

CHAPTER 8

REAL OPTIONS


The approaches that we have described in the last three chapters for assessing the effects of risk, for the most part, are focused on the negative effects of risk. Put another way, they are all focused on the downside of risk and they miss the opportunity component that provides the upside. The real options approach is the only one that gives prominence to the upside potential for risk, based on the argument that uncertainty can sometimes be a source of additional value, especially to those who are poised to take advantage of it.

We begin this chapter by describing in very general terms the argument behind the real options approach, noting its foundations in two elements – the capacity of individuals or entities to learn from what is happening around them and their willingness and the ability to modify behavior based upon that learning. We then describe the various forms that real options can take in practice and how they can affect the way we assess the value of investments and our behavior. In the last section, we consider some of the potential pitfalls in using the real options argument and how it can be best incorporated into a portfolio of risk assessment tools.

The Essence of Real Options

To understand the basis of the real options argument and the reasons for its allure, it is easiest to go back to risk assessment tool that we unveiled in chapter 6 – decision trees. Consider a very simple example of a decision tree in figure 8.1:

Figure 8.1: Simple Decision Tree




Given the equal probabilities of up and down movements, and the larger potential loss, the expected value for this investment is negative.

$$\text{Expected Value} = 0.50 (100) + 0.5 (-120) = -\$10$$

Now contrast this with the slightly more complicated two-phase decision tree in figure 8.2:

Figure 8.2: Two-phase Decision Tree



Note that the total potential profits and losses over the two phases in the tree are identical to the profit and loss of the simple tree in figure 8.1; your total gain is \$ 100 and your total loss is \$120. Note also that the cumulative probabilities of success and failure remain at the 50% that we used in the simple tree. When we compute the expected value of this tree, though, the outcome changes:

$$\text{Expected Value} = (2/3) (-10) + 1/3 [10+(2/3)(90) + (1/3)(-110)] = \$4.44$$

What is it about the second decision tree that makes a potentially bad investment (in the first tree) into a good investment (in the second)? We would attribute the change to two factors. First, by allowing for an initial phase where you get to observe the cashflows on a first and relatively small try at the investment, we allow for learning. Thus, getting a bad outcome in the first phase (-10 instead of +10) is an indicator that the overall investment is more likely to be money losing than money making. Second, you act on the learning by abandoning the investment, if the outcome from the first phase is negative; we will call this adaptive behavior.

In essence, the value of real options stems from the fact that when investing in risky assets, we can learn from observing what happens in the real world and adapting our behavior to increase our potential upside from the investment and to decrease the possible downside. Consider again the Chinese symbol for risk, as a combination of danger and opportunity that we used in chapter 1. In the real options framework, we use updated knowledge or information to expand opportunities while reducing danger. In the context of a risky investment, there are three potential actions that can be taken based

upon this updated knowledge. The first is that you build on good fortune to increase your possible profits; this is the option to expand. For instance, a market test that suggests that consumers are far more receptive to a new product than you expected them to be could be used as a basis for expanding the scale of the project and speeding its delivery to the market. The second is to scale down or even abandon an investment when the information you receive contains bad news; this is the option to abandon and can allow you to cut your losses. The third is to hold off on making further investments, if the information you receive suggests ambivalence about future prospects; this is the option to delay or wait. You are, in a sense, buying time for the investment, hoping that product and market developments will make it attractive in the future.

We would add one final piece to the mix that is often forgotten but is just as important as the learning and adaptive behavior components in terms of contributing to the real options arguments. The value of learning is greatest, when you and only you have access to that learning and can act on it. After all, the expected value of knowledge that is public, where anyone can act on that knowledge, will be close to zero. We will term this third condition “exclusivity” and use it to scrutinize when real options have the most value.

Real Options, Risk Adjusted Value and Probabilistic Assessments

Before we embark on a discussion of the options to delay, expand and abandon, it is important that we consider how the real options view of risk differs from how the approaches laid out in the last three chapters look at risk, and the implications for the valuation of risky assets.

When computing the risk-adjusted value for risky assets, we generally discount back the expected cash flows using a discount rate adjusted to reflect risk. We use higher discount rates for riskier assets and thus assign a lower value for any given set of cash flows. In the process, we are faced with the task of converting all possible outcomes in the future into one expected number. The real options critique of discounted cash flow valuation can be boiled down simply. The expected cash flows for a risky asset, where the holder of the asset can learn from observing what happens in early periods and adapting behavior, will be understated because it will not capture the diminution of the

downside risk from the option to abandon and the expansion of upside potential from the options to expand and delay. To provide a concrete example, assume that you are valuing an oil company and that you estimate the cash flows by multiplying the number of barrels of oil that you expect the company to produce each year by the expected oil price per barrel. While you may have reasonable and unbiased estimates of both these numbers (the expected number of barrels produced and the expected oil price), what you are missing in your expected cash flows is the interplay between these numbers. Oil companies can observe the price of oil and adjust production accordingly; they produce more oil when oil prices are high and less when oil prices are low. In addition, their exploration activity will ebb and flow as the oil price moves. As a consequence, their cash flows computed across all oil price scenarios will be greater than the expected cash flows used in the risk adjusted value calculation, and the difference will widen as the uncertainty about oil prices increases. So, what would real options proponents suggest? They would argue that the risk adjusted value, obtained from conventional valuation approaches, is too low and that a premium should be added to it to reflect the option to adjust production inherent in these firms.

The approach that is closest to real options in terms of incorporating adaptive behavior is the decision tree approach, where the optimal decisions at each stage are conditioned on outcomes at prior stages. The two approaches, though, will usually yield different values for the same risky asset for two reasons. The first is that the decision tree approach is built on probabilities and allows for multiple outcomes at each branch, whereas the real option approach is more constrained in its treatment of uncertainty. In its binomial version, there can be only two outcomes at each stage and the probabilities are not specified. The second is that the discount rates used to estimate present values in decision trees, at least in conventional usage, tend to be risk adjusted and not conditioned on which branch of the decision tree you are looking at. When computing the value of a diabetes drug in a decision tree, in chapter 6, we used a 10% cost of capital as the discount rate for all cash flows from the drug in both good and bad outcomes. In the real options approach, the discount rate will vary depending upon the branch of the tree being analyzed. In other words, the cost of capital for an oil companies if oil prices increase may very well be different from the cost of capital when oil prices decrease. Copeland

and Antikarov provide a persuasive proof that the value of a risky asset will be the same under real options and decision trees, if we allow for path-dependent discount rates.¹

Simulations and real options are not so much competing approaches for risk assessment, as they are complementary. Two key inputs into the real options valuation – the value of the underlying asset and the variance in that value – are often obtained from simulations. To value a patent, for instance, we need to assess the present value of cash flows from developing the patent today and the variance in that value, given the uncertainty about the inputs. Since the underlying product is not traded, it is difficult to get either of these inputs from the market. A Monte Carlo simulation can provide both values.

Real Option Examples

As we noted in the introductory section, there are three types of options embedded in investments – the option to expand, delay and abandon an investment. In this section, we will consider each of these options and how they made add value to an investment, as well as potential implications for valuation and risk management.

The Option to Delay an Investment

Investments are typically analyzed based upon their expected cash flows and discount rates at the time of the analysis; the net present value computed on that basis is a measure of its value and acceptability at that time. The rule that emerges is a simple one: negative net present value investments destroy value and should not be accepted. Expected cash flows and discount rates change over time, however, and so does the net present value. Thus, a project that has a negative net present value now may have a positive net present value in the future. In a competitive environment, in which individual firms have no special advantages over their competitors in taking projects, this may not seem significant. In an environment in which a project can be taken by only one firm

¹ Copeland, T.E. and V. Antikarov, 2003, *Real Options: A Practitioner's Guide*, Texere. For an alternate path to the same conclusion, see Brandao, L.E., J.S. Dyer and W.J. Huhn, 2005, *Using Binomial Decision Trees to Solve Real-Option Valuation Problems*, *Decision Analysis*, v2, 69-88. They use the risk-neutral probabilities from the option pricing model in the decision tree to solve for the option's value.

(because of legal restrictions or other barriers to entry to competitors), however, the changes in the project's value over time give it the characteristics of a call option.

Basic Setup

In the abstract, assume that a project requires an initial up-front investment of X , and that the present value of expected cash inflows computed right now is V . The net present value of this project is the difference between the two:

$$NPV = V - X$$

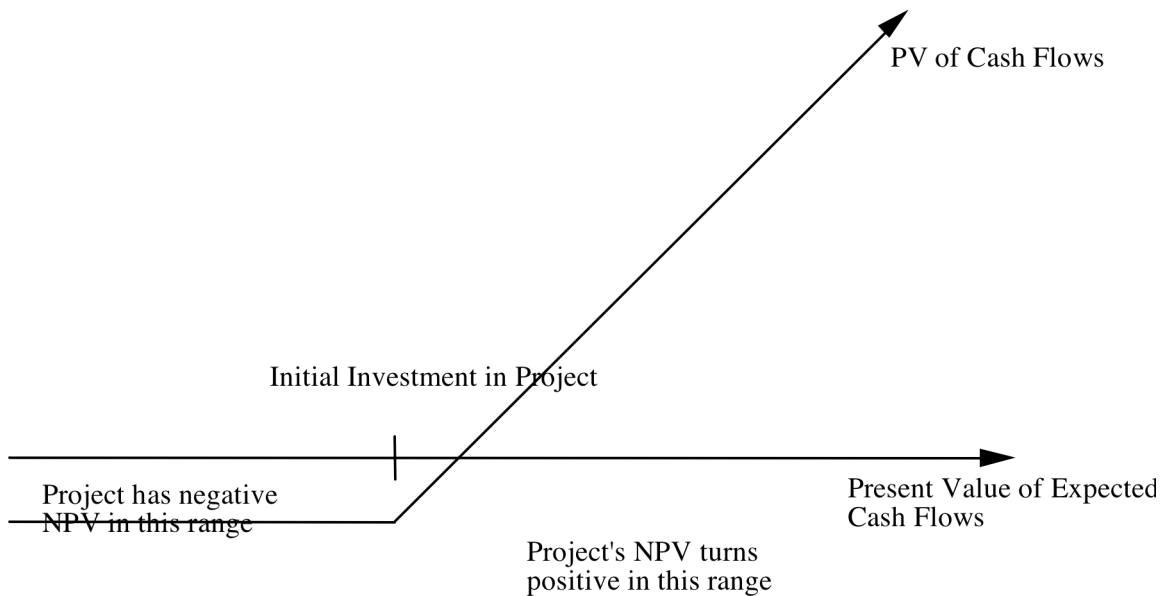
Now assume that the firm has exclusive rights to this project for the next n years, and that the present value of the cash inflows may change over that time, because of changes in either the cash flows or the discount rate. Thus, the project may have a negative net present value right now, but it may still be a good project if the firm waits. Defining V again as the present value of the cash flows, the firm's decision rule on this project can be summarized as follows:

If	$V > X$	Take the project: Project has positive net present value
	$V < X$	Do not take the project: Project has negative net present value

If the firm does not invest in the project, it incurs no additional cash flows, though it will lose what it originally invested in the project. This relationship can be presented in a payoff diagram of cash flows on this project, as shown in Figure 8.3, assuming that the firm holds out until the end of the period for which it has exclusive rights to the project:²

² McDonald, R. and D. Siegel, 2002, The Value of Waiting to Invest, *Quarterly Journal of Economics*, v101, 707-728.

Figure 8.3: The Option to Delay an Investment



Note that this payoff diagram is that of a call option — the underlying asset is the investment, the strike price of the option is the initial outlay needed to initiate the investment; and the life of the option is the period for which the firm has rights to the investment. The present value of the cash flows on this project and the expected variance in this present value represent the value and variance of the underlying asset.

Valuing an Option to Delay

On the surface, the inputs needed to value the option to delay are the same as those needed for any option. We need the value of the underlying asset, the variance in that value, the time to expiration on the option, the strike price, the riskless rate and the equivalent of the dividend yield (cost of delay). Actually estimating these inputs for a real option to delay can be difficult, however.

a. Value Of The Underlying Asset: In this case, the underlying asset is the investment itself. The current value of this asset is the present value of expected cash flows from initiating the project now, not including the up-front investment, which can be obtained by doing a standard capital budgeting analysis. There is likely to be a substantial amount of error in the cash flow estimates and the present value, however. Rather than being viewed as a problem, this uncertainty should be viewed as the reason why the project delay option has value. If the expected cash flows on the project were known with

certainty and were not expected to change, there would be no need to adopt an option pricing framework, since there would be no value to the option.

b. Variance in the value of the asset: The present value of the expected cashflows that measures the value of the asset will change over time, partly because the potential market size for the product may be unknown, and partly because technological shifts can change the cost structure and profitability of the product. The variance in the present value of cash flows from the project can be estimated in one of three ways.

- If similar projects have been introduced in the past, the variance in the cash flows from those projects can be used as an estimate. This may be the way that a consumer product company like Gillette might estimate the variance associated with introducing a new blade for its razors.
- Probabilities can be assigned to various market scenarios, cash flows estimated under each scenario and the variance estimated across present values. Alternatively, the probability distributions can be estimated for each of the inputs into the project analysis - the size of the market, the market share and the profit margin, for instance - and simulations used to estimate the variance in the present values that emerge.
- The variance in the market value of publicly traded firms involved in the same business (as the project being considered) can be used as an estimate of the variance. Thus, the average variance in firm value of firms involved in the software business can be used as the variance in present value of a software project.

The value of the option is largely derived from the variance in cash flows - the higher the variance, the higher the value of the project delay option. Thus, the value of an option to delay a project in a stable business will be less than the value of a similar option in an environment where technology, competition and markets are all changing rapidly.

c. Exercise Price On Option: A project delay option is exercised when the firm owning the rights to the project decides to invest in it. The cost of making this investment is the exercise price of the option. The underlying assumption is that this cost remains constant (in present value dollars) and that any uncertainty associated with the product is reflected in the present value of cash flows on the product.

d. Expiration Of The Option And The Riskless Rate The project delay option expires when the rights to the project lapse; investments made after the project rights expire are

assumed to deliver a net present value of zero as competition drives returns down to the required rate. The riskless rate to use in pricing the option should be the rate that corresponds to the expiration of the option. While this input can be estimated easily when firms have the explicit right to a project (through a license or a patent, for instance), it becomes far more difficult to obtain when firms only have a competitive advantage to take a project.

d. Cost of Delay (Dividend Yield): There is a cost to delaying taking a project, once the net present value turns positive. Since the project rights expire after a fixed period, and excess profits (which are the source of positive present value) are assumed to disappear after that time as new competitors emerge, each year of delay translates into one less year of value-creating cash flows.³ If the cash flows are evenly distributed over time, and the exclusive rights last n years, the cost of delay can be written as:

$$\text{Annual cost of delay} = \frac{1}{n}$$

Thus, if the project rights are for 20 years, the annual cost of delay works out to 5% a year. Note, though, that this cost of delay rises each year, to 1/19 in year 2, 1/18 in year 3 and so on, making the cost of delaying exercise larger over time.

Practical Considerations

While it is quite clear that the option to delay is embedded in many investments, there are several problems associated with the use of option pricing models to value these options. First, the *underlying asset in this option, which is the project, is not traded*, making it difficult to estimate its value and variance. We would argue that the value can be estimated from the expected cash flows and the discount rate for the project, albeit with error. The variance is more difficult to estimate, however, since we are attempting the estimate a variance in project value over time.

Second, the *behavior of prices over time may not conform to the price path assumed by the option pricing models*. In particular, the assumption that prices move in small increments continuously (an assumption of the Black-Scholes model), and that the

³ A value-creating cashflow is one that adds to the net present value because it is in excess of the required return for investments of equivalent risk.

variance in value remains unchanged over time, may be difficult to justify in the context of a real investment. For instance, a sudden technological change may dramatically change the value of a project, either positively or negatively.

Third, *there may be no specific period for which the firm has rights to the project.* For instance, a firm may have significant advantages over its competitors, which may, in turn, provide it with the virtually exclusive rights to a project for a period of time. The rights are not legal restrictions, however, and could erode faster than expected. In such cases, the expected life of the project itself is uncertain and only an estimate. Ironically, uncertainty about the expected life of the option can increase the variance in present value, and through it, the expected value of the rights to the project.

Applications of Option to Delay

The option to delay provides interesting perspectives on two common investment problems. The first is in the valuation of patents, especially those that are not viable today but could be viable in the future; by extension, this will also allow us to look at whether R&D expenses are delivering value. The second is in the analysis of natural resource assets – vacant land, undeveloped oil reserves etc.

Patents

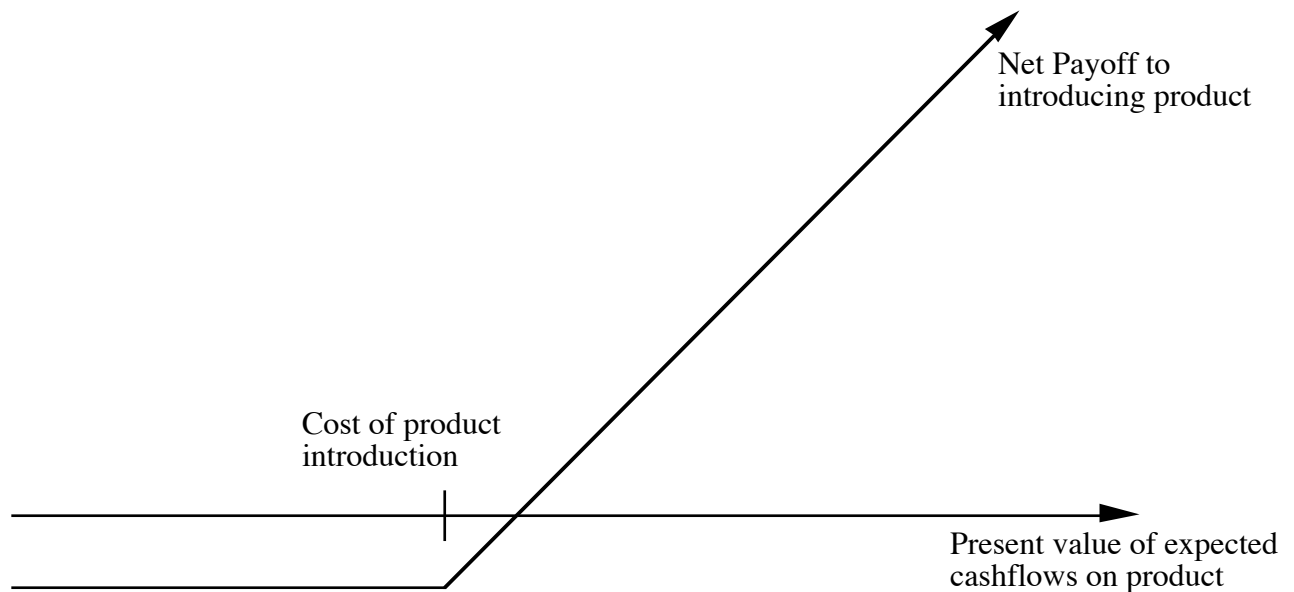
A product patent provides a firm with the right to develop and market a product. The firm will do so only if the present value of the expected cash flows from the product sales exceed the cost of development, however, as shown in Figure 8.4. If this does not occur, the firm can shelve the patent and not incur any further costs. If I is the present value of the costs of developing the product, and V is the present value of the expected cash flows from development, the payoffs from owning a product patent can be written as:

$$\begin{aligned} \text{Payoff from owning a product patent} &= V - I && \text{if } V > I \\ &= 0 && \text{if } V \leq I \end{aligned}$$

Thus, a product patent can be viewed as a call option, where the product itself is the underlying asset.⁴

⁴ Schwartz, E., 2002, Patents and R&D as Real Options, Working Paper, Anderson School at UCLA.

Figure 8.4: Payoff to Introducing Product



We will illustrate the use of option pricing to value Avonex, a drug to treat multiple sclerosis, right after it had received FDA approval in 1997, but before its parent company, Biogen, had decided whether to commercialize the drug or not. We arrived at the following estimates for use in the option pricing model:

- An internal analysis of the drug at the time, based upon the potential market and the price that the firm can expect to charge, yielded a present value of expected cash flows of \$ 3.422 billion, prior to considering the initial development cost.
- The initial cost of developing the drug for commercial use was estimated to be \$2.875 billion, if the drug was introduced immediately.
- The firm had the patent on the drug for the next 17 years, and the 17-year Treasury bond rate was 6.7%.
- The average historical variance in market value for publicly traded bio-technology firms was 0.224.
- It was assumed that the potential for excess returns exists only during the patent life, and that competition will wipe out excess returns beyond that period. Thus, any delay in introducing the drug, once it is viable, will cost the firm one year of patent-protected excess returns. (For the initial analysis, the cost of delay will be 1/17, the following year, it will be 1/16, the year after, 1/15 and so on.)

Based on these assumptions, we obtained the following inputs to the option pricing model.

Present Value of Cash Flows from Introducing Drug Now = $S = \$ 3.422$ billion

Initial Cost of Developing Drug for commercial use = $K = \$ 2.875$ billion

Patent life = $t = 17$ years Riskless Rate = $r = 6.7\%$ (17-year T.Bond rate)

Variance in Expected Present Values = $\sigma^2 = 0.224$

Expected Cost of Delay = $y = 1/17 = 5.89\%$

Using these inputs in an option pricing model, we derived a value of \$907 million for the option,⁵ and this can be considered to be the real options value attached to the patent on Avonex. To provide a contrast, the net present value of this patent is only \$ 547 million:

$$\text{NPV} = \$3,422 \text{ million} - \$ 2,875 \text{ million} = \$ 547 \text{ million}$$

The time premium of \$ 360 million (\$907 million -\$547 million) on this option suggests that the firm will be better off waiting rather than developing the drug immediately, the cost of delay notwithstanding. However, the cost of delay will increase over time, and make exercise (development) more likely. Note also that we are assuming that the firm is protected from all competition for the life of the patent. In reality, there are other pharmaceutical firms working on their own drugs to treat multiple sclerosis and that can affect both the option value and the firm's behavior. In particular, if we assume that Upjohn or Pfizer has a competing drug working through the FDA pipeline and that the drug is expected to reach the market in 6 years, the cost of delay will increase to 16.67% (1/6) and the option value will dissipate.

The implications of viewing patents as options can be significant. First, it implies that non-viable patents will continue to have value, especially in businesses where there is substantial volatility. Second, it indicates that firms may hold off on developing viable patents, if they feel that they gain more from waiting than they lose in terms of cash flows; this behavior will be more common if there is no significant competition on the horizon. Third, the value of patents will be higher in risky businesses than in safe businesses, since option value increases with volatility. If we consider R&D to be the expense associated with acquiring these patents, this would imply that research should

have its biggest payoff when directed to areas where less is known and there is more uncertainty. Consequently, we should expect pharmaceutical firms to spend more of their R&D budgets on gene therapy than on flu vaccines.⁶

Natural Resource Options

In a natural resource investment, the underlying asset is the natural resource and the value of the asset is based upon two variables - (1) the estimated quantity, and (2) the price of the resource. Thus, in a gold mine, for example, the value of the underlying asset is the value of the estimated gold reserves in the mine, based upon the current price of gold. In most such investments, there is an initial cost associated with developing the resource; the difference between the value of the asset extracted and the cost of the development is the profit to the owner of the resource (see Figure 8.5). Defining the cost of development as X , and the estimated value of the developed resource as V , the potential payoffs on a natural resource option can be written as follows:

$$\begin{aligned} \text{Payoff on natural resource investment} &= V - X && \text{if } V > X \\ &= 0 && \text{if } V \leq X \end{aligned}$$

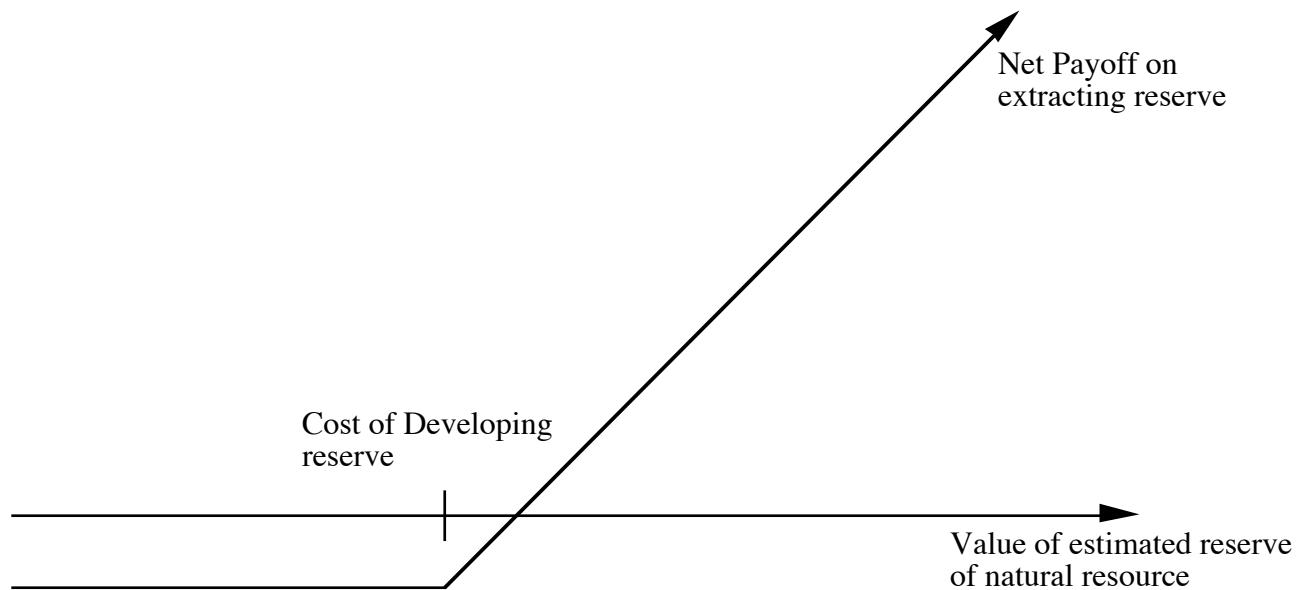
Thus, the investment in a natural resource option has a payoff function similar to a call option.⁷

⁵ This value was derived from using a Black Scholes model with these inputs. With a binomial model, the estimated value increases slightly to \$915 million.

⁶ Pakes, A., 1986, Patents as Options: Some Estimates of the Value of Holding European Patent Stocks, *Econometrica*, v54, 755-784. While this paper does not explicitly value patents as options, it examines the returns investors would have earned investing in companies that derive their value from patents. The return distribution resembles that of a portfolio of options, with most investments losing money but the winners providing disproportionate gains.

⁷ Brennan, M. and E. Schwartz, 1985, Evaluating Natural Resource Investments, *The Journal of Business*, v58, 135-157.

Figure 8.5: Payoff from Developing Natural Resource Reserves



To value a natural resource investment as an option, we need to make assumptions about a number of variables:

1. Available reserves of the resource: Since this is not known with certainty at the outset, it has to be estimated. In an oil tract, for instance, geologists can provide reasonably accurate estimates of the quantity of oil available in the tract.

2. Estimated cost of developing the resource: The estimated development cost is the exercise price of the option. Again, a combination of knowledge about past costs and the specifics of the investment have to be used to come up with a reasonable measure of development cost.

3. Time to expiration of the option: The life of a natural resource option can be defined in one of two ways. First, if the ownership of the investment has to be relinquished at the end of a fixed period of time, that period will be the life of the option. In many offshore oil leases, for instance, the oil tracts are leased to the oil company for several years. The second approach is based upon the inventory of the resource and the capacity output rate, as well as estimates of the number of years it would take to exhaust the inventory. Thus, a gold mine with a mine inventory of 3 million ounces and a capacity output rate of 150,000 ounces a year will be exhausted in 20 years, which is defined as the life of the natural resource option.

4. Variance in value of the underlying asset: The variance in the value of the underlying asset is determined by two factors – (1) variability in the price of the resource, and (2) variability in the estimate of available reserves. In the special case where the quantity of the reserve is known with certainty, the variance in the underlying asset's value will depend entirely upon the variance in the price of the natural resource. In the more realistic case where the quantity of the reserve and the oil price can change over time, the option becomes more difficult to value; here, the firm may have to invest in stages to exploit the reserves.

5. Cost of Delay: The net production revenue as a percentage of the market value of the reserve is the equivalent of the dividend yield and is treated the same way in calculating option values. An alternative way of thinking about this cost is in terms of a cost of delay. Once a natural resource option is in-the-money (Value of the reserves > Cost of developing these reserves), the firm, by not exercising the option, is costing itself the production revenue it could have generated by developing the reserve.

An important issue in using option pricing models to value natural resource options is the effect of development lags on the value of these options. Since the resources cannot be extracted instantaneously, a time lag has to be allowed between the decision to extract the resources and the actual extraction. A simple adjustment for this lag is to reduce the value of the developed reserve to reflect the loss of cash flows during the development period. Thus, if there is a one-year lag in development, the current value of the developed reserve will be discounted back one year at the net production revenue/asset value ratio⁸ (which we also called the dividend yield above).⁹

To illustrate the use of option pricing to value natural reserves, consider an offshore oil property with an estimated reserve of 50 million barrels of oil; the cost of developing the reserve is expected to be \$ 600 million, and the development lag is two years. The firm has the rights to exploit this reserve for the next 20 years, and the

⁸ Intuitively, it may seem like the discounting should occur at the riskfree rate. The simplest way of explaining why we discount at the dividend yield is to consider the analogy with a listed option on a stock. Assume that on exercising a listed option on a stock, you had to wait six months for the stock to be delivered to you. What you lose is the dividends you would have received over the six-month period by holding the stock. Hence, the discounting is at the dividend yield.

⁹ Brennan, M.J., and E.S. Schwartz, 1985, Evaluating Natural Resource Investments, Journal of Business 58, pp. 135-157.

marginal value per barrel of oil is \$12 currently¹⁰ (price per barrel - marginal cost per barrel). Once developed, the net production revenue each year will be 5% of the value of the reserves. The riskless rate is 8%, and the variance in $\ln(\text{oil prices})$ is 0.03. Given this information, the inputs to the option pricing model can be estimated as follows:

Current Value of the asset = $S = \text{Value of the developed reserve discounted back the length of the development lag at the dividend yield} = \$12 * 50 / (1.05)^2 = \$544.22$

If development is started today, the oil will not be available for sale until two years from now. The estimated opportunity cost of this delay is the lost production revenue over the delay period; hence, the discounting of the reserve back at the dividend yield.

Exercise Price = Cost of developing reserve = \$ 600 million (assumed to be both known and fixed over time)

Time to expiration on the option = 20 years

In this example, we assume that the only uncertainty is in the price of oil, and the variance therefore becomes the variance in oil prices.

Variance in the value of the underlying asset (oil) = 0.03

Riskless rate = 8%

Dividend Yield = Net production revenue / Value of reserve = 5%

Based upon these inputs, the option pricing model yields an estimate of value of \$97.08 million.¹¹ This oil reserve, though not viable at current prices, is still a valuable property because of its potential to create value if oil prices go up.¹²

The same type of analysis can be extended to any other commodity company (gold and copper reserves, for instance) and even to vacant land or real estate properties.

¹⁰ For simplicity, we will assume that while this marginal value per barrel of oil will grow over time, the present value of the marginal value will remain unchanged at \$ 12 per barrel. If we do not make this assumption, we will have to estimate the present value of the oil that will be extracted over the extraction period.

¹¹ This is the estimate from a Black-Scholes model, with a dividend yield adjustment. Using a binomial model yields an estimate of value of \$ 101 million.

¹² Paddock, J.L. & D. R. Siegel & J.L. Smith (1988): "Option Valuation of Claims on Real Assets: The Case of Offshore Petroleum Leases", *Quarterly Journal of Economics*, August 1988, pp.479-508. This paper provides a detailed examination of the application of real options to value oil reserves. They applied the model to examine the prices paid for offshore oil leases in the US in 1980 and concluded that companies over paid (relative to the option value).

The owner of vacant land in Manhattan can choose whether and when to develop the land and will make that decision based upon real estate values.¹³

What are the implications of viewing natural resource reserves as options? The first is that the value of a natural resource company can be written as a sum of two values: the conventional risk adjusted value of expected cash flows from developed reserves and the option value of undeveloped reserves. While both will increase in value as the price of the natural resource increases, the latter will respond positively to increases in price volatility. Thus, the values of oil companies should increase if oil prices become more volatile, even if oil prices themselves do not go up. The second is that conventional discounted cash flow valuation will understate the value of natural resource companies, even if the expected cash flows are unbiased and reasonable because it will miss the option premium inherent in their undeveloped reserves. The third is that development of natural resource reserves will slow down as the volatility in prices increases; the time premium on the options will increase, making exercise of the options (development of the reserves) less likely.

Mining and commodity companies have been at the forefront in using real options in decision making and their usage of the technology predates the current boom in real options. One reason is that natural resource options come closest to meeting the pre-requisites for the use of option pricing models. Firms can learn a great deal by observing commodity prices and can adjust their behavior (in terms of development and exploration) quickly. In addition, if we consider exclusivity to be a pre-requisite for real options to have value, that exclusivity for natural resource options derives from their natural scarcity; there is, after all, only a finite amount of oil and gold under the ground and vacant land in Manhattan. Finally, natural resource reserves come closest to meeting the arbitrage/replication requirements that option pricing models are built upon; both the underlying asset (the natural resource) and the option can often be bought and sold.

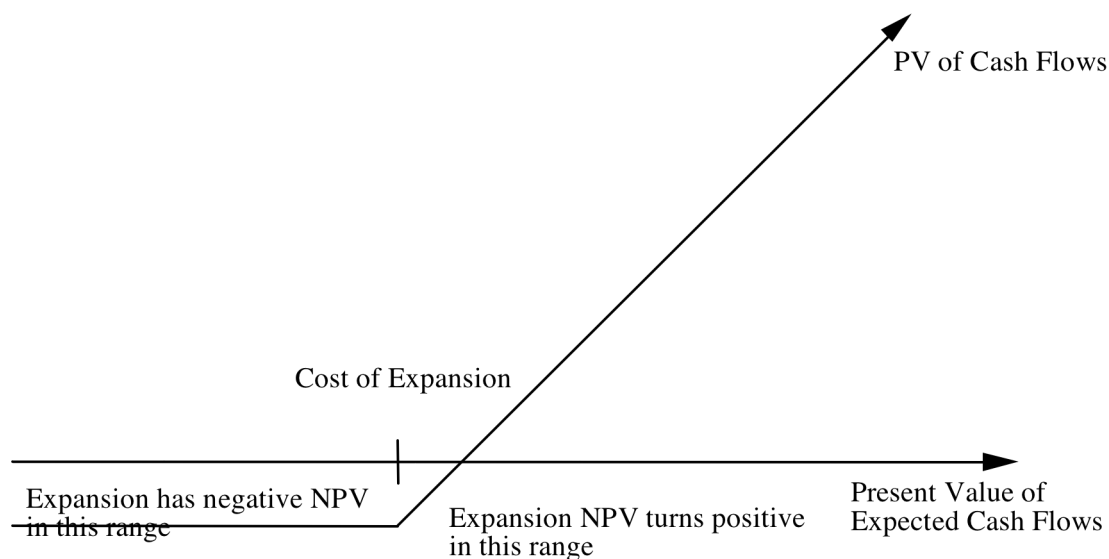
¹³ Quigg, L, 1993] Empirical Testing of Real Option-Pricing Models », *Journal of Finance*, vol.48, 621-640. The author examined transaction data on 2700 undeveloped and 3200 developed real estate properties between 1976-79 and found evidence of a premium arising from the option to wait in the former.

The Option to Expand

In some cases, a firm will take an investment because doing so allows it either to make other investments or to enter other markets in the future. In such cases, it can be argued that the initial investment provides the firm with an option to expand, and the firm should therefore be willing to pay a price for such an option. Consequently, a firm may be willing to lose money on the first investment because it perceives the option to expand as having a large enough value to compensate for the initial loss.

To examine this option, assume that the present value of the expected cash flows from entering the new market or taking the new project is V , and the total investment needed to enter this market or take this project is X . Further, assume that the firm has a fixed time horizon, at the end of which it has to make the final decision on whether or not to take advantage of this opportunity. Finally, assume that the firm cannot move forward on this opportunity if it does not take the initial investment. This scenario implies the option payoffs shown in Figure 8.6.

Figure 8.6: The Option to Expand a Project



As you can see, at the expiration of the fixed time horizon, the firm will enter the new market or take the new investment if the present value of the expected cash flows at that point in time exceeds the cost of entering the market.

Consider a simple example of an option to expand. Disney is considering starting a Spanish version of the Disney Channel in Mexico and estimates the net present value of this investment to be -\$150 million. While the negative net present value would normally suggest that rejecting the investment is the best course, assume that if the Mexican venture does better than expected, Disney plans to expand the network to the rest of South America at a cost of \$ 500 million. Based on its current assessment of this market, Disney believes that the present value of the expected cash flows on this investment is only \$ 400 million (making it a negative net present value investment as well). The saving grace is that the latter present value is an estimate and Disney does not have a firm grasp of the market; a Monte Carlo simulation of the investments yields a standard deviation of 50% in value. Finally, assume that Disney will have to make this expansion decision within 5 years of the Mexican investment, and that the five-year riskfree rate is 4%. The value of the expansion option can now be computed using the inputs:

$S = \text{Present value of expansion cash flows} = \$ 400 \text{ million}$

$K = \text{Cost of expansion} = \$ 500 \text{ million}$

$\sigma = \text{Standard deviation in value (from simulation)} = 50\%$

$t = 5 \text{ years}$

$r = 4\%$

The resulting option value is \$167 million.¹⁴

The practical considerations associated with estimating the value of the option to expand are similar to those associated with valuing the option to delay. In most cases, firms with options to expand have no specific time horizon by which they have to make an expansion decision, making these open-ended options, or, at best, options with arbitrary lives. Even in those cases where a life can be estimated for the option, neither the size nor the potential market for the product may be known, and estimating either can be problematic. To illustrate, consider the Disney example discussed above. While we adopted a period of five years, at the end of which the Disney has to decide one way or another on its future expansion into South America, it is entirely possible that this time frame is not specified at the time the store is opened. Furthermore, we have assumed that

¹⁴ This value was computed using the Black-Scholes model. A binomial model yields a similar value.

both the cost and the present value of expansion are known initially. In reality, the firm may not have good estimates for either before making the first investment, since it does not have much information on the underlying market.

Implications

The option to expand is implicitly used by firms to rationalize taking investments that have negative net present value, but provide significant opportunities to tap into new markets or sell new products. While the option pricing approach adds rigor to this argument by estimating the value of this option, it also provides insight into those occasions when it is most valuable. In general, the option to expand is clearly more valuable for more volatile businesses with higher returns on projects (such as biotechnology or computer software), than in stable businesses with lower returns (such as housing, utilities or automobile production). Specifically, the option to expand is at the basis of arguments that an investment should be made because of strategic considerations or that large investments should be broken up into smaller phases. It can also be considered a rationale for why firms may accumulate cash or hold back on borrowing, thus preserving financial flexibility.

Strategic Considerations

In many acquisitions or investments, the acquiring firm believes that the transaction will give it competitive advantages in the future. These competitive advantages range the gamut, and include:

- Entrée into a Growing or Large Market: An investment or acquisition may allow the firm to enter a large or potentially large market much sooner than it otherwise would have been able to do so. A good example of this would be the acquisition of a Mexican retail firm by a US firm, with the intent of expanding into the Mexican market.
- Technological Expertise: In some cases, the acquisition is motivated by the desire to acquire a proprietary technology, that will allow the acquirer to expand either its existing market or into a new market.

- Brand Name: Firms sometime pay large premiums over market price to acquire firms with valuable brand names, because they believe that these brand names can be used for expansion into new markets in the future.

While all of these potential advantages may be used to justify initial investments that do not meet financial benchmarks, not all of them create valuable options. The value of the option is derived from the degree to which these competitive advantages, assuming that they do exist, translate into sustainable excess returns. As a consequence, these advantages can be used to justify premiums only in cases where the acquiring firm believes that it has some degree of exclusivity in the targeted market or technology. Two examples can help illustrate this point. A telecommunications firm should be willing to pay a premium for Chinese telecomm firm, if the latter has exclusive rights to service a large segment of the Chinese market; the option to expand in the Chinese market could be worth a significant amount.¹⁵ On the other hand, a developed market retailer should be wary about paying a real option premium for an Indian retail firm, even though it may believe that the Indian market could grow to be a lucrative one. The option to expand into this lucrative market is open to all entrants and not just to existing retailers and thus may not translate into sustainable excess returns.

Multi-Stage Projects/ Investments

When entering new businesses or making new investments, firms sometimes have the option to enter the business in stages. While doing so may reduce potential upside, it also protects the firm against downside risk, by allowing it, at each stage, to gauge demand and decide whether to go on to the next stage. In other words, a standard project can be recast as a series of options to expand, with each option being dependent on the previous one. There are two propositions that follow:

- Some projects that do not look good on a full investment basis may be value creating if the firm can invest in stages.
- Some projects that look attractive on a full investment basis may become even more attractive if taken in stages.

The gain in value from the options created by multi-stage investments has to be weighed off against the cost. Taking investments in stages may allow competitors who decide to enter the market on a full scale to capture the market. It may also lead to higher costs at each stage, since the firm is not taking full advantage of economies of scale.

There are several implications that emerge from viewing this choice between multi-stage and one-time investments in an option framework. The projects where the gains will be largest from making the investment in multiple stages include:

- (1) Projects where there are significant barriers to entry from competitors entering the market, and taking advantage of delays in full-scale production. Thus, a firm with a patent on a product or other legal protection against competition pays a much smaller price for starting small and expanding as it learns more about the product
- (2) Projects where there is significant uncertainty about the size of the market and the eventual success of the project. Here, starting small and expanding allows the firm to reduce its losses if the product does not sell as well as anticipated, and to learn more about the market at each stage. This information can then be useful in subsequent stages in both product design and marketing. Hsu argues that venture capitalists invest in young companies in stages, partly to capture the value of option of waiting/learning at each stage and partly to reduce the likelihood that the entrepreneur will be too conservative in pursuing risky (but good) opportunities.¹⁶
- (3) Projects where there is a substantial investment needed in infrastructure (large fixed costs) and high operating leverage. Since the savings from doing a project in multiple stages can be traced to investments needed at each stage, they are likely to be greater in firms where those costs are large. Capital intensive projects as well as projects that require large initial marketing expenses (a new brand name product for a consumer product company) will gain more from the options created by taking the project in multiple stages.

¹⁵ A note of caution needs to be added here. If the exclusive rights to a market come with no pricing power – in other words, the Government will set the price you charge your customers – it may very well translate into zero excess returns (and no option value).

¹⁶ Hsu, Y., 2002, Staging of Venture Capital Investment: A Real Options Analysis, Working paper, University of Cambridge.

Growth Companies

In the stock market boom in the 1990s, we witnessed the phenomenon of young, start-up, internet companies with large market capitalizations but little to show in terms of earnings, cash flows or even revenues. Conventional valuation models suggested that it would be difficult, if not impossible, to justify these market valuations with expected cash flows. In an interesting twist on the option to expand argument, there were some who argued that investors in these companies were buying options to expand and be part of a potentially huge e-commerce market, rather than conventional stock.¹⁷

While the argument is alluring and serves to pacify investors in growth companies who may feel that they are paying too much, there are clearly dangers in making this stretch. The biggest one is that the “exclusivity” component that is necessary for real options to have value is being given short shrift. Consider investing in an internet stock in 1999 and assume that you are paying a premium to be part of a potentially large online market in 2008. Assume further that this market comes to fruition. Could you partake in this market without paying that upfront premium a dot-com company? We don’t see why not. After all, GE and Nokia are just as capable of being part of this online market, as are any number of new entrants into the market.¹⁸

Financial Flexibility

When making decisions about how much to borrow and how much cash to return to stockholders (in dividends and stock buybacks), managers should consider the effects such decisions will have on their capacity to make new investments or meet unanticipated contingencies in future periods. Practically, this translates into firms maintaining excess debt capacity or larger cash balances than are warranted by current needs, to meet unexpected future requirements. While maintaining this financing flexibility has value to firms, it also has costs; the large cash balances might earn below market returns, and excess debt capacity implies that the firm is giving up some value by maintaining a higher cost of capital.

¹⁷ Schwartz, E.S. and M. Moon, 2001, Rational Pricing of Internet Companies Revisited, *The Financial Review* 36, pp. 7-26. A simpler version of the same argument was made in Mauboussin, M., 1998, *Get Real: Using Real Options in Security Analysis*, CSFB Publication, June 23, 1999.

¹⁸ This argument is fleshed out in my book, “*The Dark Side of Valuation*”, published by Prentice-Hall.

Using an option framework, it can be argued that a firm that maintains a large cash balance and preserves excess debt capacity is doing so to have the option to take unexpected projects with high returns that may arise in the future. To value financial flexibility as an option, consider the following framework: A firm has expectations about how much it will need to reinvest in future periods, based upon its own past history and current conditions in the industry. On the other side of the ledger, a firm also has expectations about how much it can raise from internal funds and its normal access to capital markets in future periods. Assume that there is actual reinvestment needs can be very different from the expected reinvestment needs; for simplicity, we will assume that the capacity to generate funds is known to the firm. The advantage (and value) of having excess debt capacity or large cash balances is that the firm can meet any reinvestment needs in excess of funds available using its excess debt capacity and surplus cash. The payoff from these projects, however, comes from the excess returns that the firm expects to make on them.

Looking at financial flexibility as an option yields valuable insights on when financial flexibility is most valuable. Using the framework developed above, for instance, we would argue that:

- Other things remaining equal, firms operating in businesses where projects earn substantially higher returns than their hurdle rates should value flexibility more than those that operate in stable businesses where excess returns are small. This would imply that firms that earn large excess returns on their projects can use the need for financial flexibility as the justification for holding large cash balances and excess debt capacity.
- Since a firm's ability to fund these reinvestment needs is determined by its capacity to generate internal funds, other things remaining equal, financial flexibility should be worth less to firms with large and stable earnings, as a percent of firm value. Young and growing firms that have small or negative earnings, and therefore much lower capacity to generate internal funds, will value flexibility more. As supporting evidence, note that technology firms usually borrow very little and accumulate large cash balances.

- Firms with limited internal funds can still get away with little or no financial flexibility if they can tap external markets for capital – bank debt, bonds and new equity issues. Other things remaining equal, the greater the capacity (and willingness) of a firm to raise funds from external capital markets, the less should be the value of flexibility. This may explain why private or small firms, which have far less access to capital, will value financial flexibility more than larger firms. The existence of corporate bond markets can also make a difference in how much flexibility is valued. In markets where firms cannot issue bonds and have to depend entirely upon banks for financing, there is less access to capital and a greater need to maintain financial flexibility.
- The need for and the value of flexibility is a function of how uncertain a firm is about future reinvestment needs. Firms with predictable reinvestment needs should value flexibility less than firms in sectors where reinvestment needs are volatile on a period-to-period basis.

In conventional corporate finance, the optimal debt ratio is the one that minimizes the cost of capital and there is little incentive for firms to accumulate cash balances. This view of the world, though, flows directly from the implicit assumption we make that capital markets are open and can be accessed with little or no cost. Introducing external capital constraints, internal or external, into the model leads to a more nuanced analysis where rational firms may borrow less than optimal and hold back on returning cash to stockholders.

The Option to Abandon an Investment


The final option to consider here is the option to abandon a project when its cash flows do not measure up to expectations. One way to reflect this value is through decision trees, as evidenced in chapter 6. The decision tree has limited applicability in most real world investment analyses; it typically works only for multi-stage projects, and it requires inputs on probabilities at each stage of the project. The option pricing approach provides a more general way of estimating and building in the value of abandonment into investment analysis. To illustrate, assume that V is the remaining value on a project if it continues to the end of its life, and L is the liquidation or abandonment value for the same

project at the same point in time. If the project has a life of n years, the value of continuing the project can be compared to the liquidation (abandonment) value. If the value from continuing is higher, the project should be continued; if the value of abandonment is higher, the holder of the abandonment option could consider abandoning the project .

$$\begin{aligned} \text{Payoff from owning an abandonment option} &= 0 && \text{if } V > L \\ &= L - V && \text{if } V \leq L \end{aligned}$$

These payoffs are graphed in Figure 8.8, as a function of the expected value from continuing the investment.

Figure 8.8: The Option to Abandon a Project



Unlike the prior two cases, the option to abandon takes on the characteristics of a put option.

Consider a simple example. Assume that a firm is considering taking a 10-year project that requires an initial investment of \$ 100 million in a real estate partnership, where the present value of expected cash flows is \$ 110 million. While the net present value of \$ 10 million is small, assume that the firm has the option to abandon this project anytime in the next 10 years, by selling its share of the ownership to the other partners in the venture for \$ 50 million. Assume that the variance in the present value of the expected cash flows from being in the partnership is 0.09.

The value of the abandonment option can be estimated by determining the characteristics of the put option:

Value of the Underlying Asset (S) = PV of Cash Flows from Project
 = \$ 110 million

Strike Price (K) = Salvage Value from Abandonment = \$ 50 million

Variance in Underlying Asset's Value = 0.06

Time to expiration = Period for which the firm has abandonment option = 10 years

The project has a 25-year life and is expected to lose value each year; for simplicity, we will assume that the loss is linear (4% a year).

Loss in value each year = $1/n = 1/25 = 4\%$

Assume that the ten-year riskless rate is 6%. The value of the put option can be estimated as follows:

Call Value = $110 \exp^{(-.04)(10)} (0.9737) - 50 (\exp^{(-0.06)(10)} (0.8387)) = \$ 84.09$ million

Put Value = $\$ 84.09 - 110 + 50 \exp^{(-0.06)(10)} = \$ 1.53$ million

The value of this abandonment option has to be added on to the net present value of the project of \$ 10 million, yielding a total net present value with the abandonment option of \$ 11.53 million. Note though that abandonment becomes a more and more attractive option as the remaining project life decreases, since the present value of the remaining cash flows will decrease.

In the above analysis, we assumed, rather unrealistically, that the abandonment value was clearly specified up front and that it did not change during the life of the project. This may be true in some very specific cases, in which an abandonment option is built into the contract. More often, however, the firm has the option to abandon, and the salvage value from doing so has to be estimated (with error) up front. Further, the abandonment value may change over the life of the project, making it difficult to apply traditional option pricing techniques. Finally, it is entirely possible that abandoning a project may not bring in a liquidation value, but may create costs instead; a manufacturing firm may have to pay severance to its workers, for instance. In such cases, it would not make sense to abandon, unless the present value of the expected cash flows from continuing with the investment are even more negative.

Implications

The fact that the option to abandon has value provides a rationale for firms to build in operating flexibility to scale back or terminate projects if they do not measure up to expectations. It also indicates that firms that focus on generating more revenues by offering their customers the option to walk away from commitments may be giving up more than they gain, in the process.

1. Escape Clauses

When a firm enters into a long term risky investment that requires a large up front investment, it should do so with the clear understanding that it may regret making this investment fairly early in its life. Being able to get out of such long-term commitments that threaten to drain more resources in the future is at the heart of the option to abandon. It is true that some of this flexibility is determined by the business that you are in; getting out of bad investments is easier to do in service businesses than in heavy infrastructure businesses. However, it is also true that there are actions that firms can take at the time of making these investments that give them more choices, if things do not go according to plan.

The first and most direct way is to build operating flexibility contractually with those parties that are involved in the investment. Thus, contracts with suppliers may be written on an annual basis, rather than long term, and employees may be hired on a temporary basis, rather than permanently. The physical plant used for a project may be leased on a short-term basis, rather than bought, and the financial investment may be made in stages rather than as an initial lump sum. While there is a cost to building in this flexibility, the gains may be much larger, especially in volatile businesses. The initial capital investment can be shared with another investor, presumably with deeper pockets and a greater willingness to stay with the investment, even if it turns sour. This provides a rationale for joint venture investing, especially for small firms that have limited resources; finding a cash-rich, larger company to share the risk may well be worth the cost.

None of these actions are costless. Entering into short term agreements with suppliers and leasing the physical plant may be more expensive than committing for the

life of the investment, but that additional cost has to be weighed off against the benefit of maintaining the abandonment option.

2. Customer Incentives

Firms that are intent on increasing revenues sometimes offer abandonment options to customers to induce them to buy their products and services. As an example, consider a firm that sells its products on multi-year contracts and offers customers the option to cancel their contracts at any time, with no cost. While this may sweeten the deal and increase sales, there is likely to be a substantial cost. In the event of a recession, customers that are unable to meet their obligations are likely to cancel their contracts. In effect, the firm has made its good times better and its bad times worse; the cost of this increased volatility in earnings and revenues has to be measured against the potential gain in revenue growth to see if the net effect is positive.

This discussion should also act as a cautionary note for those firms that are run with marketing objectives such as maximizing market share or posting high revenue growth. Those objectives can often be accomplished by giving valuable options to customers – sales people will want to meet their sales targets and are not particularly concerned about the long term costs they may create with their commitments to customers – and the firm may be worse off as a consequence.

3. Switching Options

While the abandonment option considers the value of shutting an investment down entirely, there is an intermediate alternative that is worth examining. Firms can sometimes alter production levels in response to demand and being able to do so can make an investment more valuable. Consider, for instance, a power company that is considering a new plant to generate electricity and assume that the company can run the plant at full capacity and produce 1 million kilowatt hours of power or at half capacity (and substantially less cost) and produce 500,000 kilowatt hours of power. In this case, the company can observe both the demand for power and the revenues per kilowatt-hour and decide whether it makes sense to run at full or half capacity. The value of this switching option can then be compared to the cost of building in this flexibility in the first place.

The airline business provides an interesting case study in how different companies manage their cost structure and the payoffs to their strategies. One reason that Southwest Airlines has been able to maintain its profitability in a deeply troubled sector is that the company has made cost flexibility a central component in its decision process. From its choice of using only one type of aircraft for its entire fleet¹⁹ to its refusal, for the most part, to fly into large urban airports (with high gate costs), the company's operations have created the most flexible cost structure in the business. Thus, when revenues dip (as they inevitably do at some point in time when the economy weakens), Southwest is able to trim its costs and stay profitable while other airlines teeter on the brink of bankruptcy.

Caveats on Real Options

The discussion on the potential applications of real options should provide a window into why they are so alluring to practitioners and businesses. In essence, we are ignoring that the time honored rules of capital budgeting, which include rejecting investments that have negative net present value, when real options are present. Not only does the real options approach encourage you to make investments that do not meet conventional financial criteria, it also makes it more likely that you will do so, the less you know about the investment. Ignorance, rather than being a weakness, becomes a virtue because it pushes up the uncertainty in the estimated value and the resulting option value. To prevent the real options process from being hijacked by managers who want to rationalize bad (and risky) decisions, we have to impose some reasonable constraints on when it can be used and when it is used, how to estimate its value.

First, not all investments have options embedded in them, and not all options, even if they do exist, have value. To assess whether an investment creates valuable options that need to be analyzed and valued, three key questions need to be answered affirmatively.

- Is the first investment a pre-requisite for the later investment/expansion? If not, how necessary is the first investment for the later investment/expansion? Consider our earlier analysis of the value of a patent or the value of an undeveloped oil reserve as options. A firm cannot generate patents without investing in research or paying

¹⁹ From its inception until recently, Southwest used the Boeing 737 as its workhorse, thus reducing its need

another firm for the patents, and it cannot get rights to an undeveloped oil reserve without bidding on it at a government auction or buying it from another oil company. Clearly, the initial investment here (spending on R&D, bidding at the auction) is required for the firm to have the second option. Now consider the Disney expansion into Mexico. The initial investment in a Spanish channel provides Disney with information about market potential, without which presumably it is unwilling to expand into the larger South American market. Unlike the patent and undeveloped reserves illustrations, the initial investment is not a pre-requisite for the second, though management might view it as such. The connection gets even weaker when we look at one firm acquiring another to have the option to be able to enter a large market. Acquiring an internet service provider to have a foothold in the internet retailing market or buying a Brazilian brewery to preserve the option to enter the Brazilian beer market would be examples of such transactions.

- Does the firm have an exclusive right to the later investment/expansion? If not, does the initial investment provide the firm with significant competitive advantages on subsequent investments? The value of the option ultimately derives not from the cash flows generated by then second and subsequent investments, but from the excess returns generated by these cash flows. The greater the potential for excess returns on the second investment, the greater the value of the option in the first investment. The potential for excess returns is closely tied to how much of a competitive advantage the first investment provides the firm when it takes subsequent investments. At one extreme, again, consider investing in research and development to acquire a patent. The patent gives the firm that owns it the exclusive rights to produce that product, and if the market potential is large, the right to the excess returns from the project. At the other extreme, the firm might get no competitive advantages on subsequent investments, in which case, it is questionable as to whether there can be any excess returns on these investments. In reality, most investments will fall in the continuum between these two extremes, with greater competitive advantages being associated with higher excess returns and larger option values.

- How sustainable are the competitive advantages? In a competitive market place, excess returns attract competitors, and competition drives out excess returns. The more sustainable the competitive advantages possessed by a firm, the greater will be the value of the options embedded in the initial investment. The sustainability of competitive advantages is a function of two forces. The first is the nature of the competition; other things remaining equal, competitive advantages fade much more quickly in sectors where there are aggressive competitors and new entry into the business is easy. The second is the nature of the competitive advantage. If the resource controlled by the firm is finite and scarce (as is the case with natural resource reserves and vacant land), the competitive advantage is likely to be sustainable for longer periods. Alternatively, if the competitive advantage comes from being the first mover in a market or technological expertise, it will come under assault far sooner. The most direct way of reflecting this in the value of the option is in its life; the life of the option can be set to the period of competitive advantage and only the excess returns earned over this period counts towards the value of the option.

Second, when real options are used to justify a decision, the justification has to be in more than qualitative terms. In other words, managers who argue for taking a project with poor returns or paying a premium on an acquisition on the basis of real options, should be required to value these real options and show, in fact, that the economic benefits exceed the costs. There will be two arguments made against this requirement. The first is that real options cannot be easily valued, since the inputs are difficult to obtain and often noisy. The second is that the inputs to option pricing models can be easily manipulated to back up whatever the conclusion might be. While both arguments have some basis, an estimate with error is better than no estimate at all, and the process of quantitatively trying to estimate the value of a real option is, in fact, the first step to understanding what drives its value.

There is one final note of caution that we should add about the use of option pricing models to assess the value of real options. Option pricing models, be they of the binomial or Black Scholes variety, are based on two fundamental precepts – replication and arbitrage. For either to be feasible, you have to be able to trade on the underlying asset and on the option. This is easy to accomplish with a listed option on a traded stock;

you can trade on both the stock and the listed option. It is much more difficult to pull off when valuing a patent or an investment expansion opportunity; neither the underlying asset (the product that emerges from the patent) nor the option itself are traded. This does not mean that you cannot estimate the value of a patent as an option but it does indicate that monetizing this value will be much more difficult to do. In the Avonex example from earlier in the chapter, the option value for the patent was \$907 million whereas the conventional risk adjusted value was only \$547 million. Much as you may believe in the former as the right estimate of value, it is unlikely that any potential buyer of the patent will come close to paying that amount.

Real Options in a Risk Management Framework

Given the different perspective on risk brought into the picture by real options, how do we fit this approach into the broader set of risk assessment tools and what role, if any, should it play in risk management? While there are some real options purists who view it as the answer to all of the problems that we face in managing risk, a more nuanced conclusion is merited.

Real options have made an important contribution to the risk management debate by bringing in the potential upside in risk to offset the hand wringing generated by the downside. It can also be viewed as a bridge between corporate finance and corporate strategy. Historically, the former has been focused on how best to assess the value of risky assets in the interests of maximizing firm value, and the latter on the sources of competitive advantages and market potential. The real option framework allows us to bring the rigors of financial analysis to corporate strategic analysis and link it up with value creation and maximization. Finally, the real options approach reveals the value of maintaining flexibility in both operating and financial decisions. By preserving the flexibility to both scale up an investment, in good scenarios, and to scale down or abandon the same investment, in down scenarios, a firm may be able to turn a bad investment into a good one.

As we noted earlier in the chapter, though, the value of real options is greatest when you have exclusivity and dissipates quickly in competitive environments. Consequently, real options will be most useful to firms that have significant competitive

advantages and can therefore assume that they will be able to act alone or at least much earlier than their competition in response to new information. It should come as no surprise that the real options approach has been used longest and with the most success, by mining and commodity companies. The danger with extending the real options framework to all firms is that it will inevitably be used to justify bad investments and decisions.

If you decide to use the real option approach in risk management, it should not replace risk adjusted values or Monte Carlo simulations but should be viewed more as a supplement or a complement to these approaches.²⁰ After all, to assess the value of Avonex, we began with the risk adjusted present value of the expected cash flows from the drug. Similarly, to analyze the Disney expansion opportunity in South America, we drew on the output from Monte Carlo simulations.

Conclusion

In contrast to the approaches that focus on downside risk – risk adjusted value, simulations and Value at Risk – the real options approach brings an optimistic view to uncertainty. While conceding that uncertainty can create losses, it argues that uncertainty can also be exploited for potential gains and that updated information can be used to augment the upside and reduce the downside risks inherent in investments. In essence, you are arguing that the conventional risk adjustment approaches fail to capture this flexibility and that you should be adding an option premium to the risk adjusted value.

In this chapter, we considered three potential real options and applications of each. The first is the option to delay, where a firm with exclusive rights to an investment has the option of deciding when to take that investment and to delay taking it, if necessary. The second is the option to expand, where a firm may be willing to lose money on an initial investment, in the hope of expanding into other investments or markets further down the road. The third is the option to abandon an investment, if it looks like a money loser, early in the process.

²⁰ For an example of how simulations and real options complement each other, see Gamba, A., 2002, Real Options Valuation: A Monte Carlo Approach, Working Paper, SSRN.


While it is clearly appropriate to attach value to real options in some cases – patents, reserves of natural resources or exclusive licenses – the argument for an option premium gets progressively weaker as we move away from the exclusivity inherent in each of these cases. In particular, a firm that invests into an emerging market in a money-losing enterprise, using the argument that that market is a large and potentially profitable one, could be making a serious mistake. After all, the firm could be right in its assessment of the market, but absent barriers to entry, it may not be able to earn excess returns in that market or keep the competition out. Not all opportunities are options and not all options have significant economic value.

Appendix: Basics of Options and Option Pricing

An option provides the holder with the right to buy or sell a specified quantity of an underlying asset at a fixed price (called a strike price or an exercise price) at or before the expiration date of the option. Since it is a right and not an obligation, the holder can choose not to exercise the right and allow the option to expire. There are two types of options - call options and put options.


Option Payoffs

A call option gives the buyer of the option the right to buy the underlying asset at a fixed price, called the strike or the exercise price, at any time prior to the expiration date of the option: the buyer pays a price for this right. If at expiration, the value of the asset is less than the strike price, the option is not exercised and expires worthless. If, on the other hand, the value of the asset is greater than the strike price, the option is exercised - the buyer of the option buys the stock at the exercise price and the difference between the asset value and the exercise price comprises the gross profit on the investment. The net profit on the investment is the difference between the gross profit and the price paid for the call initially. A payoff diagram illustrates the cash payoff on an option at expiration. For a call, the net payoff is negative (and equal to the price paid for the call) if the value of the underlying asset is less than the strike price. If the price of the underlying asset exceeds the strike price, the gross payoff is the difference between the value of the underlying asset and the strike price, and the net payoff is the difference between the gross payoff and the price of the call. This is illustrated in the figure 8A.1:

Figure 8A.1: Payoff on Call Option

A put option gives the buyer of the option the right to sell the underlying asset at a fixed price, again called the strike or exercise price, at any time prior to the expiration date of the option. The buyer pays a price for this right. If the price of the underlying asset is greater than the strike price, the option will not be exercised and will expire worthless. If on the other hand, the price of the underlying asset is less than the strike price, the owner of the put option will exercise the option and sell the stock at the strike price, claiming the difference between the strike price and the market value of the asset as the gross profit. Again, netting out the initial cost paid for the put yields the net profit from the transaction. A put has a negative net payoff if the value of the underlying asset exceeds the strike price, and has a gross payoff equal to the difference between the strike price and the value of the underlying asset if the asset value is less than the strike price. This is summarized in figure 8A.2.

Figure 8A.2: Payoff on Put Option



There is one final distinction that needs to be made. Options are usually categorized as American or European options. A primary distinction between two is that American options can be exercised at any time prior to its expiration, while European options can be exercised only at expiration. The possibility of early exercise makes American options more valuable than otherwise similar European options; it also makes them more difficult to value. There is one compensating factor that enables the former to be valued using models designed for the latter. In most cases, the time premium associated with the remaining life of an option and transactions costs makes early exercise sub-optimal. In other words, the holders of in-the-money options will generally get much more by selling the option to someone else than by exercising the options.²¹

Determinants of Option Value

The value of an option is determined by a number of variables relating to the underlying asset and financial markets.

²¹ While early exercise is not optimal generally, there are at least two exceptions to this rule. One is a case where the underlying asset pays large dividends, thus reducing the value of the asset, and any call options on that asset. In this case, call options may be exercised just before an ex-dividend date, if the time premium on the options is less than the expected decline in asset value as a consequence of the dividend payment. The other exception arises when an investor holds both the underlying asset and deep in-the-money puts on that asset at a time when interest rates are high. In this case, the time premium on the put may be less than the potential gain from exercising the put early and earning interest on the exercise price.

1. Current Value of the Underlying Asset : Options are assets that derive value from an underlying asset. Consequently, changes in the value of the underlying asset affect the value of the options on that asset. Since calls provide the right to buy the underlying asset at a fixed price, an increase in the value of the asset will increase the value of the calls. Puts, on the other hand, become less valuable as the value of the asset increase.

2. Variance in Value of the Underlying Asset: The buyer of an option acquires the right to buy or sell the underlying asset at a fixed price. The higher the variance in the value of the underlying asset, the greater the value of the option. This is true for both calls and puts. While it may seem counter-intuitive that an increase in a risk measure (variance) should increase value, options are different from other securities since buyers of options can never lose more than the price they pay for them; in fact, they have the potential to earn significant returns from large price movements.

3. Dividends Paid on the Underlying Asset: The value of the underlying asset can be expected to decrease if dividend payments are made on the asset during the life of the option. Consequently, the value of a call on the asset is a decreasing function of the size of expected dividend payments, and the value of a put is an increasing function of expected dividend payments. A more intuitive way of thinking about dividend payments, for call options, is as a cost of delaying exercise on in-the-money options. To see why, consider a option on a traded stock. Once a call option is in the money, i.e, the holder of the option will make a gross payoff by exercising the option, exercising the call option will provide the holder with the stock, and entitle him or her to the dividends on the stock in subsequent periods. Failing to exercise the option will mean that these dividends are foregone.

4. Strike Price of Option: A key characteristic used to describe an option is the strike price. In the case of calls, where the holder acquires the right to buy at a fixed price, the value of the call will decline as the strike price increases. In the case of puts, where the holder has the right to sell at a fixed price, the value will increase as the strike price increases.

5. Time To Expiration On Option: Both calls and puts become more valuable as the time to expiration increases. This is because the longer time to expiration provides more time for the value of the underlying asset to move, increasing the value of both types of

options. Additionally, in the case of a call, where the buyer has to pay a fixed price at expiration, the present value of this fixed price decreases as the life of the option increases, increasing the value of the call.

6. Riskless Interest Rate Corresponding To Life Of Option: Since the buyer of an option pays the price of the option up front, an opportunity cost is involved. This cost will depend upon the level of interest rates and the time to expiration on the option. The riskless interest rate also enters into the valuation of options when the present value of the exercise price is calculated, since the exercise price does not have to be paid (received) until expiration on calls (puts). Increases in the interest rate will increase the value of calls and reduce the value of puts.

Table 8A.1 below summarizes the variables and their predicted effects on call and put prices.

Table 8A.1: Summary of Variables Affecting Call and Put Prices

<i>Factor</i>	<i>Effect on</i>	
	<i>Call Value</i>	<i>Put Value</i>
Increase in underlying asset's value	Increases	Decreases
Increase in Strike Price	Decreases	Increases
Increase in variance of underlying asset	Increases	Increases
Increase in time to expiration	Increases	Increases
Increase in interest rates	Increases	Decreases
Increase in dividends paid	Decreases	Increases


Option Pricing Models

Option pricing theory has made vast strides since 1972, when Black and Scholes published their path-breaking paper providing a model for valuing dividend-protected European options. Black and Scholes used a “replicating portfolio” — a portfolio composed of the underlying asset and the risk-free asset that had the same cash flows as the option being valued— to come up with their final formulation. While their derivation is mathematically complicated, there is a simpler binomial model for valuing options that draws on the same logic.

The Binomial Model

The binomial option pricing model is based upon a simple formulation for the asset price process, in which the asset, in any time period, can move to one of two possible prices. The general formulation of a stock price process that follows the binomial is shown in figure 8A.3.

Figure 8A.3: General Formulation for Binomial Price Path



In this figure, S is the current stock price; the price moves up to S_u with probability p and down to S_d with probability $1-p$ in any time period.

The objective in creating a replicating portfolio is to use a combination of risk-free borrowing/lending and the underlying asset to create the same cash flows as the option being valued. The principles of arbitrage apply here, and the value of the option must be equal to the value of the replicating portfolio. In the case of the general formulation above, where stock prices can either move up to S_u or down to S_d in any time period, the replicating portfolio for a call with strike price K will involve borrowing $\$B$ and acquiring Δ of the underlying asset, where:

$$\Delta = \text{Number of units of the underlying asset bought} = (C_u - C_d)/(S_u - S_d)$$

where,


C_u = Value of the call if the stock price is S_u

C_d = Value of the call if the stock price is S_d

In a multi-period binomial process, the valuation has to proceed iteratively; i.e., starting with the last time period and moving backwards in time until the current point in time. The portfolios replicating the option are created at each step and valued, providing the values for the option in that time period. The final output from the binomial option pricing model is a statement of the value of the option in terms of the replicating portfolio, composed of Δ shares (option delta) of the underlying asset and risk-free borrowing/lending.

Value of the call = Current value of underlying asset * Option Delta - Borrowing needed to replicate the option

Consider a simple example. Assume that the objective is to value a call with a strike price of 50, which is expected to expire in two time periods, on an underlying asset whose price currently is 50 and is expected to follow a binomial process:



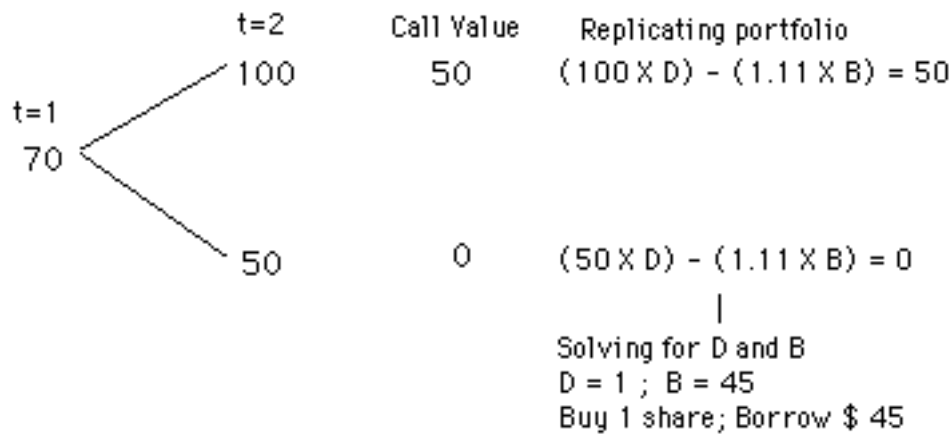
Now assume that the interest rate is 11%. In addition, define

Δ = Number of shares in the replicating portfolio

B = Dollars of borrowing in replicating portfolio

The objective is to combine Δ shares of stock and B dollars of borrowing to replicate the cash flows from the call with a strike price of \$ 50. This can be done iteratively, starting with the last period and working back through the binomial tree.


Step 1: Start with the end nodes and work backwards:



Thus, if the stock price is \$70 at t=1, borrowing \$45 and buying one share of the stock will give the same cash flows as buying the call. The value of the call at t=1, if the stock price is \$70, is therefore:

$$\text{Value of Call} = \text{Value of Replicating Position} = 70 \Delta - B = 70 - 45 = 25$$

Considering the other leg of the binomial tree at t=1,



If the stock price is 35 at t=1, then the call is worth nothing.

Step 2: Move backwards to the earlier time period and create a replicating portfolio that will provide the cash flows the option will provide.

price process is a continuous one. If as t approaches 0, price changes remain large, the limiting distribution is the Poisson distribution, i.e., a distribution that allows for price jumps. The Black-Scholes model applies when the limiting distribution is the normal distribution,²² and it explicitly assumes that the price process is continuous.

The Model

The original Black and Scholes model was designed to value European options, which were dividend-protected. Thus, neither the possibility of early exercise nor the payment of dividends affects the value of options in this model. The value of a call option in the Black-Scholes model can be written as a function of the following variables:

S = Current value of the underlying asset

K = Strike price of the option

t = Life to expiration of the option

r = Riskless interest rate corresponding to the life of the option

σ^2 = Variance in the $\ln(\text{value})$ of the underlying asset

The model itself can be written as:

$$\text{Value of call} = S N(d_1) - K e^{-rt} N(d_2)$$

where

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right) t}{\sigma \sqrt{t}}$$

$$d_2 = d_1 - \sigma \sqrt{t}$$

The process of valuation of options using the Black-Scholes model involves the following steps:

Step 1: The inputs to the Black-Scholes are used to estimate d_1 and d_2 .

Step 2: The cumulative normal distribution functions, $N(d_1)$ and $N(d_2)$, corresponding to these standardized normal variables are estimated.

²² Stock prices cannot drop below zero, because of the limited liability of stockholders in publicly listed firms. Hence, stock prices, by themselves, cannot be normally distributed, since a normal distribution requires some probability of infinitely negative values. The distribution of the natural logs of stock prices is assumed to be log-normal in the Black-Scholes model. This is why the variance used in this model is the variance in the log of stock prices.

Step 3: The present value of the exercise price is estimated, using the continuous time version of the present value formulation:

$$\text{Present value of exercise price} = K e^{-rt}$$

Step 4: The value of the call is estimated from the Black-Scholes model.

The determinants of value in the Black-Scholes are the same as those in the binomial - the current value of the stock price, the variability in stock prices, the time to expiration on the option, the strike price, and the riskless interest rate. The principle of replicating portfolios that is used in binomial valuation also underlies the Black-Scholes model. In fact, embedded in the Black-Scholes model is the replicating portfolio.

$$\text{Value of call} = \underbrace{S N(d_1)}_{\text{Buy } \overline{N(d_1)} \text{ shares}} - \underbrace{K e^{-rt} N(d_2)}_{\text{Borrow this amount}}$$

$N(d_1)$, which is the number of shares that are needed to create the replicating portfolio is called the option delta. This replicating portfolio is self-financing and has the same value as the call at every stage of the option's life.

Model Limitations and Fixes

The version of the Black-Scholes model presented above does not take into account the possibility of early exercise or the payment of dividends, both of which impact the value of options. Adjustments exist, which while not perfect, provide partial corrections to value.

1. Dividends

The payment of dividends reduces the stock price. Consequently, call options will become less valuable and put options more valuable as dividend payments increase. One approach to dealing with dividends to estimate the present value of expected dividends paid by the underlying asset during the option life and subtract it from the current value of the asset to use as “S” in the model. Since this becomes impractical as the option life becomes longer, we would suggest an alternate approach. If the dividend yield ($y = \text{dividends} / \text{current value of the asset}$) of the underlying asset is expected to remain unchanged during the life of the option, the Black-Scholes model can be modified to take dividends into account.

$$C = S e^{-yt} N(d_1) - K e^{-rt} N(d_2)$$

where

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + (r - y + \frac{\sigma^2}{2})t}{\sigma\sqrt{t}}$$

$$d_2 = d_1 - \sigma\sqrt{t}$$

From an intuitive standpoint, the adjustments have two effects. First, the value of the asset is discounted back to the present at the dividend yield to take into account the expected drop in value from dividend payments. Second, the interest rate is offset by the dividend yield to reflect the lower carrying cost from holding the stock (in the replicating portfolio). The net effect will be a reduction in the value of calls, with the adjustment, and an increase in the value of puts.

2. *Early Exercise*

The Black-Scholes model is designed to value European options, whereas most options that we consider are American options, which can be exercised anytime before expiration. Without working through the mechanics of valuation models, an American option should always be worth at least as much and generally more than a European option because of the early exercise option. There are three basic approaches for dealing with the possibility of early exercise. The first is to continue to use the unadjusted Black-Scholes, and regard the resulting value as a floor or conservative estimate of the true value. The second approach is to value the option to each potential exercise date. With options on stocks, this basically requires that we value options to each ex-dividend day and chooses the maximum of the estimated call values. The third approach is to use a modified version of the binomial model to consider the possibility of early exercise.

While it is difficult to estimate the prices for each node of a binomial, there is a way in which variances estimated from historical data can be used to compute the expected up and down movements in the binomial. To illustrate, if σ^2 is the variance in $\ln(\text{stock prices})$, the up and down movements in the binomial can be estimated as follows:

$$u = \text{Exp} [(r - \sigma^2/2)(T/m) + \sqrt{(\sigma^2 T/m)}]$$

$$d = \text{Exp} [(r - \sigma^2/2)(T/m) - \sqrt{(\sigma^2 T/m)}]$$

where u and d are the up and down movements per unit time for the binomial, T is the life of the option and m is the number of periods within that lifetime. Multiplying the stock price at each stage by u and d will yield the up and the down prices. These can then be used to value the asset.

3. *The Impact Of Exercise On The Value Of The Underlying Asset*

The derivation of the Black-Scholes model is based upon the assumption that exercising an option does not affect the value of the underlying asset. This may be true for listed options on stocks, but it is not true for some types of options. For instance, the exercise of warrants increases the number of shares outstanding and brings fresh cash into the firm, both of which will affect the stock price.²³ The expected negative impact (dilution) of exercise will decrease the value of warrants compared to otherwise similar call options. The adjustment for dilution in the Black-Scholes to the stock price is fairly simple. The stock price is adjusted for the expected dilution from the exercise of the options. In the case of warrants, for instance:

$$\text{Dilution-adjusted } S = (S n_s + W n_w) / (n_s + n_w)$$

where

S = Current value of the stock
 n_w = Number of warrants outstanding

W = Market value of warrants outstanding
 n_s = Number of shares outstanding

When the warrants are exercised, the number of shares outstanding will increase, reducing the stock price. The numerator reflects the market value of equity, including both stocks and warrants outstanding. The reduction in S will reduce the value of the call option.

There is an element of circularity in this analysis, since the value of the warrant is needed to estimate the dilution-adjusted S and the dilution-adjusted S is needed to estimate the value of the warrant. This problem can be resolved by starting the process off with an estimated value of the warrant (say, the exercise value), and then iterating with the new estimated value for the warrant until there is convergence.

²³ Warrants are call options issued by firms, either as part of management compensation contracts or to raise equity.

Valuing Puts

The value of a put can be derived from the value of a call with the same strike price and the same expiration date through an arbitrage relationship that specifies that:

$$C - P = S - K e^{-rt}$$

where C is the value of the call and P is the value of the put (with the same life and exercise price).

This arbitrage relationship can be derived fairly easily and is called put-call parity. To see why put-call parity holds, consider creating the following portfolio:

- (a) Sell a call and buy a put with exercise price K and the same expiration date "t"
- (b) Buy the stock at current stock price S

The payoff from this position is riskless and always yields K at expiration (t). To see this, assume that the stock price at expiration is S*:

<i>Position</i>	<i>Payoffs at t if S* > K</i>	<i>Payoffs at t if S* < K</i>
Sell call	-(S*-K)	0
Buy put	0	K-S*
Buy stock	S*	S*
<i>Total</i>	<i>K</i>	<i>K</i>

Since this position yields K with certainty, its value must be equal to the present value of K at the riskless rate ($K e^{-rt}$).

$$S + P - C = K e^{-rt}$$

$$C - P = S - K e^{-rt}$$

This relationship can be used to value puts. Substituting the Black-Scholes formulation for the value of an equivalent call,

$$\text{Value of put} = S e^{-yt} (N(d_1) - 1) - K e^{-rt} (N(d_2) - 1)$$

where

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + \left(r - y + \frac{\sigma^2}{2}\right) t}{\sigma \sqrt{t}}$$

$$d_2 = d_1 - \sigma \sqrt{t}$$
