

Real Options

► **When you use** discounted cash flow (DCF) to value a project, you implicitly assume that your firm will hold the project passively. In other words, you are ignoring the *real options* attached to the project—options that sophisticated managers can take advantage of. You could say that DCF does not reflect the value of management. Managers who hold real options do not have to be passive; they can make decisions to capitalize on good fortune or to mitigate loss. The opportunity to make such decisions clearly adds value whenever project outcomes are uncertain.

Chapter 10 introduced the four main types of real options:

- The option to expand if the immediate investment project succeeds.
- The option to wait (and learn) before investing.
- The option to shrink or abandon a project.
- The option to vary the mix of output or the firm's production methods.

Chapter 10 gave several simple examples of real options. We also showed you how to use decision trees to set out possible future outcomes and decisions. But we did not show you how to value real options. That is our task in this chapter. We apply the concepts and valuation principles you learned in Chapter 21.

For the most part we work with simple numerical examples. The art and science of valuing real options are illustrated just as well with simple calculations as complex ones. But we also describe several more realistic examples, including:

- A strategic investment in the computer business.
- The valuation of an aircraft purchase option.
- The option to develop commercial real estate.
- The decision to operate or mothball an oil tanker.

These examples show how financial managers can value real options in real life. We also show how managers can *create* real options, adding value by adding flexibility to the firm's investments and operations.

We should start with a warning. Setting out the possible future choices that the firm may encounter usually calls for a strong dose of judgment. Therefore, do not expect precision when valuing real options. Often managers do not even try to put a figure on the value of the option, but simply draw on their experience to decide whether it is worth paying for additional flexibility. Thus they might say, "We just don't know whether gargle blasters will catch on, but it probably makes sense to spend an extra \$200,000 now to allow for an extra production line in the future."

22-1 The Value of Follow-on Investment Opportunities

It is 1982. You are assistant to the chief financial officer (CFO) of Blitzen Computers, an established computer manufacturer casting a profit-hungry eye on the rapidly developing personal computer market. You are helping the CFO evaluate the proposed introduction of the Blitzen Mark I Micro.

The Mark I's forecasted cash flows and NPV are shown in Table 22.1. Unfortunately the Mark I can't meet Blitzen's customary 20% hurdle rate and has a \$46 million negative NPV, contrary to top management's strong gut feeling that Blitzen ought to be in the personal computer market.

The CFO has called you in to discuss the project:

"The Mark I just can't make it on financial grounds," the CFO says. "But we've got to do it for strategic reasons. I'm recommending we go ahead."

"But you're missing the all-important financial advantage, Chief," you reply.

"Don't call me 'Chief.' What financial advantage?"

"If we don't launch the Mark I, it will probably be too expensive to enter the micro market later, when Apple, IBM, and others are firmly established. If we go ahead, we have the opportunity to make follow-on investments that could be extremely profitable. The Mark I gives not only its own cash flows but also a call option to go on with a Mark II micro. That call option is the real source of strategic value."

"So it's strategic value by another name. That doesn't tell me what the Mark II investment's worth. The Mark II could be a great investment or a lousy one—we haven't got a clue."

"That's exactly when a call option is worth the most," you point out perceptively. "The call lets us invest in the Mark II if it's great and walk away from it if it's lousy."

"So what's it worth?"

"Hard to say precisely, but I've done a back-of-the-envelope calculation, which suggests that the value of the option to invest in the Mark II could more than offset the Mark I's \$46 million negative NPV. [The calculations are shown in Table 22.2.] If the option to invest is worth \$55 million, the total value of the Mark I is its own NPV, -\$46 million, plus the \$55 million option attached to it, or +\$9 million."

"You're just overestimating the Mark II," the CFO says gruffly. "It's easy to be optimistic when an investment is three years away."

"No, no," you reply patiently. "The Mark II is expected to be no more profitable than the Mark I—just twice as big and therefore twice as bad in terms of discounted cash flow. I'm forecasting it to have a negative NPV of about \$100 million. But there's a chance the Mark II could be extremely valuable. The call option allows Blitzen to cash in on those upside outcomes. The chance to cash in could be worth \$55 million."

"Of course, the \$55 million is only a trial calculation, but it illustrates how valuable follow-on investment opportunities can be, especially when uncertainty is high and the product market is growing rapidly. Moreover, the Mark II will give us a call on the Mark III, the Mark III on the Mark IV, and so on. My calculations don't take subsequent calls into account."

"I think I'm beginning to understand a little bit of corporate strategy," mumbles the CFO.

	Year					
	1982	1983	1984	1985	1986	1987
After-tax operating cash flow (1)		+110	+159	+295	+185	0
Capital investment (2)	450	0	0	0	0	0
Increase in working capital (3)	0	50	100	100	-125	-125
Net cash flow (1) - (2) - (3)	-450	+60	+59	+195	+310	+125
NPV at 20% = - \$46.45, or about - \$46 million						

TABLE 22.1

Summary of cash flows and financial analysis of the Mark I microcomputer (\$ millions).

TABLE 22.2

Valuing the option to invest in the Mark II microcomputer.

Assumptions	
1.	The decision to invest in the Mark II must be made after three years, in 1985.
2.	The Mark II investment is double the scale of the Mark I (note the expected rapid growth of the industry). Investment required is \$900 million (the exercise price), which is taken as fixed.
3.	Forecasted cash inflows of the Mark II are also double those of the Mark I, with present value of \$807 million in 1985 and $807/(1.2)^3 = \$467$ million in 1982.
4.	The future value of the Mark II cash flows is highly uncertain. This value evolves as a stock price does with a standard deviation of 35% per year. (Many high-technology stocks have standard deviations higher than 35%.)
5.	The annual interest rate is 10%.
Interpretation	
The opportunity to invest in the Mark II is a three-year call option on an asset worth \$467 million with a \$900 million exercise price.	
Valuation	
$PV(\text{exercise price}) = \frac{900}{(1.1)^3} = 676$ $\text{Call value} = [N(d_1) \times P] - [N(d_2) \times PV(\text{EX})]$ $d_1 = \frac{\log[P/PV(\text{EX})]/\sigma\sqrt{t} + \sigma\sqrt{t}/2}{\sigma\sqrt{t}}$ $= \frac{\log[.691]/.606 + .606/2}{.606} = -.3072$ $d_2 = d_1 - \sigma\sqrt{t} = -.3072 - .606 = -.9134$ $N(d_1) = .3793, N(d_2) = .1805$ $\text{Call value} = [.3793 \times 467] - [.1805 \times 676] = \55.1 million	

Questions and Answers about Blitzen's Mark II

Question: I know how to use the Black-Scholes formula to value traded call options, but this case seems harder. What number do I use for the stock price? I don't see any traded shares.

Answer: With traded call options, you can see the value of the *underlying asset* that the call is written on. Here the option is to buy a nontraded real asset, the Mark II. We can't observe the Mark II's value; we have to compute it.

The Mark II's forecasted cash flows are set out in Table 22.3. The project involves an initial outlay of \$900 million in 1985. The cash inflows start in the following year and have a present value of \$807 million in 1985, equivalent to \$467 million in 1982 as shown in Table 22.3. So the real option to invest in the Mark II amounts to a three-year call on an underlying asset worth \$467 million, with a \$900 million exercise price.

TABLE 22.3

Cash flows of the Mark II microcomputer, as forecasted from 1982 (\$ millions).

	Year						
	1982	1985	1986	1987	1988	1989	1990
After-tax operating cash flow			+220	+318	+590	+370	0
Increase in working capital			100	200	200	-250	-250
Net cash flow			+120	+118	+390	+620	+250
Present value at 20%	+467	←	+807				
Investment, PV at 10% (PV in 1982)	676	←	900				
Forecasted NPV in 1985			-	93			

Notice that real options analysis does *not* replace DCF. You typically need DCF to value the underlying asset.

Question: Table 22.2 uses a standard deviation of 35% per year. Where does that number come from?

Answer: We recommend you look for *comparables*, that is, traded stocks with business risks similar to the investment opportunity.¹ For the Mark II, the ideal comparables would be growth stocks in the personal computer business, or perhaps a broader sample of high-tech growth stocks. Use the average standard deviation of the comparable companies' returns as the benchmark for judging the risk of the investment opportunity.²

Question: Table 22.3 discounts the Mark II's cash flows at 20%. I understand the high discount rate, because the Mark II is risky. But why is the \$900 million investment discounted at the risk-free interest rate of 10%? Table 22.3 shows the present value of the investment in 1982 of \$676 million.

Answer: Black and Scholes assumed that the exercise price is a fixed, certain amount. We wanted to stick with their basic formula. If the exercise price is uncertain, you can switch to a slightly more complicated valuation formula.³

Question: Nevertheless, if I had to decide in 1982, once and for all, whether to invest in the Mark II, I wouldn't do it. Right?

Answer: Right. The NPV of a commitment to invest in the Mark II is negative:

$$\text{NPV}(1982) = \text{PV}(\text{cash inflows}) - \text{PV}(\text{investment}) = \$467 - 676 = -\$209 \text{ million}$$

The option to invest in the Mark II is "out of the money" because the Mark II's value is far less than the required investment. Nevertheless, the option is worth +\$55 million. It is especially valuable because the Mark II is a risky project with lots of upside potential. Figure 22.1 shows the probability distribution of the possible present values of the Mark II in 1985. The expected (mean or average) outcome is our forecast of \$807,⁴ but the actual value could exceed \$2 billion.

Question: Could it also be far below \$807 million—\$500 million or less?

Answer: The downside is irrelevant, because Blitzen won't invest unless the Mark II's actual value turns out higher than \$900 million. The net option payoffs for all values less than \$900 million are zero.

In a DCF analysis, you discount the expected outcome (\$807 million), which averages the downside against the upside, the bad outcomes against the good. The value of a call option depends only on the upside. You can see the danger of trying to value a future investment option with DCF.

¹ You could also use scenario analysis, which we described in Chapter 10. Work out "best" and "worst" scenarios to establish a range of possible future values. Then find the annual standard deviation that would generate this range over the life of the option. For the Mark II, a range from \$300 million to \$2 billion would cover about 90% of the possible outcomes. This range, shown in Figure 22.1, is consistent with an annual standard deviation of 35%.

² Be sure to "unlever" the standard deviations, thereby eliminating volatility created by debt financing. Chapter 17 covered unlevering procedures for beta. The same principles apply for standard deviation: You want the standard deviation of a portfolio of all the debt and equity securities issued by the comparable firm.

³ If the required investment is uncertain, you have, in effect, an option to exchange one risky asset (the future value of the exercise price) for another (the future value of the Mark II's cash inflows). See W. Margrabe, "The Value of an Option to Exchange One Asset for Another," *Journal of Finance* 33 (March 1978), pp. 177–186.

⁴ We have drawn the future values of the Mark II as a lognormal distribution, consistent with the assumptions of the Black–Scholes formula. Lognormal distributions are skewed to the right, so the average outcome is greater than the most likely outcome. The most likely outcome is the highest point on the probability distribution.

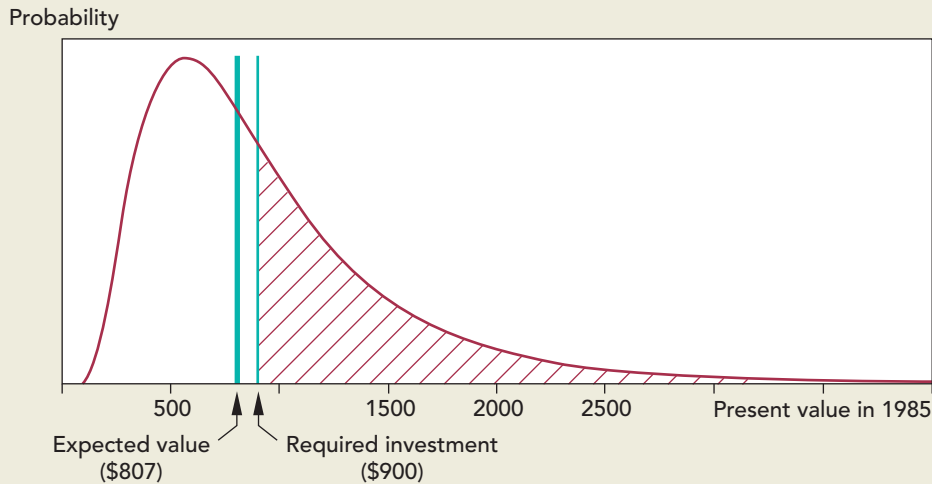


FIGURE 22.1

This distribution shows the range of possible present values for the Mark II project in 1985. The expected value is about \$800 million, less than the required investment of \$900 million. The option to invest pays off in the shaded area above \$900 million.

Question: What's the decision rule?

Answer: Adjusted present value. The Mark I project costs \$46 million ($NPV = -\46 million), but accepting it creates the expansion option for the Mark II. The expansion option is worth \$55 million, so:

$$APV = -46 + 55 = +\$9 \text{ million}$$

Of course we haven't counted other follow-on opportunities. If the Mark I and Mark II are successes, there will be an option to invest in the Mark III, possibly the Mark IV, and so on.

Other Expansion Options

You can probably think of many other cases where companies spend money today to create opportunities to expand in the future. A mining company may acquire rights to an ore body that is not worth developing today but could be very profitable if product prices increase. A real estate developer may invest in worn-out farmland that could be turned into a shopping mall if a new highway is built. A pharmaceutical company may acquire a patent that gives the right but not the obligation to market a new drug. In each case the company is acquiring a real option to expand.

22-2 The Timing Option

The fact that a project has a positive NPV does not mean that you should go ahead today. It may be better to wait and see how the market develops.

Suppose that you are contemplating a now-or-never opportunity to build a malted herring factory. In this case you have an about-to-expire call option on the present value of the factory's future cash flows. If the present value exceeds the cost of the factory, the call option's payoff is the project's NPV. But if NPV is negative, the call option's payoff is zero, because in that case the firm will not make the investment.

Now suppose that you can delay construction of the plant. You still have the call option, but you face a trade-off. If the outlook is highly uncertain, it is tempting to wait and see whether the malted herring market takes off or decays. On the other hand, if the project is truly profitable, the sooner you can capture the project's cash flows, the better. If the cash flows are high enough, you will want to exercise your option right away.

The cash flows from an investment project play the same role as dividend payments on a stock. When a stock pays no dividends, an American call is always worth more alive than dead and should never be exercised early. But payment of a dividend before the option matures reduces the ex-dividend price and the possible payoffs to the call option at maturity. Think of the extreme case: If a company pays out all its assets in one bumper dividend, the stock price must be zero and the call worthless. Therefore, any in-the-money call would be exercised just before this liquidating dividend.

Dividends do not always prompt early exercise, but if they are sufficiently large, call option holders capture them by exercising just before the ex-dividend date. We see managers acting in the same way: When a project's forecasted cash flows are sufficiently large, managers capture the cash flows by investing right away. But when forecasted cash flows are small, managers are inclined to hold on to their call rather than to invest, even when project NPV is positive.⁵ This explains why managers are sometimes reluctant to commit to positive-NPV projects. This caution is rational as long as the option to wait is open and sufficiently valuable.

Valuing the Malted Herring Option

Figure 22.2 shows the possible cash flows and end-of-year values for the malted herring project. If you commit and invest \$180 million, you have a project worth \$200 million. If demand turns out to be low in year 1, the cash flow is only \$16 million and the value of the project falls to \$160 million. But if demand is high in year 1, the cash flow is \$25 million and value rises to \$250 million. Although the project lasts indefinitely, we assume that investment cannot be postponed beyond the end of the first year, and therefore we show only the cash flows for the first year and the possible values at the end of the year. Notice that if you undertake the investment right away, you capture the first year's cash flow (\$16 million or \$25 million); if you delay, you miss out on this cash flow, but you will have more information on how the project is likely to work out.

We can use the binomial method to value this option. The first step is to pretend that investors are risk neutral and to calculate the probabilities of high and low demand in this risk-neutral world. If demand is high in the first year, the malted herring plant

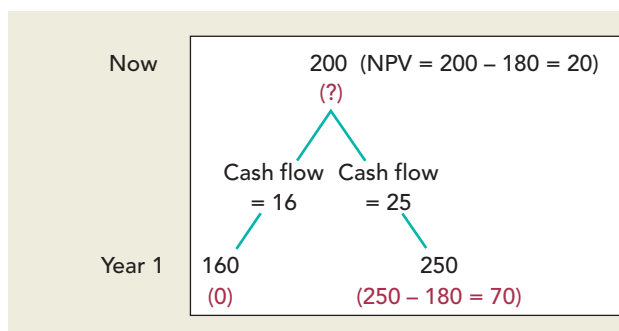


FIGURE 22.2

Possible cash flows and end-of-period values for the malted herring project are shown in black. The project costs \$180 million, either now or later. The red figures in parentheses show payoffs from the option to wait and to invest later if the project is positive-NPV at year 1. Waiting means loss of the first year's cash flows. The problem is to figure out the current value of the option.

⁵ We have been a bit vague about forecasted project cash flows. If competitors can enter and take away cash that you could have earned, the meaning is clear. But what about the decision to, say, develop an oil well? Here delay doesn't waste barrels of oil in the ground; it simply postpones production and the associated cash flow. The cost of waiting is the decline in today's *present value* of revenues from production. Present value declines if the cash flow from production increases more slowly than the cost of capital.

has a cash flow of \$25 million and a year-end value of \$250 million. The total return is $(25 + 250)/200 - 1 = .375$, or 37.5%. If demand is low, the plant has a cash flow of \$16 million and a year-end value of \$160 million. Total return is $(16 + 160)/200 - 1 = -.12$, or -12%. In a *risk-neutral* world, the expected return would be equal to the interest rate, which we assume is 5%:

$$\text{Expected return} = \left(\begin{array}{c} \text{Probability of} \\ \text{high demand} \end{array} \right) \times 37.5 + \left(\begin{array}{c} 1 - \text{probability of} \\ \text{high demand} \end{array} \right) \times (-12) = 5\%$$

Therefore the (pretend) probability of high demand is 34.3%.

We want to value a call option on the malted herring project with an exercise price of \$180 million. We begin as usual at the end and work backward. The bottom row of Figure 22.2 shows the possible values of this option at the end of the year. If project value is \$160 million, the option to invest is worthless. At the other extreme, if project value is \$250 million, option value is $250 - 180 = \$70$ million.

To calculate the value of the option today, we work out the expected payoffs in a risk-neutral world and discount at the interest rate of 5%. Thus, the value of your option to invest in the malted herring plant is:

$$\frac{(.343 \times 70) + (.657 \times 0)}{1.05} = \$22.9 \text{ million}$$

But here is where we need to recognize the opportunity to exercise the option immediately. The option is worth \$22.9 million if you keep it open, and it is worth the project's immediate NPV ($200 - 180 = \$20$ million) if exercised now. Therefore we decide to wait, and then to invest next year only if demand turns out high.

We have of course simplified the malted herring calculations. You won't find many actual investment-timing problems that fit into a one-step binomial tree. (We work through a more realistic, 8-step tree in the next section.) But the example delivers an important practical point: A positive NPV is not a sufficient reason for investing. It may be better to wait and see.

Optimal Timing for Real Estate Development

Sometimes it pays to wait for a long time, even for projects with large positive NPVs. Suppose you own a plot of vacant land in the suburbs.⁶ The land can be used for a hotel or an office building, but not for both. A hotel could be later converted to an office building, or an office building to a hotel, but only at significant cost. You are therefore reluctant to invest, even if both investments have positive NPVs.

In this case you have two options to invest, but only one can be exercised. You therefore learn two things by waiting. First, you learn about the general *level* of cash flows from development, for example, by observing changes in the value of developed properties near your land. Second, you can update your estimates of the *relative* size of the hotel's future cash flows versus the office building's.

Figure 22.3 shows the conditions in which you would finally commit to build either the hotel or the office building. The horizontal axis shows the current cash flows that a hotel would generate. The vertical axis shows current cash flows for an office building. For simplicity, we assume that each investment would have an NPV of exactly zero at a current cash flow of 100. Thus, if you were forced to invest today, you would choose the building with the higher cash flow, assuming the cash flow is greater than 100. (What if you were forced to decide today and each building could generate the same cash flow, say, 150? You would flip a coin.)

⁶ The following example is based on P. D. Childs, T. J. Riddiough, and A. J. Triantis, "Mixed Uses and the Redevelopment Option," *Real Estate Economics* 24 (Fall 1996), pp. 317-339.




FIGURE 22.3

Development option for vacant land, assuming two mutually exclusive uses, either hotel or office building. The developer should “wait and see” unless the hotel’s and office building’s cash flows end up in one of the shaded areas.

Source: Adapted from Figure 1 in P. D. Childs, T. J. Riddiough, and A. J. Triantis, “Mixed Uses and the Redevelopment Option,” *Real Estate Economics* 24 (Fall 1996), pp. 317–339. © 1996 Blackwell Publishers.

If the two buildings’ cash flows plot in the colored area at the lower right of Figure 22.3, you build the hotel. To fall in this area, the hotel’s cash flows have to beat two hurdles. First, they must exceed a minimum level of about 240. Second, they must exceed the office building’s cash flows by a sufficient amount. If the situation is reversed, with office building cash flows above the minimum level of 240, and also sufficiently above the hotel’s, then you build the office building. In this case, the cash flows plot in the colored area at the top left of the figure.

Notice how the “Wait and see” region extends upward along the 45-degree line in Figure 22.3. When the cash flows from the hotel and office building are nearly the same, you become *very* cautious before choosing one over the other.

You may be surprised at how high cash flows have to be in Figure 22.3 to justify investment. There are three reasons. First, building the office building means not building the hotel, and vice versa. Second, the calculations underlying Figure 22.3 assumed cash flows that were small, but growing; therefore, the costs of waiting to invest were small. Third, the calculations did not consider the threat that someone might build a competing hotel or office building right next door. In that case the “relax and wait” area of Figure 22.3 would shrink dramatically.

22-3 The Abandonment Option

Expansion value is important. When investments turn out well, the quicker and easier the business can be expanded, the better. But suppose bad news arrives, and cash flows are far below expectations. In that case it is useful to have the option to bail out and recover the value of the project’s plant, equipment, or other assets. The option to abandon is equivalent to a put option. You exercise that abandonment option if the value recovered from the project’s assets is greater than the present value of continuing the project for at least one more period.

The binomial method is tailor-made for most abandonment options. Here is an example.

The Zircon Subductor Project

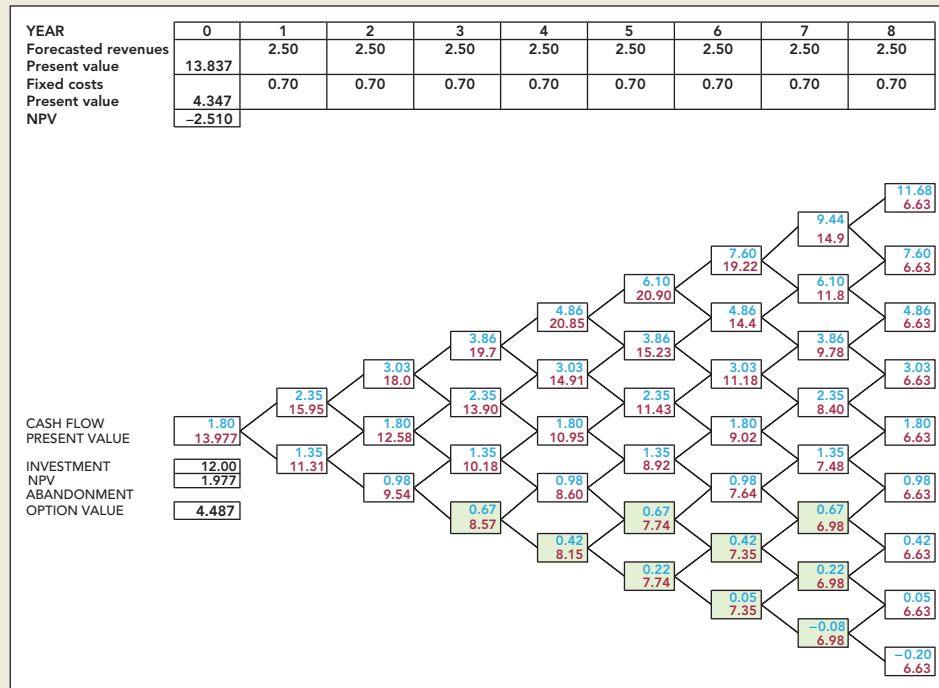
Dawn East, the chief financial officer of Maine Subductor Corp., has to decide whether to start production of zircon subductors. The investment required is \$12 million—\$2 million for roads and site preparation and \$10 million for equipment. The equipment costs \$700,000 per year (\$.7 million) to operate (a fixed cost). For simplicity, we ignore other costs and taxes.

FIGURE 22.4

Binomial tree for the Subductor project. Cash flow (top number) and end-of-period present value are shown for each node in millions of dollars. Abandonment occurs if cash flows drop into the shaded boxes. Beginning present value is about \$14 million.

eXcel

Visit us at
www.mhhe.com/bma



At today's prices, the project would generate revenues of \$2.5 million per year. Annual output will be constant, so revenue is proportional to price. If the mine were operating today, cash flow would be $\$2.5 - .7 = \1.8 million.

Calculate the Present Value of the Project The first step in a real options analysis is to value the underlying asset, that is, the project if it had no options attached. Usually this is done by discounted cash flow (DCF). In this case the chief source of uncertainty is the future selling price of zircon subductors. Therefore Ms. East starts by calculating the present value of future revenues. She perceives no upward trend in subductor prices, and ends up forecasting stable prices for the next 8 years. Fixed costs are constant at \$.7 million. The top panel of Figure 22.4 shows these cash-flow forecasts and calculates present values: about \$13.8 million for revenues, after discounting at a risk-adjusted rate of 9%, and \$4.3 million for fixed costs, after discounting at a risk-free rate of 6%.⁷ The NPV of the project, assuming no salvage value or abandonment over its 10-year life, is:

$$\begin{aligned} \text{NPV} &= \text{PV}(\text{revenues}) - \text{PV}(\text{fixed costs}) - \text{investment required} \\ &= \$13.84 - 4.35 - 12.00 = -\$2.51 \text{ million} \end{aligned}$$

This NPV is slightly negative, but Ms. East has so far made no provision for abandonment.

Build a Binomial Tree Now Ms. East constructs a binomial tree for revenues and PV(revenues). She notes that subductor prices have followed a random walk with an annual standard deviation of about 20%. She constructs a binomial tree with one step per year. The "up" values for revenues are 122% of the prior year's revenues. The "down" values are 82%

⁷ Why calculate present values for revenues and fixed costs separately? Because it's easier to construct a binomial tree for revenues, which can be assumed to follow a random walk with constant standard deviation. We will subtract fixed costs after the binomial tree is constructed.

of prior revenues.⁸ Thus, the up and down revenues for year 1 are $\$2.5 \times 1.22 = \3.05 and $\$2.5 \times .82 = \2.05 million, respectively. After deduction of fixed costs, the up and down cash flows are $\$2.35$ and $\$1.35$ million, respectively. The first two years of the resulting tree are shown below (figures in millions of dollars).

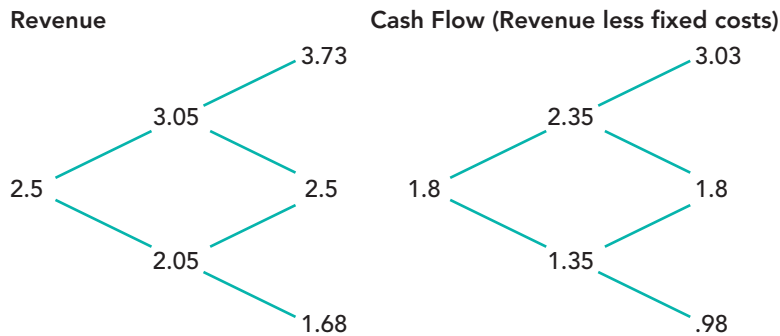


Figure 22.4 shows the whole tree, starting with cash flows in year 1. (Maine Subductor can't generate any revenues in year 0 because it hasn't started production yet.) The top number at each node is cash flow. The bottom number is the *end-of-year* present value of *all* subsequent cash flows, including the value of the production equipment when the project ends or is abandoned. We will see in a moment how these present values are calculated.

Finally, Ms. East calculates the risk-neutral probabilities of up and down changes in revenues, p and $1 - p$, respectively. Here she must pause to make sure that each year's revenue is valued properly. Remember that we have discounted revenues at a risk-adjusted rate of 9%. Thus the present value of year 1 revenues is not \$2.5 million, but only

$$PV = \frac{2.5}{1.09} = \$2.29 \text{ million}$$

Therefore, Ms. East needs to calculate the risk-neutral probabilities that generate an expected return equal to the 6% risk-free rate.⁹

$$\text{Expected return} = \frac{[3.05p + 2.05(1 - p)]}{2.29} - 1 = .06$$

$$\text{Probability of up change} = .382$$

$$\text{Probability of down change} = .618$$

Ms. East can use these probabilities at every node of the binomial tree, because the proportional up and down moves are the same at each node.

Solve for Optimal Abandonment and Project Value Ms. East has assumed a project life of 8 years. At that time the production equipment, which normally depreciates by about 5% per year, should be worth \$6.63 million. This salvage value represents what the equipment could be sold for, or its value to Maine Subductor if shifted to another use.

Now let's calculate this project's value in the binomial tree. We start at the far right of Figure 22.4 (year 8) and work back to the present. The company will abandon for sure in year 8, when the ore body is exhausted. Thus we enter the ending salvage value (\$6.63 million) as the end-of-year value in year 8. Then we step back to year 7.

Suppose that Maine Subductor ends up in the best possible place in that year, where cash flow is \$9.44 million. The upside payoff if the company does not abandon is the "up" node

⁸ The formula (given in Section 21-2) for the up percentage is $u = e^{\sigma\sqrt{h}}$ where σ is the standard deviation per year and h is the interval as a fraction of a year. In this case, $h = 1$ and $e^{\sigma} = e^{.2} = 1.22$. The down value is $d = 1/u = .82$.

⁹ Notice that the "up" revenues are 122% of today's revenue level, but 133% of the present value of next year's forecasted revenues. Thus the "up" probability required to generate a 6% average return is relatively small.

in year 8: $11.68 + 6.63 = \$18.31$ million. The downside payoff is $7.60 + 6.63 = \$14.23$ million. The present value, using the risk-neutral probabilities, is:

$$PV = \frac{(.382 \times 18.31) + (.618 \times 14.23)}{1.06} = \$14.90 \text{ million}$$

The company could abandon at the end of year 7, realizing salvage value of \$6.98 million, but continuing is better. We therefore enter \$14.90 million as the end-of-year value at the top node for year 7 in Figure 22.4.

We can fill in the end-of-period values for the other nodes in year 7 by the same procedure. But at some point, as we step down to lower and lower cash flows, there will come a node where it's better to bail out than continue. This occurs when cash flow is \$.67 million. The present value of continuing is only:

$$PV = \frac{.382 \times (.98 + 6.63) + .618 \times (.42 + 6.63)}{1.06} = \$6.85 \text{ million}$$

The payoff to abandonment is \$6.98 million, so that payoff is entered as the value in year 7 at all nodes with cash flows equal to or less than \$.67 million.

The cash flows and end-of-year values for year 7 are the payoffs to continuing from year 6. We then calculate values in year 6, checking at each node whether to abandon. Repeat this drill for year 5, then year 4, and so on back to year 0. In this example, Maine Subductor should abandon the project if cash flows drop to \$.67 million or less in each year. We have colored the nodes in Figure 22.4 where abandonment occurs.

Solving back through the binomial tree, we get a present value of \$13.977 million at year 0, and net present value of $\$13.977 - 12.0 = \1.977 million.¹⁰ If there were no option to abandon, the DCF valuation would be $-\$2.51$ million. Therefore the option to abandon is worth $\$1.977 + 2.510 = \4.487 million.¹¹ In an APV format,

$$\begin{aligned} APV &= NPV \text{ with no abandonment} + \text{abandonment option value} \\ &= -2.51 + 4.487 = +\$1.977 \text{ million} \end{aligned}$$

The project looks good, although Ms. East may wish to check out the timing option. She could decide to wait.

Abandonment Value and Project Life

Ms. East assumed that the zircon subductor project had a definite 8-year life. But most projects' economic lives are not known at the start. Cash flows from a new product may last only a year or so if the product fails in the marketplace. But if it succeeds, that product, or variations or improvements of it, could be produced for decades.

A project's economic life can be just as hard to predict as the project's cash flows. Yet in standard DCF capital-budgeting analysis, that life is assumed to end at a fixed future date. Real options analysis allows us to relax that assumption. Here is the procedure.¹²

1. Forecast the range of possible cash flows well beyond your best guess of the project's economic life. Suppose, for example, that your guess is 8 years. You could prepare a binomial tree like Figure 22.4 stretching out 20 years into the future.

¹⁰ We will spare you the calculations. You can check them, however. The "live" spreadsheet for Figure 22.4 is on this book's Web site (www.mhhe.com/bma).

¹¹ It turns out, however, that the value of *early* abandonment in this example is relatively small. Suppose that Maine Subductor could recover salvage value of \$6.63 million in year 8, but not before. The present value of this recovery in year 0, using a 6% discount rate, is \$4.16 million. APV in this case is $-2.51 + 4.16 = \$1.65$ million, a little less than the APV of \$1.977 million if early abandonment is allowed.

¹² See S. C. Myers and S. Majd, "Abandonment Value and Project Life," in *Advances in Futures and Options Research*, ed. F. J. Fabozzi (Greenwich, CT: JAI Press, 1990).

- Then value the project, including its abandonment value. In the best upside scenarios, project life will be 20 years—it will never make sense to abandon before year 20. In the worst downside scenarios, project life will be much shorter, because the project will be more valuable dead than alive. If your original guess about project life is right, then in intermediate scenarios, where actual cash flows match expectations, abandonment will occur around year 8.

This procedure links project life to the performance of the project. It does not impose an arbitrary ending date, except in the far distant future.

Temporary Abandonment

Companies are often faced with complex options that allow them to abandon a project *temporarily*, that is, to mothball it until conditions improve. Suppose you own an oil tanker operating in the short-term spot market. (In other words, you charter the tanker voyage by voyage, at whatever short-term charter rates prevail at the start of the voyage.) The tanker costs \$50 million a year to operate and at current tanker rates it produces charter revenues of \$52.5 million per year. The tanker is therefore profitable but scarcely cause for celebration. Now tanker rates dip by about 10%, forcing revenues down to \$47 million. Do you immediately lay off the crew and mothball the tanker until prices recover? The answer is clearly yes if the tanker business can be turned on and off like a faucet. But that is unrealistic. There is a fixed cost to mothballing the tanker. You don't want to incur this cost only to regret your decision next month if rates rebound to their earlier level. The higher the costs of mothballing and the more variable the level of charter rates, the greater the loss that you will be prepared to bear before you call it quits and lay up the boat.

Suppose that eventually you do decide to take the boat off the market. You lay up the tanker temporarily.¹³ Two years later your faith is rewarded; charter rates rise, and the revenues from operating the tanker creep above the operating cost of \$50 million. Do you reactivate immediately? Not if there are costs to doing so. It makes more sense to wait until the project is well in the black and you can be fairly confident that you will not regret the cost of bringing the tanker back into operation.

These choices are illustrated in Figure 22.5. The blue line shows how the value of an operating tanker varies with the level of charter rates. The black line shows the value of the




FIGURE 22.5

An oil tanker should be mothballed when tanker rates fall to M, where the tanker's value if mothballed is enough above its value in operation to cover mothballing costs. The tanker is reactivated when rates recover to R.

¹³ We assume it makes sense to keep the tanker in mothballs. If rates fall sufficiently, it will pay to scrap the tanker.

tanker when mothballed.¹⁴ The level of rates at which it pays to mothball is given by M and the level at which it pays to reactivate is given by R . The higher the costs of mothballing and reactivating and the greater the variability in tanker rates, the further apart these points will be. You can see that it will pay for you to mothball as soon as the value of a mothballed tanker reaches the value of an operating tanker plus the costs of mothballing. It will pay to reactivate as soon as the value of a tanker that is operating in the spot market reaches the value of a mothballed tanker plus the costs of reactivating. If the level of rates falls below M , the value of the tanker is given by the black line; if the level is greater than R , value is given by the blue line. If rates lie between M and R , the tanker's value depends on whether it happens to be mothballed or operating.

22-4 Flexible Production

Flexible production means the ability to vary production inputs or outputs in response to fluctuating demand or prices. Take the case of CT (combustion-turbine) generating plants, which are designed to deliver short bursts of peak-load electrical power. CTs can't match the thermal efficiency of coal or nuclear power plants, but CTs can be turned on or off on short notice. The coal plants and "nukes" are efficient only if operated on "base load" for long periods.

The profits from operating a CT depend on the *spark spread*, that is, on the difference between the price of electricity and the cost of the natural gas used as fuel. CTs are money-losers at average spark spreads, but the spreads are volatile and can spike to very high levels when demand is high and generating capacity tight. Thus a CT delivers a series of call options that can be exercised day by day (even hour by hour) when spark spreads are sufficiently high. The call options are normally out-of-the-money (CTs typically operate only about 5% of the time), but the money made at peak prices makes investment in the CTs worthwhile.

The volatility of spark spreads depends on the correlation between the price of electricity and the price of natural gas used as fuel. If the correlation were 1.0, so that electricity and natural gas prices moved together dollar for dollar, the spark spread would barely move from its average value, and the options to operate the gas turbine would be worthless. But in fact the correlation is less than 1.0, so the options are valuable. In addition, some CTs are set up to give a further option, because they can be run on oil as well as natural gas.¹⁵

In this example, the output is the same (electricity); option value comes from the ability to vary the amount produced and the raw materials used (natural gas or oil). In other cases, option value comes from the flexibility to switch from product to product using the same production facilities. For example, textile firms have invested heavily in computer-controlled knitting machines, which allow production to shift from product to product, or from design to design, as demand and fashion dictate.

Flexibility in *procurement* can also have option value. For example, a computer manufacturer planning next year's production must also plan to buy components, such as disk drives and microprocessors, in large quantities. Should it strike a deal today with the component manufacturer? This locks in the quantity, price, and delivery dates. But it also gives

¹⁴ Dixit and Pindyck estimate these thresholds for a medium-sized tanker and show how they depend on costs and the volatility of freight rates. See A. K. Dixit and R. S. Pindyck, *Investment under Uncertainty* (Princeton, NJ: Princeton University Press, 1994), Chapter 7. Brennan and Schwartz provide an analysis of a mining investment that also includes an option to shut down temporarily. See M. Brennan and E. Schwartz, "Evaluating Natural Resource Investments," *Journal of Business* 58 (April 1985), pp. 135–157.

¹⁵ Industrial steam and heating systems can also be designed to switch between fuels, depending on relative fuel costs. See N. Kulatilaka, "The Value of Flexibility: The Case of a Dual-Fuel Industrial Steam Boiler," *Financial Management* 22 (Autumn 1993), pp. 271–280.

Valuing Flexibility

With the help of faculty from Stanford University, Hewlett-Packard has experimented with real options since the beginning of the 1990s. Example: In the '80s, HP customized inkjet printers for foreign markets at the factory, then shipped them in finished form to warehouses. Customizing at the factory is cheaper than customizing in the field. But HP kept guessing wrong on demand and ending up with, say, too many printers configured for French customers but not enough for Germans.

Executives realized that it would be smarter to ship partially assembled printers and then customize them at the warehouse, once it had firm orders. True, local

customization costs more. But even though production costs rose, HP saved \$3 million a month by more effectively matching supply to demand, says Corey A. Billington, a former Stanford professor who directs HP's Strategic Planning & Modeling group.

Common sense? Sure. But you can also view it as a neat solution of a real-options problem. Increasing the cost of production—anathema to your average engineer—was in effect the price HP paid for the option to delay configuration choices until the optimal time.

Source: P. Coy, "Exploiting Uncertainty." Reprinted from June 7, 1999 issue of *Business Week* by special permission, © 1999 by The McGraw-Hill Companies, Inc.

up flexibility, for example, the ability to switch suppliers next year or buy at a "spot" price if next year's prices are lower.

The Finance in Practice box features another example of the value of flexibility in production or procurement.

22-5 Aircraft Purchase Options

For our final example, we turn to the problem confronting airlines that order new airplanes for future use. In this industry lead times between an order and delivery can extend to several years. Long lead times mean that airlines that order planes today may end up not needing them. You can see why an airline might negotiate for an aircraft purchase *option*.

In Section 10-3, we used aircraft purchase options to illustrate the option to expand. What we said there was the truth, but not the whole truth. Let's take another look. Suppose an airline forecasts a need for a new Airbus A320 four years hence.¹⁶ It has at least three choices.

- *Commit now.* It can commit now to buy the plane, in exchange for Airbus's offer of locked-in price and delivery date.
- *Acquire option.* It can seek a purchase option from Airbus, allowing the airline to decide later whether to buy. A purchase option fixes the price and delivery date if the option is exercised.
- *Wait and decide later.* Airbus will be happy to sell another A320 at any time in the future if the airline wants to buy one. However, the airline may have to pay a higher price and wait longer for delivery, especially if the airline industry is flying high and many planes are on order.

¹⁶ The following example is based on J. E. Stonier, "What Is an Aircraft Purchase Option Worth? Quantifying Asset Flexibility Created through Manufacturer Lead-Time Reductions and Product Commonality," in *Handbook of Airline Finance*, ed. G. F. Butler and M. R. Keller. © 1999 Aviation Week Books.

	<u>Year 0</u>	<u>Year 3</u>	<u>Year 4</u>	<u>Year 5 or later</u>
Buy option	Airline and manufacturer set price and delivery date	Exercise? (Yes or no)	Aircraft delivered if option exercised	
Wait	Wait and decide later	Buy now? If yes, negotiate price and wait for delivery.		Aircraft delivered if purchased at year 3.

FIGURE 22.6

This aircraft purchase option, if exercised at year 3, guarantees delivery at year 4 at a fixed price. Without the option, the airline can still order the plane at year 3, but the price is uncertain and the wait for delivery longer.

Source: Adapted from Figure 17–17 in J. Stonier, “What Is an Aircraft Purchase Option Worth? Quantifying Asset Flexibility Created through Manufacturer Lead-Time Reductions and Product Commonality,” *Handbook of Airline Finance*, ed. G.F. Butler and M.R. Keller. Copyright 1999 Aviation Week Books; Reprinted with permission from The McGraw-Hill Companies, Inc.

The top half of Figure 22.6 shows the terms of a typical purchase option for an Airbus A320. The option must be exercised at year 3, when final assembly of the plane will begin. The option fixes the purchase price and the delivery date in year 4. The bottom half of the figure shows the consequences of “wait and decide later.” We assume that the decision will come at year 3. If the decision is “buy,” the airline pays the year-3 price and joins the queue for delivery in year 5 or later.

The payoffs from “wait and decide later” can never be better than the payoffs from an aircraft purchase option, since the airline can discard the option and negotiate afresh with Airbus if it wishes. In most cases, however, the airline will be better off in the future with the option than without it; the airline is at least guaranteed a place in the production line, and it may have locked in a favorable purchase price. But how much are these advantages worth today, compared to the wait-and-see strategy?

Figure 22.7 illustrates Airbus’s answers to this problem. It assumes a three-year purchase option with an exercise price equal to the current A320 price of \$45 million. The present value of the purchase option depends on both the NPV of purchasing an A320 at that price and on the forecasted wait for delivery if the airline does *not* have a purchase option but nevertheless decides to place an order in year 3. The longer the wait in year 3, the more valuable it is to have the purchase option today. (Remember that the purchase option holds a place in the A320 production line and guarantees delivery in year 4.)

If the NPV of buying an A320 today is very high (the right-hand side of Figure 22.7), future NPV will probably be high as well, and the airline will want to buy regardless of whether it has a purchase option. In this case the value of the purchase option comes mostly from the value of guaranteed delivery in year 4.¹⁷ If the NPV is very low, then the option has low value because the airline is unlikely to exercise it. (Low NPV today probably means low NPV in year 3.) The purchase option is worth the most, compared to the wait-and-decide-later strategy, when NPV is around zero. In this case the airline can exercise the option, getting a good price and early delivery, if future NPV is higher than expected; alternatively, it can walk away from the option if NPV disappoints. Of course,

¹⁷ The Airbus real-options model assumes that future A320 prices will be increased when demand is high, but only to an upper bound. Thus the airline that waits and decides later may still have a positive-NPV investment opportunity if future demand and NPV are high. Figure 22.7 plots the *difference* between the value of the purchase option and this wait-and-see opportunity. This difference can shrink when NPV is high, especially if forecasted waiting times are short.

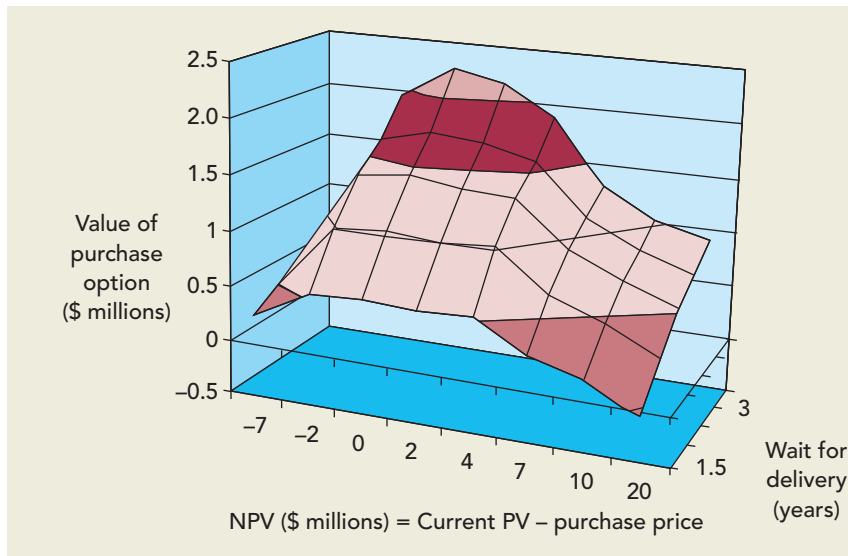


FIGURE 22.7

Value of aircraft purchase option—the extra value of the option versus waiting and possibly negotiating a purchase later. (See Figure 22.6.) The purchase option is worth most when NPV of purchase now is about zero and the forecasted wait for delivery is long.

Source: Adapted from Fig. 17–20 in J. Stonier, “What Is an Aircraft Purchase Option Worth? Quantifying Asset Flexibility Created Through Manufacturer Lead-Time Reductions and Product Commonality,” in *Handbook of Aviation Finance*, ed. G. F. Butler and M. R. Keller. © 1999 Aviation Week Books; Reprinted with permission from The McGraw-Hill Companies, Inc.

if it walks away, it may still wish to negotiate with Airbus for delivery at a price lower than the option’s exercise price.

We have cruised by many of the technical details of Airbus’s valuation model for purchase options. But the example does illustrate how real-options models are being built and used. By the way, Airbus offers more than just plain-vanilla purchase options. Airlines can negotiate “rolling options,” which lock in price but do not guarantee a place on the production line. (Exercise of the rolling option means that the airline joins the end of the queue.) Airbus also offers a purchase option that includes the right to switch from delivery of an A320 to an A319, a somewhat smaller plane.

22-6 A Conceptual Problem?

In this chapter we have suggested that option pricing models can help to value the real options in capital investment decisions. But that raises a question.

When we introduced option pricing models in Chapter 21, we suggested that the trick is to construct a package of the underlying asset and a loan that would give exactly the same payoffs as the option. If the two investments do not sell for the same price, then there are arbitrage possibilities. But many assets are not freely traded. This means that we can no longer rely on arbitrage arguments to justify the use of option models.

The risk-neutral method still makes practical sense, however. It’s really just an application of the *certainty-equivalent* method introduced in Chapter 9.¹⁸ The key assumption—implicit till now—is that the company’s *shareholders* have access to assets with the same risk characteristics (e.g., the same beta) as the capital investments being evaluated by the firm.

Think of each real investment opportunity as having a “double,” a security or portfolio with identical risk. Then the expected rate of return offered by the double is also the cost of capital for the real investment and the discount rate for a DCF valuation of the investment project. Now what would investors pay for a real *option* based on the project? The same as for an identical traded option written on the double. This traded option does not

¹⁸ Use of risk-neutral probabilities converts future cash flows to certainty equivalents, which are then discounted to present value at a risk-free rate.

have to exist; it is enough to know how it would be valued by investors, who could employ either the arbitrage or the risk-neutral method. The two methods give the same answer, of course.

When we value a real option by the risk-neutral method, we are calculating the option's value if it could be traded. This exactly parallels standard capital budgeting. Shareholders would vote unanimously to accept any capital investment whose market value *if traded* exceeds its cost, as long as they can buy traded securities with the same risk characteristics as the project. This key assumption supports the use of both DCF and real-option valuation methods.

Practical Challenges

The challenges in applying real-options analysis are not conceptual but practical. It isn't always easy. We can tick off some of the reasons why.

First, real options can be complex, and valuing them can absorb a lot of analytical and computational horsepower. Whether you want to invest in that horsepower is a matter for business judgment. Sometimes an approximate answer now is more useful than a "perfect" answer later, particularly if the perfect answer comes from a complicated model that other managers will regard as a black box. One advantage of real options analysis, if you keep it simple, is that it's relatively easy to explain. Complex decision trees can often be described as the payoffs to one or two simple call or put options.

The second problem is lack of *structure*. To quantify the value of a real option, you have to specify its possible payoffs, which depend on the range of possible values of the underlying asset, exercise prices, timing of exercise, etc. In this chapter we have taken well-structured examples where it is easy to see the road map of possible outcomes. In other cases you may not have a road map. For example, reading this book can enhance your personal call option to work in financial management, yet we suspect that you would find it hard to write down how that option would change the binomial tree of your entire future career.

A third problem can arise when your *competitors* have real options. This is not a problem in industries where products are standardized and no single competitor can shift demand and prices. But when you face just a few key competitors, all with real options, then the options can interact. If so, you can't value your options without thinking of your competitors' moves. Your competitors will be thinking in the same fashion.

An analysis of competitive interactions would take us into other branches of economics, including game theory. But you can see the danger of assuming passive competitors. Think of the timing option. A simple real-options analysis will often tell you to wait and learn before investing in a new market. Be careful that you don't wait and learn that a competitor has moved first.¹⁹

Given these hurdles, you can understand why systematic, quantitative valuation of real options is restricted mostly to well-structured problems like the examples in this chapter. The qualitative implications of real options are widely appreciated, however. Real options give the financial manager a conceptual framework for strategic planning and thinking about capital investments. If you can identify and understand real options, you will be a more sophisticated consumer of DCF analysis and better equipped to invest your company's money wisely.

Understanding real options also pays off when you can *create* real options, adding value by adding flexibility to the company's investments and operations. For example, it may be better to design and build a series of modular production plants, each with capacity of 50,000 tons per year of magnesium alloy, than to commit to one large plant with

¹⁹ Being the first mover into a new market is not always the best strategy, of course. Sometimes later movers win. For a survey of real options and product-market competition, see H. Smit and L. Trigeorgis, *Strategic Investment, Real Options and Games* (Princeton, NJ: Princeton University Press, 2004).

capacity of 150,000 tons per year. The larger plant will probably be more efficient because of economies of scale. But with the smaller plants, you retain the flexibility to expand in step with demand and to defer investment when demand growth is disappointing.

Sometimes valuable options can be created simply by “overbuilding” in the initial round of investment. For example, oil-production platforms are typically built with vacant deck space to reduce the cost of adding equipment later. Undersea oil pipelines from the platforms to shore are often built with larger diameters and capacity than production from the platform will require. The additional capacity is then available at low cost if additional oil is found nearby. The extra cost of a larger-diameter pipeline is much less than the cost of building a second pipeline later.

In Chapter 21 you learned the basics of option valuation. In this chapter we described four important real options:

1. *The option to make follow-on investments.* Companies often cite “strategic” value when taking on negative-NPV projects. A close look at the projects’ payoffs reveals call options on follow-on projects in addition to the immediate projects’ cash flows. Today’s investments can generate tomorrow’s opportunities.
2. *The option to wait (and learn) before investing.* This is equivalent to owning a call option on the investment project. The call is exercised when the firm commits to the project. But rather than exercising the call immediately, it’s often better to defer a positive-NPV project in order to keep the call alive. Deferral is most attractive when uncertainty is great and immediate project cash flows—which are lost or postponed by waiting—are small.
3. *The option to abandon.* The option to abandon a project provides partial insurance against failure. This is a put option; the put’s exercise price is the value of the project’s assets if sold or shifted to a more valuable use.
4. *The option to vary the firm’s output or its production methods.* Firms often build flexibility into their production facilities so that they can use the cheapest raw materials or produce the most valuable set of outputs. In this case they effectively acquire the option to exchange one asset for another.

We should offer here a healthy warning: The real options encountered in practice are often complex. Each real option brings its own issues and trade-offs. Nevertheless the tools that you have learned in this and previous chapters can be used in practice. The Black–Scholes formula often suffices to value expansion options. Problems of investment timing and optimal abandonment can be tackled with binomial trees.

Binomial trees are cousins of decision trees. You work back through binomial trees from future payoffs to present value. Whenever a future decision needs to be made, you figure out the value-maximizing choice, using the principles of option pricing theory, and record the resulting value at the appropriate node of the tree.

Don’t jump to the conclusion that real-option-valuation methods can replace discounted cash flow (DCF). First, DCF works fine for safe cash flows. It also works for “cash cow” assets—that is, for assets or businesses whose value depends primarily on forecasted cash flows, not on real options. Second, the starting point in most real-option analyses is the present value of an underlying asset. To value the underlying asset, you typically have to use DCF.

Real options are rarely traded assets. When we value a real option, we are estimating its value if it could be traded. This is the standard approach in corporate finance, the same approach taken in DCF valuations. The key assumption is that shareholders can buy traded securities or



SUMMARY