoid singularities in the calculations.

the project rights are initially worth 1.05. There can be substantial uncertainty, but if there is no capacity to learn more about the project’s true value over time (the right-hand column, where $h \rightarrow \infty$) then the project rights will also be worth 1.05. However, whenever there is some uncertainty about the project’s actual value and the manager is able to reduce this uncertainty while waiting to invest, the table shows that the project rights are worth more than 1.05. In some cases, the increase in value is substantial. The market value of the project rights increases when there is more uncertainty about the project’s actual value (that is, as we move down the columns in Table 3) or that uncertainty can be reduced more rapidly (that is, as we move left along the rows). In both cases, the short-term volatility of the investment payoff is relatively high. In the first case, a single piece of new information has a greater effect on the manager’s prediction of the project’s value when initial uncertainty is high because, as equation (6) shows, the weight attached to the new observation is relatively high. In the second case, a single piece of new information has a greater effect on the manager’s predicted project value when there is relatively little noise associated with that observation: equation (6) shows that the weight attached to the new observation is relatively high when θ (and hence h) is small. The increased volatility raises the value of the option to wait and see what happens to the anticipated investment payoff, leading to a more stringent investment hurdle, and a higher investment threshold.

The discussion in this section shows that the ability to learn more about a project while investment is delayed can have a substantial impact on the market value of the project rights, as well as on the optimal policy for exercising those rights. All that is required to analyze the impact of this learning option is a simple modification of the standard binomial option pricing model.

5 Conclusion

This paper modifies the standard binomial option pricing approach to real options analysis so that it can incorporate learning options. In the modified approach, volatility in the anticipated market value of the underlying project has two sources. As in standard real options models, fluctuations in market-wide factors such as input and output prices affect the project’s market value. However, the decision-maker’s ability to learn more about the inherent profitability of the project prior to investment introduces additional volatility. This source of volatility is greatest when much about the project is still unknown, so that the project’s overall volatility falls over time even though the rate of information gathering is constant in the model. This poses problems for the standard binomial option pricing model, as it performs best—specifically, the tree for the state variable recombines—when volatility is constant over time. The main contribution of this paper is to show how a recombining tree can be constructed by making the time steps relatively short early in the life of the investment option and longer later. The paper describes a simple scheme for calculating the lengths of these steps and the risk-neutral probabilities that are needed to calculate arbitrage-free asset prices. This allows us to easily build trees with an arbitrarily large number of periods, which is important in practical applications where very large trees are used to ensure that the results converge to some continuous-time limit.

This paper makes it possible to implement a straightforward real options analysis of situations in which firms may benefit from waiting in order to learn more about a project’s inherent profitability. The model of information-gathering is highly stylized, and so will not be appropriate in every situation. For example, information gathering may not occur at a constant rate: there will be some stages of development where much information is revealed in a short period of time and other stages where very little is learnt. When analyzing such situations our first choice would be to build a model that is closely tailored to the problem at hand. Techniques are available for this purpose, although the data requirements can be substantial (Guthrie, 2009,

Chapter 11). When, as will often be the case, such situation-specific analysis is impractical, the generic approach described in this paper may provide an effective method for obtaining an approximate indication of the importance of learning options.

References

- Childs PD, Ott SH, and Riddiough TJ. 2001. Valuation and information acquisition policy for claims written on noisy real assets. *Financial Management* **30(2)**:45–75.
- Childs PD, Ott SH, and Riddiough TJ. 2002a. Optimal valuation of claims on noisy real assets: Theory and an application. *Real Estate Economics* **30(3)**:415–443. doi: 10.1111/1540-6229.00045
- Childs PD, Ott SH, and Riddiough TJ. 2002b. Optimal valuation of noisy real assets. *Real Estate Economics* **30(3)**:385–414. doi: 10.1111/1080-8620.t01-1-00010
- Copeland T and Antikarov V. 2003. *Real Options: A Practitioner's Guide*. Texere: New York, NY.
- Epstein D, Mayor N, Schönbucher P, Whalley E, and Wilmott P. 1999. The value of market research when a firm is learning: Real options pricing and optimal filtering. In Trigeorgis L (ed). *Real Options and Business Strategy: Applications to Decision Making*. Risk Books: London.
- Guthrie G. 2007. Missed opportunities: Optimal investment timing when information is costly. *Journal of Financial and Quantitative Analysis* **42(2)**:467–488. doi:10.1017/S0022109000003355
- Guthrie G. 2009. *Real Options in Theory and Practice*. Oxford University Press: New York, NY.
- Øksendal B. 1998. *Stochastic Differential Equations: An Introduction with Applications* (5th edition). Springer: Berlin.
- Rees M. 2008. *Financial Modelling in Practice: A Concise Guide for Intermediate and Advanced Level*. John Wiley & Sons: Chichester, England.
- Shockley RL. 2006. *An Applied Course in Real Options Valuation*. South-Western College Publishing: Mason, OH.

Appendix

A.1 A heuristic derivation of equations (5) and (6)

Suppose that at date t_n the manager believes that X is normally distributed with mean X_{t_n} and variance $a_{t_n}^2$. The manager then observes $Y_n = X + \varepsilon_n$, where $\varepsilon_n \sim N(0, \theta^2/(t_{n+1} - t_n))$. We calculate the linear combination of X_{t_n} and Y_n that minimizes the mean squared forecast error

$$\begin{aligned} E_{t_n} [(bX_{t_n} + cY_n - X)^2] &= E_{t_n} [(bX_{t_n} + c\varepsilon_n - (1-c)X)^2] \\ &= b^2 X_{t_n}^2 + c^2 E_{t_n}[\varepsilon_n^2] + (1-c)^2 E_{t_n}[X^2] - 2b(1-c)X_{t_n} E_{t_n}[X] \\ &= (b+c-1)^2 X_{t_n}^2 + (1-c)^2 a_{t_n}^2 + \frac{c^2 \theta^2}{t_{n+1} - t_n}. \end{aligned}$$

Evaluating the first order conditions for b and c , and then solving the resulting pair of equations, shows that the optimal weights are

$$b = \frac{\frac{\theta^2}{t_{n+1}-t_n}}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}} \quad \text{and} \quad c = \frac{a_{t_n}^2}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}}.$$

For these values of b and c , the mean squared forecast error equals

$$\frac{a_{t_n}^2 \cdot \frac{\theta^2}{t_{n+1}-t_n}}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}}.$$

Therefore, immediately after the new information is obtained, the manager believes that X is normally distributed with mean

$$X_{t_{n+1}} = \left(\frac{\frac{\theta^2}{t_{n+1}-t_n}}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}} \right) X_{t_n} + \left(\frac{a_{t_n}^2}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}} \right) Y_n$$

and variance

$$a_{t_{n+1}}^2 = \frac{a_{t_n}^2 \cdot \frac{\theta^2}{t_{n+1}-t_n}}{a_{t_n}^2 + \frac{\theta^2}{t_{n+1}-t_n}}.$$

It follows that

$$\frac{1}{a_{t_{n+1}}^2} = \frac{1}{\frac{\theta^2}{t_{n+1}-t_n}} + \frac{1}{a_{t_n}^2} = \frac{t_{n+1}-t_n}{\theta^2} + \frac{1}{a_{t_n}^2},$$

which, after recursive substitution, implies that

$$\frac{1}{a_{t_n}^2} = \frac{t_n}{\theta^2} + \frac{1}{a_0^2}.$$

Equation (5) follows immediately.

A.2 Derivation of equations (9) and (10)

Equation (3) implies that

$$\log V_{t_{n+1}} - \log V_{t_n} = (\log P_{t_{n+1}} - \log P_{t_n}) + (X_{t_{n+1}} - X_{t_n}) + \frac{1}{2} (a_{t_{n+1}}^2 - a_{t_n}^2),$$

which reduces to

$$\log V_{t_{n+1}} - \log V_{t_n} = \left(\mu - \frac{1}{2} \phi^2 \right) (t_{n+1} - t_n) + \frac{1}{2} (a_{t_{n+1}}^2 - a_{t_n}^2) + Z + W_{n+1},$$

upon substituting in the expressions for the changes in $\log P$ and X . Using equation (5) to eliminate $a_{t_n}^2$ and $a_{t_{n+1}}^2$ shows that

$$\log V_{t_{n+1}} - \log V_{t_n} = \left(\mu - \frac{1}{2} \left(\phi^2 + \frac{\theta^2}{\left(t_{n+1} + \left(\frac{\theta}{a_0} \right)^2 \right) \left(t_n + \left(\frac{\theta}{a_0} \right)^2 \right)} \right) \right) (t_{n+1} - t_n) + Z + W_{n+1}.$$

The results follow from the observation that $Z + W_{n+1}$ is normally distributed with mean zero and variance σ_n^2 .

A.3 Constructing a recombining tree for V_t

All that we need to prove is that the binomial tree described by the up-move size in equation (12) and the dates in equation (13) satisfies the condition in equation (11).

We begin by noting that t_n given in (13) satisfies the quadratic equation

$$\phi^2 t_n^2 + \left(a_0^2 + \left(\frac{\phi\theta}{a_0} \right)^2 - n(\log U)^2 \right) t_n = n(\log U)^2 \left(\frac{\theta}{a_0} \right)^2,$$

which can be rearranged to give

$$t_n \left(\phi^2 + \frac{a_0^2}{t_n + \left(\frac{\theta}{a_0} \right)^2} \right) = n(\log U)^2. \quad (\text{A-1})$$

Replacing every appearance of n in equation (A-1) with $n + 1$ implies that

$$t_{n+1} \left(\phi^2 + \frac{a_0^2}{t_{n+1} + \left(\frac{\theta}{a_0} \right)^2} \right) = (n + 1)(\log U)^2.$$

Subtracting equation (A-1) from this implies that

$$t_{n+1} \left(\phi^2 + \frac{a_0^2}{t_{n+1} + \left(\frac{\theta}{a_0} \right)^2} \right) - t_n \left(\phi^2 + \frac{a_0^2}{t_n + \left(\frac{\theta}{a_0} \right)^2} \right) = (\log U)^2,$$

which reduces to

$$(t_{n+1} - t_n) \left(\phi^2 + \frac{\theta^2}{\left(t_n + \left(\frac{\theta}{a_0} \right)^2 \right) \left(t_{n+1} + \left(\frac{\theta}{a_0} \right)^2 \right)} \right) = (\log U)^2.$$

Since the left-hand side equals $(t_{n+1} - t_n)\sigma_n^2$, this implies that equation (11) holds, as required.