Chapter 10

Valuing a Start-up

Portlandia Ale was two guys and a dream. The company was started by two brewers who wanted to develop their own products. The company needed $4 million to begin product development and manufacturing and another $12 million in two years for the market launch. The entrepreneurs were very optimistic about their business opportunity, despite considerable uncertainty about the value of the market opportunity they were chasing. In other words, Portlandia Ale was a typical start-up. How much was it worth?

Figure 10.1 shows Portlandia's simple start-up plan. The company needed $4 million to get started and an additional $12 million to launch the product. Prospective investors were worried about the marketplace. Once distribution and sales began, Portlandia would earn the same margins as other microbrewers and would be expected to be able to sustain its future growth from profits. The microbrewery segment of the beer market was becoming increasingly crowded, however, and sales migrated across products. Portlandia knew that it would need to continue to develop new products as an established microbrewery, but that the value of a first product introduction was uncertain.

The brewers had tried to value the company using a discounted cash flow analysis, but the result had been a negative value. As they searched for financing, the brewers had shown prospective investors a business plan with a single very optimistic forecast for future sales that everyone knew did not accurately reflect the risks of the start-up. Portlandia's value came largely from its growth option: Today's $4 million investment creates the opportunity to make the $12 million investment in two years to become an established microbrewer.

Figure 10.1 Portlandia Ale's Business Plan
Portlandia's business plan calls for the microbrewer to raise $4M in initial funding and $12M in follow-on funding 2 years later. The follow-on investment will be made only if the market value of an established brewer exceeds the required investment.

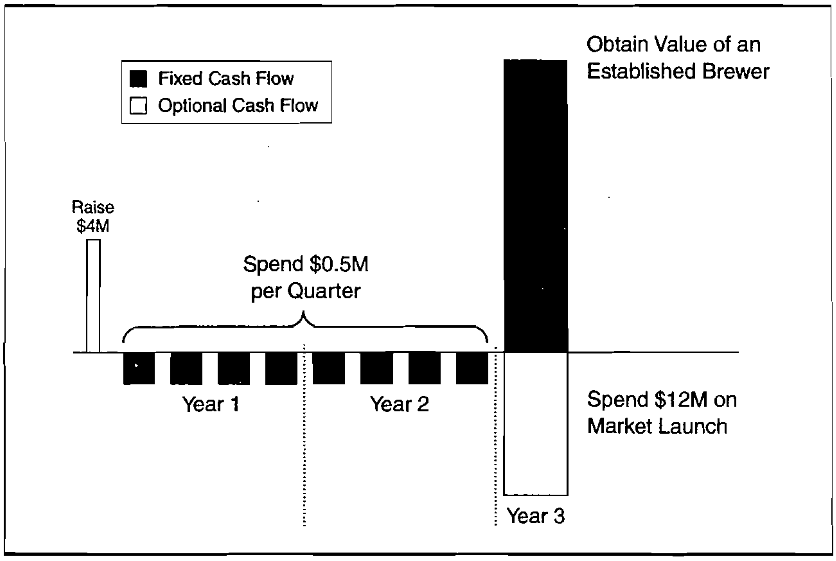

Parallels to Other Applications

Many growth opportunities derive their value from future contingent investment decisions and are nearly impossible to value with discounted cash flow and other conventional tools. This case study shows how the value of a start-up is derived from growth options. Although the firm in this case is private, the results obtained by the real options approach are aligned with financial market valuations of established firms in the industry.

Many projects and investment opportunities are similar to a start-up venture, and this simple model of a growth option covers a wide range of applications, including corporate R&D, "intraprenurial" ventures, and platform projects. The real options approach shows that using realistic inputs, investments that embed growth options will have higher expected returns, conditional on survival, than the underlying asset. This helps to explain differences in performance across firms in the same industry.

The Questions

What is the value of Portlandia Ale that correctly reflects its growth option? What company-specific and financial market information is used in this valuation?

The Application Frame

If all goes well, in two years Portlandia will have a viable business with a complete product line moving through its distributor. The company will have arrived as an established microbrewer with a business model similar to that of other microbrewers. Discounted cash flow, or use of comparables such as P/E ratios, can be used to value the firm as an established microbrewer because the value is largely based on cash flow.

Before establishing itself, Portlandia has a growth option. The option's exercise is triggered by the value of the firm as an established microbrewer, which is calculated as follows:

- Forecast Portlandia's sales level in two years, as an established business. The brewers expected that this would be $6 million per year.
- Calculate the average of the current market value–to–sales ratio for established publicly traded microbrewers whose current sales are in the same range. (Make sure not to select firms with valuable growth options, but rather ones that are representative of the business model of the industry for established firms.) The average market value-to-sales ratio of three of these firms as of March 1997 was 3.66.
- The current value of Portlandia's market opportunity, the value of being an established firm, is $22 million ($6 million × 3.66). This is the initial value of the underlying asset in the growth option.

The average stock price volatility of three established microbrewers is estimated to be 40%. It is expected that the microbrewer's value will fluctuate in value with financial market revisions of the market-to-sales ratio or from revisions in Portlandia's sales forecast that arise from industrywide factors, such as a change in alcoholic beverage taxes or increased bottling costs from recycling restrictions.[1]

The real options analysis recognizes the uncertainty about the value of an established microbrewer, and it also recognizes that the second investment opportunity limits some of the downside risk because it is contingent on a strong value of the market opportunity. The option to invest in two years is analogous to a European call option on the value of an established microbrewery. The Black-Scholes equation can be used to get a first-cut valuation.

The Results

For comparison, we first calculate the value of Portlandia Ale using the discounted cash flow method. The analysis assumes that Portlandia will follow a fixed strategy of expanding into a full-fledged microbrewer regardless of market conditions in two years. This results in a -$0.23 million valuation. A risk-adjusted discount rate for the microbrewer of 21% is used.[2]

A simple real options analysis includes the contingent nature of the launch decision in two years. Portlandia and its investors will commit $12 million required for launch only if the value of Portlandia as an established brewery at that time exceeds $12 million. Although the forecast in the discounted cash flow analysis was for a $22 million valuation, this is only a forecast. The actual realization may be substantially higher or lower.

The Black-Scholes equation is used to value the growth option (see Figure 8.7). The input parameters are A = $14.46 million (the present value of $22 million), X = $12 million, $\sigma = 40\%$ per year, $r = 5\%$, and $T = 2$. The result is $4.96 million, the value of the option to invest $12 million in two years' time. Hence, the value of the firm is $1.13 million after accounting for the present value of the $4 million required during the next two years ($1.13M = $4.96 – PV ($4)).

Portlandia's management and investors have even more flexibility in their investment decisions. During the first two years, they can decide at the beginning of each quarter whether to invest another $0.5 million to keep the project (with all the subsequent investment options) alive or to abandon it. The value of the firm, including this American-style option, can be obtained with the binomial option valuation model. Under these assumptions, the value of Portlandia increases to $1.74 million.

The contribution of the growth option to Portlandia's value can be seen by comparing the value of the company obtained from the real options approach and the value obtained from the discounted cash flow analysis. The difference, $1.36 million, arises from uncertainty about the value of the market opportunity and from Portlandia's contingent

P/E Ratios—Valuation Using Comparables

P/E ratios are the valuation tool favored by financial analysts on Wall Street. The P/E ratio, defined as stock price to earnings per share, is a quick summary of how much investors are willing to pay per dollar of current earnings. For a quick analysis, analysts use P/E ratios from companies comparable to the one they want to value. (Value equals earnings times the comparable company P/E ratio. Either forecasted or current earnings are used.)

The long-term stock market average P/E ratio is 13, but individual company P/E ratios vary significantly, even in the same industry. Stock prices embed growth options, and firms differ by their portfolio of growth options and their ability to execute to realize the value of these options. Also, the amount of debt issued by the firm affects its P/E ratio. Earnings are net of interest expense, so a firm with substantial debt obligations will have lower earnings and a higher P/E ratio, all else being equal.

Both P/E ratios and dynamic tracking are objective financial market information, not subjective estimates. It is hard to argue with a price determined by a set of investors who are willing to pay real dollars. However, the use of P/E ratio for valuation is based on a static selection of the comparable firms, whereas dynamic tracking defines a comparable firm as one with the same sensitivity to the same underlying asset—a much more precise basis of comparison.

Start-ups often have no earnings or earnings so small as to make a P/E ratio–based valuation appear arbitrary. Consequently, it is customary to value them using a price-to-sales ratio.

decision in two years. Not only does the real options analysis produce a more attractive valuation, it also better communicates the risks.

The real options results can also be used to calculate the minimum growth rate that Portlandia must realize if the second investment is to be made. At the end of two years, Portlandia will invest only if its value as an established brewer is greater than or equal to $12 million. ($12 million is the critical value.) Examination of the binomial tree shows that if the value of the established market was $12 million or higher, Portlandia's investors would realize a 66% return over the two-year period.

This very simple analysis has captured the essence of the contingent decisions in start-up opportunities and has established the value of a private firm from financial market values. No detailed forecasts of cash flows were used in the valuation, although pro forma forecasts will have other uses in assessing the business opportunity.

This case study valued a growth option with a single, market-priced source of uncertainty. The next case study incorporates private risk, which is present in most real options applications.

Chapter 11

Investing in a Start-up

Redwood Ventures, a venture capital fund, is evaluating the business plan of Lighthouse Services, a start-up company providing Web page design tools in the exploding Internet market. Lighthouse is asking for a $10-million investment. Venture capitalists have a few rules of thumb that they often use to evaluate prospective deals: Their investment should be doubled in the next round of financing; the size of the market opportunity must be large enough to justify the investment and the risk; and there must be an exit opportunity—a way for the venture capital fund to liquidate its holdings—within six years. Redwood uses an in-house real options tool to evaluate whether start-up business plans meet the required terms.

Lighthouse Services sent a very simple business plan to Redwood Ventures:

- Invest $10 million now.
- Other investors will put in $10 million in two years.
- Yet more investors will put in $20 million in three years, and Lighthouse will be an established firm selling an easy-to-use tool kit for Web page design.
- Investors can liquidate their holdings at a profit when the firm goes public around that time.

After evaluating the plan and meeting the management team, Redwood Ventures was seriously considering investing in Lighthouse. However, the partners at Redwood saw two major sources of uncertainty: the value of the business after the three rounds of investment

Figure 11.1 The Investment Opportunity at Lighthouse Services
Initially, Lighthouse Services faces risk from private and market-priced sources, but by the third year the firm faces only market-priced risk. Much of the firm's value comes from the opportunities to respond to unfolding events at the end of the second and third years. In this figure a binomial market outcome is shown for simplicity; in fact, there is a distribution of outcomes.

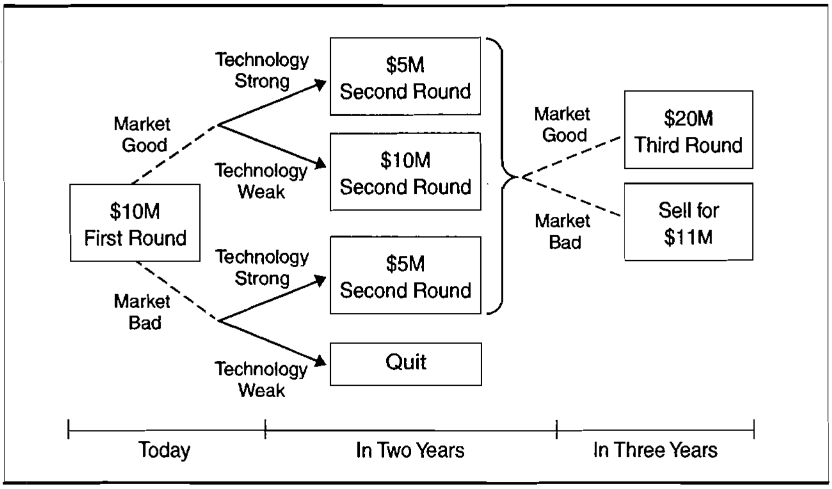

and the chance of technology success. The business plan was based on a discounted cash flow analysis, assuming that all stages of investment were made. Figure 11.1 shows the types of outcomes that everyone expected.

During the first two years, Lighthouse has private and market-priced risk. The private risk is whether the Web page tool kit could be completed at the end of two years and within the budgeted $10 million. The market-priced risk is the value of being in the Web page tool-kit business. During the third year, Lighthouse has only market-priced risk, because the technology risk is fully resolved in the first two years. The distribution of possible outcomes for the market-priced risk widens with time.

At the end of the second year, Lighthouse would need additional funding from outside investors: a $5 million investment if the technology was successfully completed and $10 million if the technology was weak. If the value of the market opportunity decreased and the technology was weak, everyone expected that the business would fold, but that decision was contingent on the value of the market opportunity in two years. At the end of three years, Lighthouse would need its last round of financing. If the business opportunity had a sufficiently high value, a final $20 million investment would be required. Otherwise, Redwood believed that Lighthouse would be sold without further investment. As a first guess, they put the value for this outcome at $11 million and planned to test whether the results were sensitive to this input.

Redwood and Lighthouse began negotiating the terms of the first $10 million investment. Discussions quickly centered around one term of the deal: What share of the equity of the firm would Redwood receive for their $10 million investment? This hinges on the value of Lighthouse after the $10 million investment is committed, V_0. (In the language of venture capitalists, V_0 is the "post-money valuation.") The value of the entrepreneur's share is smaller than V_0 minus $10 million, because the venture capitalists want to claim a bit extra equity to ensure that they obtain a 100% return in the next round of funding at the end of year 2.

The real options framework can be used to figure out the expected value of Lighthouse in year 2 and work back from this to compute the current share of the firm required to attain this rate of return.

The Questions

What is the value of Lighthouse Services? What share of the firm should Redwood receive for its $10 million investment? How should the terms of the venture capital investment be structured?

Parallels to Other Applications

Often good investment opportunities for real assets are contracts written on growth options. Venture capital investments are typical of this situation—the business opportunity is risky, but contract terms can be used to modify the risk profile of the investment. To help them navigate these murky waters, experienced investors have developed rules of thumb—proven yet simple decision criteria that are effective in characterizing investment opportunities. The real options approach provides a transparent and quantifiable framework that integrates the valuation of the growth opportunities, the contract terms, and the rules of thumb.

The framework here can be used for corporate venture capital investments; comparison of acquisitions, joint venture, and licensing opportunities; and so on.

The Application Frame

This application has two decision points, two sources of uncertainty for the first two years, and one source of uncertainty in the last year. This stylization captures the early-stage technology risk and the importance of business risk, particularly when the technology platform is complete.

In the model, private risk is treated very simply: There is a 50% chance that the technology is completed in two years and a 50% chance that it remains incomplete. Although the model could be made more "realistic" by directly modeling the factors that determine success, such as costs to complete the product, time to market, quality of marketing efforts, and so on, the importance of including more detail depends on the decisions to be made. Would the value of the firm or the terms of the investment change significantly if costs were high but quality of marketing was good compared to the case when costs were low and marketing was poor? If so, model these two private risks separately; otherwise, the simple characterization is appropriate.

Lighthouse will be selling to a retail segment of the Internet market and once fully established will have profit margins and selling expenses that are similar to those of a software publisher. To estimate the value of their market, we multiplied the total segment sales by the average market-to-sales ratio of several software publishers.[1] Lighthouse's expected share of this value is their forecasted sales divided by total segment sales. The volatility of this segment is estimated from the average stock price volatility of the software publishers. The current assessment is that the value of Lighthouse's business opportunity in three years is $130 million, which will be obtained if both stages of investment are completed.

The next step is to add the contract terms. For simplicity, only Redwood's share of equity is analyzed, but the application can be expanded to handle many other standard terms. Redwood Ventures wants a 100% rate of return on its $10 million investment if Lighthouse continues. Let y be Redwood's share of the equity and W be the expected value of Lighthouse just before the second round of financing, including the value of its growth option. The rule of thumb can be written as

$$yW \,/\, \$10 \text{ million} = 200\% \quad \text{or} \quad y = \frac{\$20 \text{ million}}{W}$$

We used a simple binomial model to get illustrative results.

The Results

The real options approach shows that Lighthouse's current value is \$38.5 million. This valuation includes a range of outcomes, including the possibility of becoming the next Yahoo! Clearly there is only a remote chance of this event; based on historical start-up success rates and Internet volatility, there are many more outcomes with moderate returns.

The \$38.5 million company valuation includes the \$10 million investment. At face value, Redwood's share of the firm might be calculated as \$10 / \$38.5 = 26%. However, Redwood wants to ensure that if a second round of financing is done, the value of the \$10 million investment must increase by 100%. The real options model was used to value the firm at each outcome. The constraint requires increasing Redwood's share of the firm to 47%. By taking a larger share, Redwood will obtain its 100% return even when Lighthouse does not grow in value by 100%. In sum, the venture capitalists invest \$10 million today in return for equity valued at \$17.86 million and the entrepreneur has 53% of the company, equity valued at \$20.14 million.

The real options framework can be "reverse-engineered" to address the size of the market opportunity rule of thumb that is used by venture capitalists.[2] Private risk reduces the sensitivity of the value of the firm to the size of the market opportunity because good news about market outcomes may be coupled with bad news about the private risk. This suggests that the rule of thumb requiring a large size of market is a good screening device; investment selections will tend to stay out of the zone of indeterminacy (good technology outcome, bad market outcome) and in the range where the firm will have substantial value.

Finally, notice the role of luck in the success of a start-up. The private risk can resolve favorably, but the value of the market opportunity can turn sour. Venture capitalists seem to recognize that failures of start-ups may not be entirely the fault of the entrepreneur. In fact, Silicon Valley has another rule of thumb: The successful entrepreneur has had two failed start-ups.

This case study has shown how two sources of uncertainty interact to change the risk profile of an investment. The real options approach provides a way to address the issues that concern managers and investors in a start-up venture: valuation and terms of investments, the role of luck, and how to structure an investment to achieve the venture capitalists' financial performance objectives.

Chapter **12**

Exploring for Oil*

Houston Oil had leased a large tract of land on the Alaskan North Slope and was evaluating alternate exploration strategies. Seismic investments would provide additional information about the amount of oil in the ground, and drilling would add information about the amount of oil and resolve whether the oil could be produced. Should Houston begin exploration? Which exploration investment strategy should they use?

Houston Oil knew the risks of oil exploration:

- Six to fifteen years needed to bring an unexplored tract into production
- The millions of dollars required for exploration and development
- The small chance, 10%, that exploration efforts would ultimately lead to oil produced from the tract

Many exploration efforts were halted because the estimates of the amount of oil in the ground and/or the geological features of the oil reservoir implied that development costs would exceed production profits or because the price of oil fell too low to justify more expenditures.

Oil exploration activities determine whether oil and gas are present and produce data to better estimate the size of the field. Development activities provide data to improve the estimate and build the

*This case study is part of a larger ongoing real options initiative at Anadarko Petroleum Corp. and reflects the joint work of Scott Albertson, Martha Amram, Ron Bain, Michael Cochran, Steve Rutherford, and Ron Simenauer.

infrastructure required for production. Most of the uncertainty about the size of the field is resolved before production begins.

Exploration and development are learning investments, a sequence of options that resolve uncertainty about the amount of oil that can be recovered from a tract. Although all options require the purchase of an opportunity to make a contingent decision, learning investments also require the holder to purchase information for decision making and to invest to keep the option alive.

The oil learning investment works as follows. The first investment creates the option to continue exploration or to move on into development. The information generated is about the amount of oil in the ground, the reserve size, and about the chance of success (COS), the physical possibility of pulling the oil out of the ground.[1] Uncertainty about the reserve size can be reduced by probing the geology with seismic waves or by drilling wells; uncertainty about the COS can be resolved with drilling. The value of the follow-on options depends on the updated geological uncertainty and the level and nature of oil price uncertainty.

Exploration decisions based on geology alone are incomplete. The value of the oil in the ground is driven by market-priced risk, oil prices. Traditional decision analysis models of oil exploration have tied geological information to one particular oil price or oil price forecast. Real options models ensure that the financial market pricing of the risk-return trade-offs is incorporated into oil exploration decisions.

Figure 12.1 shows how uncertainty about reserve size is reduced with exploration efforts. Before the first stage of exploration, estimates of the size of the reserve and uncertainty about the reserve size are obtained from other sources, such as commercially available maps or government surveys. During the first stage of exploration, the new data are analyzed and a new estimate of the size of the reserve is announced, one that can be made with additional confidence because the range of uncertainty has narrowed. The exploration efforts might increase or decrease the estimate of the reserve size. The question now is: Given the newly revised estimate, is it worthwhile to continue exploration? Figure 12.1 shows how reserve size uncertainty is reduced with four stages of exploration.

The choice between seismic and drilling investments is based on the value of the tract under each strategy, accounting for cost of exploration, updates in reserve size uncertainty, resolution of the COS (resolved with drilling, unchanged with seismic), and future exploration options. Figure 12.2 shows a sample sequence of possible oil exploration investments. The five possible decisions at the beginning of each

Figure 12.1 Learning from a Sequence of Oil Exploration Investments

Activities that generate otherwise unavailable information are learning investments. Oil exploration investments update the estimates and reduce uncertainty about the reserve size and/or the chance of success. The option to continue exploration will be exercised if the revised estimates are sufficiently good, given the reduced uncertainty about their value.

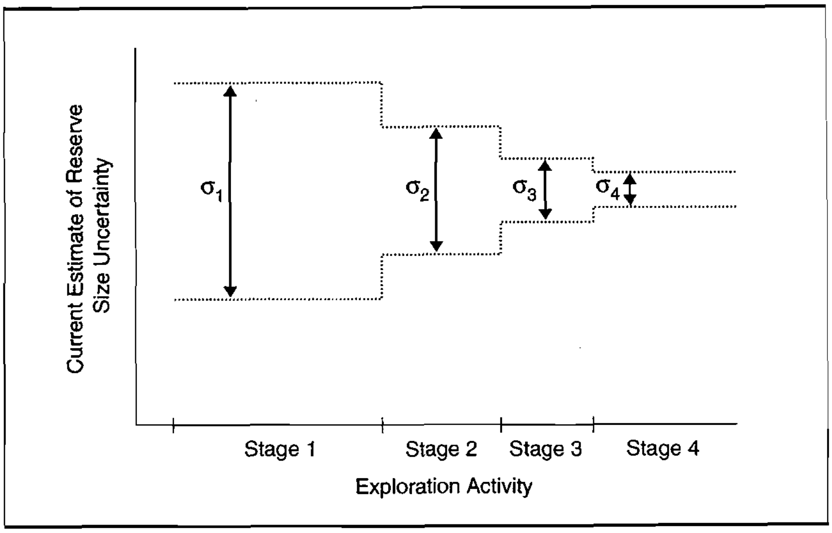

Figure 12.2 Illustrative Oil Exploration Decision Structure

Although oil exploration is a frequently cited application of the real options approach, there is considerable variation in actual implementation. For the decision structure shown, the real options outputs are the value of the reserve before exploration (accounting for the options depicted) and the optimal initial exploration strategy.

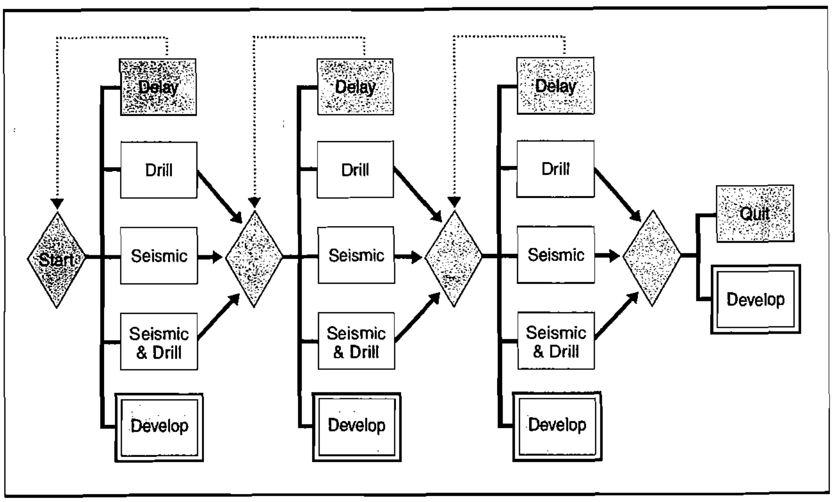

Parallels to Other Applications

Oil exploration is a clear example of a common real options application, one that involves learning investments and how much managers are willing to pay to reduce private risk. Oil prices are a market-priced risk evolving over time without investment and will influence the value of private risk information. This is one of the important links between internal decision making and financial market discipline. Examples of learning investments abound and include product development, R&D, and mining.

phase are delay, develop immediately, explore using seismic, explore using drilling, or explore using seismic and drilling. Delay is the optimal strategy when at current oil prices it is not worthwhile to invest. It is unlikely the tract would be abandoned before the end of the lease because of the chance of an increase in oil prices.

The Questions

Should the tract be explored? Which exploration strategy should be used?

The Application Frame

The oil exploration decision is made by valuing the tract under each initial exploration strategy (and all contingent follow-on strategies). The strategy that provides the highest-valued tract is chosen. The value of the tract depends on three sources of uncertainty—oil prices, reserve size, and COS—and is the value of a sequence of investment options.

The current spot price of oil is reported daily, and the volatility of oil prices is estimated as the volatility implied by option contracts on oil. The convenience yield was estimated from oil futures contracts. A common assumption in oil geology is that the distribution of reserve sizes is log-normal, and as mentioned earlier, an initial estimate of the range of uncertainty is available at the time the tract is acquired. The initial level and the standard deviation of the COS are based on historical experience in the region and experience with specific geological features. The market-priced risk is easily tracked, and the private risks are uncorrelated with any traded asset and can be accounted for by taking expectations over their range of values and discounting at the risk-free rate.

Because exploration is a learning investment, another set of inputs is needed, the reduction in uncertainty from each stage of exploration. A "learning ratio" for reserve size and one for COS can be defined as the standard deviation after the exploration investment divided by the standard deviation before the exploration investment.[2] The learning ratio is always less than one, differs by type of investment, and decreases with the stage of investment. The learning ratio for the COS with drilling equals zero, because drilling fully resolves the COS uncertainty.

This application has one source of market-priced risk, two sources of private risk, and a long sequence of options, making dynamic programming the most straightforward solution approach. At the end of the lease, the tract is either developed or abandoned. Moving back one stage, the optimal last-stage exploration decision is conditional on the optimal decision at the end of the lease. Continuing the fold back, the value of the tract under the first-stage strategy can be found.

The Results

The first result is the value of the tract under each first-stage strategy. (The subsequent investments are contingent on the first-stage results and are not yet known.) When the optimal strategy is delay, the same analysis is repeated after one year's time, with updated oil prices and a shorter time to the expiration of the lease. Immediate development will be the optimal strategy when the payoff to producing oil is so high that managers are willing to risk abandoning the tract after a sizable development expenditure.

Figure 12.3 shows the optimal first-stage exploration investment strategy as a function of the current estimate of reserve size and of current oil prices, holding the other inputs constant. This type of sensitivity analysis can be produced for any input combination and can be quite useful for decision making by committee. It might be that the group has two opinions about the value of one input, and that the disagreement is of no consequence because both input values lead to the same investment strategy.

Figure 12.3 also shows two types of waiting. In the lower left, the optimal strategy is to wait to explore because oil prices and the

Figure 12.3 Sample Strategy Space for Oil Exploration Investment
The strategy space diagram can be used to build consensus around investment strategy. In this example, holding other inputs constant, seismic is the optimal first-stage exploration strategy over a wide range of reserve size estimates.

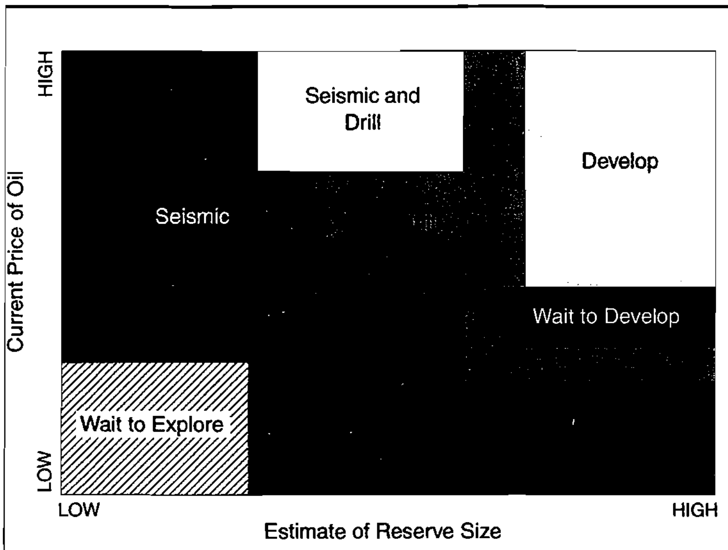

estimate of reserve size are low. In the upper right, the optimal strategy is to wait to develop because the estimated reserve size is high, but oil prices are a bit too low to justify development. This figure captures an important feature of the oil industry: an increase in oil prices brings tracts out of delay mode and into exploration, creating demand for oil services (the companies that do the exploration and development work). This raises the cost of further exploration development, reducing the value of the exploration options.

The final result is about the value of information. For oil exploration, the value of information is how much the oil company would be willing to pay to further resolve reserve size or COS uncertainty. The real options approach shows that the value of resolving uncertainty depends on the following:

- *The future:* future contingent exploration decisions, the magnitude of development costs, the range of evolution of market-priced risk, and so on
- *Current values:* the current oil price and the current estimate of the reserve size
- *The type of uncertainty:* mean-reverting, log-normal, and so on

For example, when development costs are high, the value of resolving the COS uncertainty is high, and seismic investments are not terribly valuable because they don't supply the critical piece of information. Using the real options approach, the value of information is aligned with financial market valuations.

This case study has highlighted the value of information, how much a firm would be willing to pay to resolve private risk. Although we did not emphasize this point, the framework also shows that there is a role for oil price fluctuations in exploration decisions. The next application, drug development, has a similar structure, but pharmaceutical industry practice differs from oil industry practice in that fluctuations in the value of the market opportunity have a much smaller impact on learning investments.

Chapter **13**

Developing a Drug

Recent academic research only confirmed what Mega-Pharm already knew: that the NPV of an investment in drug development was nearly zero.[1] It now costs over $300 million to develop and take a drug through the FDA approval process and nearly $500 million to market it. Drug sales vary widely; a blockbuster like Prozac could earn billions, but many drugs have life-cycle revenues of only $100 million.[2] Drug development and marketing go through well-established phases, and Mega-Pharm wondered how they could use the real options approach to better manage the process.

Developing and marketing a drug is a very expensive and lengthy process. Figure 13.1 shows a typical profile of annual expenditures over the life cycle of a drug.

For the first 14 years, a successful drug passes through various phases of preclinical and clinical testing, culminating in a filing with the Food and Drug Administration (FDA) and approval. After year 14, the drug is sold and the profile of marketing expenditures roughly follows that of sales.[3] There is considerable uncertainty about the drug's value throughout the process.

Mega-Pharm has a number of decision points during this lengthy life cycle in which it could decide to stop development or marketing efforts. If development is halted, the drug can be abandoned, licensed to another company, or sold. Drug development efforts resolve two private risks:

- Does the drug work and can it be profitably produced?
- Is the size of the market for this drug large enough to warrant further expenditures?

Figure 13.1 Basic Profile of Drug Development and Marketing Costs
Drug development can be seen as a sequence of options. The options that arise just before the peaks in spending are particularly valuable because they prevent regret. Data sources: Myers and Howe (1997) and *The Economist* (1998).

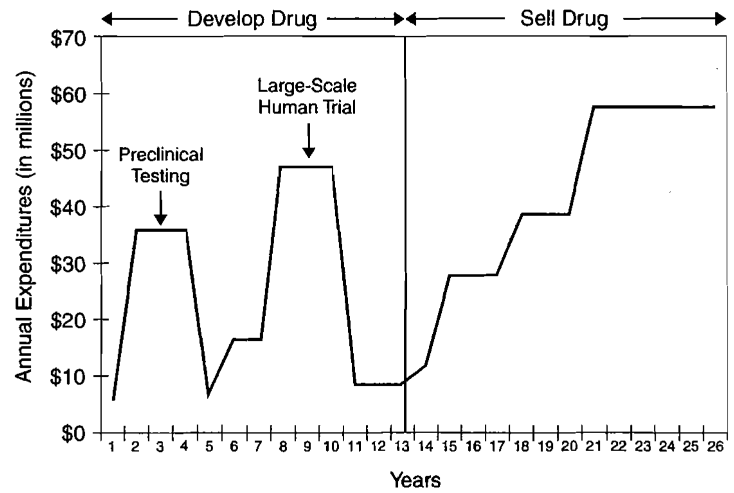

Development efforts sometimes reveal that the drug won't work on certain diseases as expected, reducing the potential sales, while early-stage marketing efforts resolve the size of the market uncertainty. At the time of the marketing launch, it is expected that the drug's sales will increase until it comes off patent protection and then slowly decrease over time. What is not known is the magnitude of peak sales, which depends on competitive entry, changes in distribution channels, and so on.

The market-priced risk for drug development is about the value of the market. This will be determined by factors influencing the entire industry, including insurance company practices, employer inclusion of pharmaceutical costs in health plans, provisions for pharmaceutical payments in government programs, and so on.

Recently Mega-Pharm had been bringing marketing managers and scientists together in an attempt to better manage the drug development process. Quantitative analyses of the drug life cycle were available in the form of NPV calculations with scenarios.

Parallels to Other Applications

Drug development is very expensive and highly uncertain, appearing to make it a terrific candidate for the real options approach to investment decision making. Drug development can be viewed as a learning investment, in which R&D investments reduce the uncertainty about the remaining costs to complete the development of the drug, and initial marketing efforts resolve uncertainty about the size of the drug's market. While market-priced risk evolves during the years of drug development, its level of uncertainty is relatively small compared to the private risk uncertainty. The real options approach to drug development relies more on evaluation of the consequences of private risk through stochastic dynamic programming than on the valuation of options by relationships that must hold to prevent arbitrage.

Real options applications that contain learning investments can be placed on a continuum, one that characterizes the relative importance of private risk to market-priced risk to option valuation. Market-priced risk is relatively important in oil exploration applications and disciplines oil exploration strategies and transacted values of reserves. Private risk is relatively important in drug development, and while the real options approach provides needed discipline for transactions to license early-stage drug projects and for joint ventures, financial market discipline has a smaller effect on transaction values.

The Questions

What are the roles of private risk and market-priced risk in drug development decisions and valuation? What are the implications of the large amount of private risk for the application frame and the solution method?

The Application Frame

The drug development and marketing process can be modeled as a sequence of learning investments and abandonment options. In each period, Mega-Pharm decides whether to spend a predetermined amount for further development or marketing or to abandon the drug. The reward for continuing is the next option. Each option contains the opportunity to make a similar decision in future periods and the possible profits late in the life cycle.

The application can be framed in terms of the following sources of uncertainty[4]:

- *Industrywide index of value.* This is a market-priced risk that captures the effects on project value from changes in regulation, changes in drug delivery by HMOs, and so on. This risk can be tracked by a stock index of drug companies, which will capture their common revisions in valuation and diversify away their firm-specific revisions (such as those arising from announcements about a particular drug or joint partner).
- *Size of market for the drug that finishes development.* This is a private risk and one that is resolved by drug development and marketing efforts. An estimate of future sales can be built up using well-established models that count the number of potential patients, disease days, quantity of drug per day, and so on. The value of the market opportunity for the drug that finishes development is the joint outcome of private and market-priced risks; the value of the market for the drug is equal to the size of the market (in units of sales) multiplied by the value of the market (dollars per unit sold). Uncertainty about the size of the market is resolved with marketing expenditures.[5]
- *Remaining life-cycle costs of the drug.* The initial value of this private risk is simply the present value of future drug development and marketing expenditures. As expenditures are made, bad news will increase the remaining costs, so this variable becomes a way to model scientific or technical risk.[6] The remaining costs can be avoided by abandoning the drug. As with the value of the market, uncertainty about the remaining life-cycle costs can be modeled as resolving with expenditures.
- *Probability of passing regulatory tests.* Based on historical experience, an estimate can be made of the probability of passing each phase of FDA tests. The cumulative probability of surviving all clinical trials and gaining FDA approval is around 12%.[7] In this application frame, these estimates are not revised during drug development.[8]

The outcome of the drug development and marketing process depends largely on the resolution of the three sources of private risk. The largest expenditures arise during marketing, and options to avoid these expenditures for a marginally profitable drug—one whose private risk resolved less than satisfactorily—can be quite valuable. Pilot marketing programs and tests of selected markets are valuable options because of the information they generate. Similarly, investments to resolve uncertainty will be most valuable just before a large expenditure, such as just before the large expenditures for preclinical testing and large-scale human trials.

The application frames for both drug development and oil exploration demonstrate the value of a sequence of learning options, but industry-specific features change the role of market-priced risk. In oil exploration, the value of exploration activities is strongly influenced by the price of oil, because oil price volatility is fairly large in magnitude and because of the option to temporarily cap a producing oil field to wait for higher oil prices. This late-stage option creates an opportunity to reverse strategies and makes current exploration decisions more sensitive to the price of oil. In contrast, the decision to stop selling a drug is irreversible. The critical value for this decision will be lower than for a reversible decision. The drug development decision is also affected by the structure of expected expenditures—a riskier later-stage investment increases the value of the option to develop a drug and makes it more likely that the drug will continue to be developed over a range of market values.

The Results

This application frame can be used to obtain a number of results for each phase of the drug life cycle. Figure 13.2 shows an illustrative set of results. Figure 13.2(a) shows that as managers use the option to abandon, the number of drugs in development falls by phase of investment. Figure 13.2(b) shows the value of a surviving drug by phase. The value of the drug increases as it passes the FDA hurdles and as uncertainty about the remaining costs and size of the market is resolved. The value of early-phase drugs is higher using a real options analysis rather than conventional tools. Figure 13.2(c) shows how the level of uncertainty about surviving drug projects decreases by phase, until all private risk is eliminated and the project uncertainty is equal to the volatility of the market-priced risk.

Figure 13.2 Illustrative Results for Drug Development
Three types of real options results are shown. The results are common to learning investments that contain valuable options to abandon.

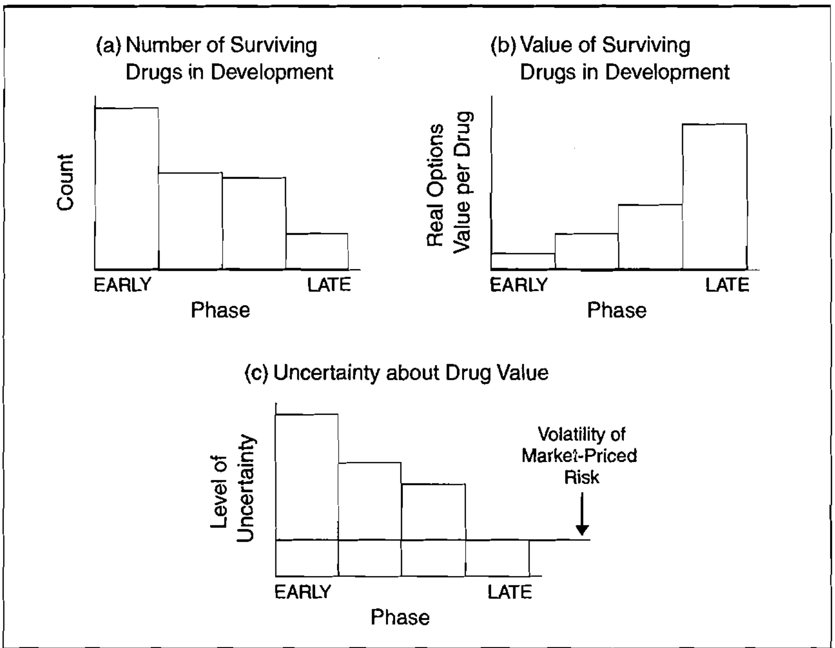

Although a discounted cash flow analysis might show that developing a drug is a zero-NPV investment, consideration of all the options during the drug's life cycle shows that the correct valuation can be much higher. For example, for the particular application frame and inputs used in one study, it was found that a drug that had a zero NPV with discounted cash flow had a value of $14.6 million using the real options approach.[9] Using the real options approach to manage a portfolio of drugs causes more drugs to start the development process and more drugs to be abandoned, a result consistent with trends in industry practice.[10] Also, realized returns to a successful drug will be quite high because they are the result of successfully executing options, while the returns across the drug portfolio should be significantly lower.

This case study has shown how industry-specific features will influence real options results. The last few cases have focused on a project-level analysis, and our next case study moves up to the next level, showing how to value the elusive company-wide contingent benefits of an infrastructure investment.

Chapter **14**

Investing in Infrastructure*

The chief information officer (CIO) of National Mortgage Trust (NMT) had been arguing with senior management for $5 million for a new document imaging system. Other managers worried about the assumptions in his analysis that supported the investment—an increase in the number of mortgages processed and a reduction in processing costs. What if the mortgage market did not grow as expected? The result was a divided team.

NMT is a small but aggressive financial institution specializing in mortgage lending. NMT currently offers several fixed or variable interest rate plans but wants to migrate from this limited menu to offering customized mortgages sold directly to the homeowner. For example, NMT hopes to market over the Internet, with customers sending in applications electronically. The CIO's view is that the imaging system would help to serve this objective by moving NMT to a paperless office environment.

Mortgage lending requires dealing with several outside customers: prospective homeowners (and the real estate agents who guide them) who apply for loans; banks and other financial institutions who assume the mortgages; and actual homeowners as they require services. NMT believed that it could become a larger and more profitable enterprise by offering better service and more customization to both homeowners and financial institutions.

---

*This chapter is based on case study prepared by E. Balasubramanian, John Henderson, Nalin Kulatilaka, Robert Materna, John Storck, and Janet Wilson.

Figure 14.1 From Strategic Vision to Information Technology Investment at NMT

The strategic vision is translated into tangible business goals and then into desired competencies. Comparison of desired and current competencies surfaces the competency gap, which is closed by physical and organizational investments. For NMT, the image processing system was the key information technology investment required to obtain the desired competencies.

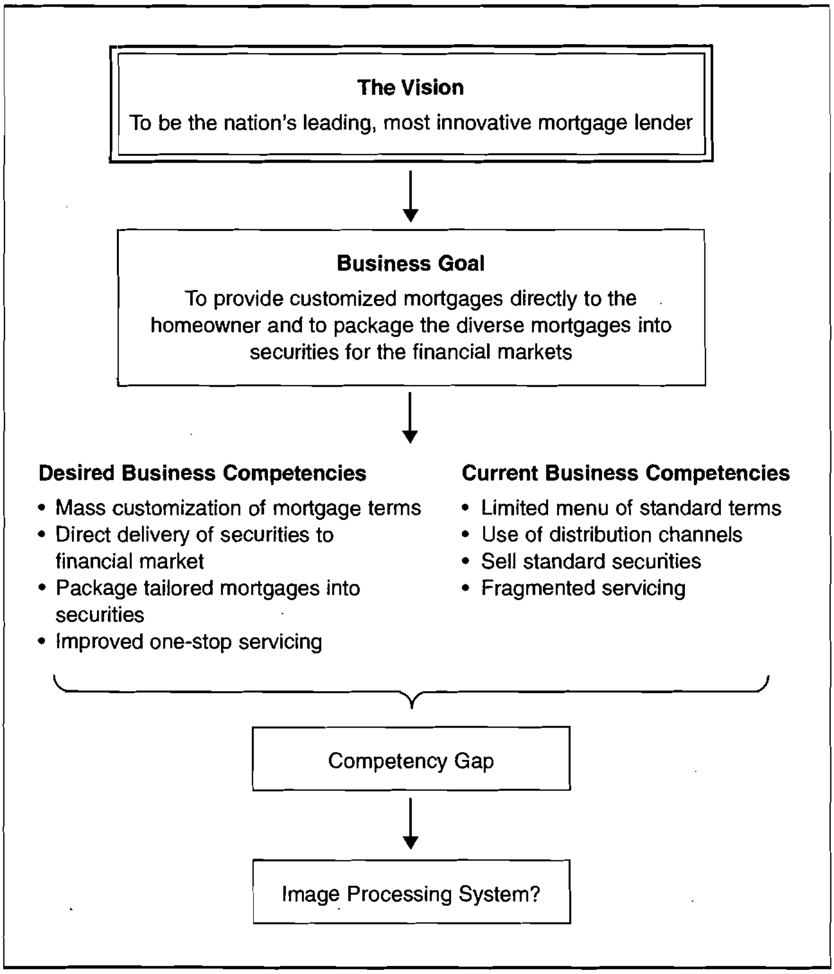

After the larger vision was in place, senior management broke down the vision into the set of desired competencies shown in Figure 14.1. Also shown is the set of current competencies. The gap between the two sets brought to the surface the need to acquire a paperless office competency, which takes more than an investment in a document imaging system. Acquisition of the competency would require investments in other information technology systems and changes in the way work is organized around that technology. The document imaging system, however, was at the center of change.

Although the CIO had originally requested the system as a standalone project, he and other senior management now saw the document imaging system as a platform investment for the other projects. Everyone recognized that the full value of the document imaging system would be obtained only if other investments—training, reworking of office procedures, and so on—were made as well. In addition, senior management reworked the capital request and divided the investment into three stages, with an option to abandon at the end of each stage:

Stage 1: Implement the document imaging system in a few locations for new mortgages.

Stage 2: Use the system for all new data in all offices and redesign workflows.

Stage 3: Implement direct delivery of mortgages to financial markets.

The imaging investment was part of a broader investment program to develop strategic business competencies. Each stage of investment was a package, and a decision to invest would be made only if mortgage market activity reached a critical level.

The Questions

Should NMT invest in the document imaging system? What do the benefits of this investment depend on? How can the imaging project be structured to best integrate it with the overall business strategy?

The Application Frame

As the senior management team reviewed the CIO's budget request, they were bothered by his assumptions that the technology would be successfully implemented. In addition to concerns about an implementation stumble, they viewed the demand for new mortgages as highly uncertain and driven by the level of interest rates and health of the economy. In addition, a proposed change in regulations might narrow NMT's "spreads," the difference between their rev-

Parallels to Other Applications

Infrastructure investments are hard to value because their benefits are elusive, spread across the company, and contingent upon follow-on investments. The real options approach can be used to take a broad look at the alignment of the investment with corporate strategy. For instance, a telecommunications firm that buys the right to use a portion of the radio frequency can realize its value only after building a radio network and offering various services. The decisions to build the network (itself done in stages) and offer services would be contingent on the degree of success of the preceding stages and the attractiveness of the realized market for the services.

enues (mortgage interest rates) and costs (costs of funds and processing expenditures.)

Figure 14.2 shows the decision to invest in the image processing system and subsequent possible events of Stage 1. The technology implementation either is successful and delivers the competency C_S or fails and delivers the competency C_F. The competency outcomes are defined in terms of three performance measures that affect cash flow: market share, fixed costs, and variable costs. A portfolio of mortgage company stock returns tracks the level of mortgage activity. Further analysis showed that overall mortgage activity and NMT's profits were only weakly related, so the use of this proxy introduces some tracking error. But using such a market-based measure results in fewer errors than would be introduced by overmodeling private risk. Total revenue from mortgages is assumed to have a log-normal distribution, and an estimate of volatility, 35%, was obtained from historical data on the mortgage market. A simple binomial representation of mortgage demand is shown in Figure 14.2 for illustrative purposes.

The remaining inputs are as follows. Initial cash flows were calculated from annual revenues of $200 million at current levels of mortgage demand, fixed costs equal to 20% of revenue, and variable costs equal to 70% of revenue. Market share was 10%. Changes in competencies and mortgage demand affect the levels of these inputs to cash flow. It was assumed that the probability of implementation failure in Stage 1 was 10%, the probability of implementation failure in Stage 2 was 20%, and that success of Stage 2 was independent of the results of Stage

Figure 14.2 Fold Out of NMT's Stage 1 Investment
The figure lays out how the information technology investment translates into competencies, and how the cash flows are determined by a combination of market and competency outcomes. Competency outcomes are linked to specific performance outcomes for variable costs, fixed costs, and market share.

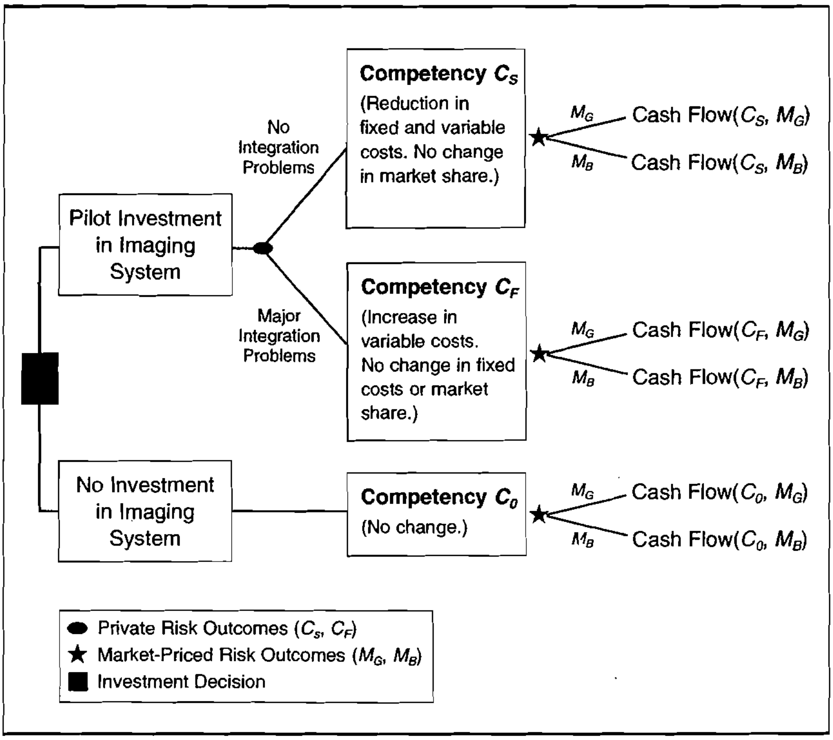

1. Specifically, the failure was most likely to come from the imaging technology in Stage 1; other organizational and coordination difficulties would occur in Stages 2 and 3. The cost of Stage 1 is $500,000, and the cost of Stage 2 is $5 million.

The Results

The real options analysis changed the Standard NPV results dramatically. The value of the project in the CIO's proposal was $500,000 (revised by senior management to a loss of $380,000), and the real options project value was $2.1 million. The optimal Stage 1 investment strategy with Standard NPV is abandon; the optimal Stage 1 investment strategy with real options is proceed with a pilot project. The real options approach to the imaging investment creates more value, yet is more incremental.

Figure 14.3 Investment Risk Profile of the Imaging Project
The risk profile of an investment shows the distribution of possible values at the project's conclusion. For NMT, the one-shot rollout has many final values below $0. The distribution for the value of the three-stage rollout is shifted to the right and has a large chance of a positive outcome.

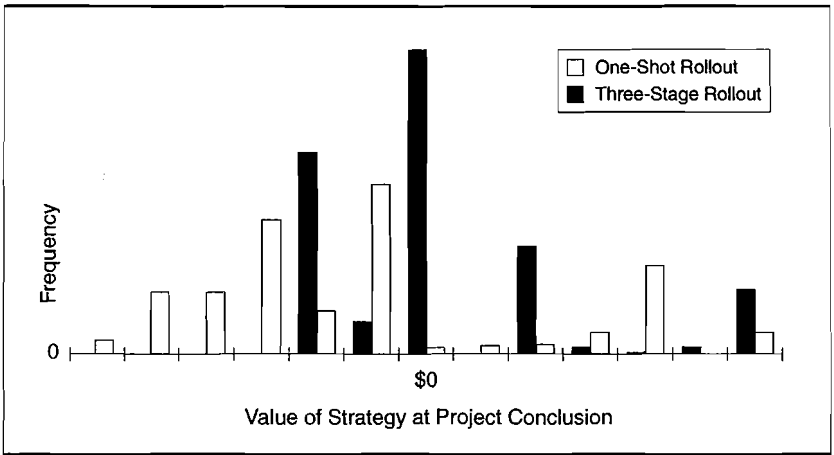

The risk profile of the imaging project is another useful output of the real options analysis and is shown in Figure 14.3. The figure plots the frequency count for the value of the project at its conclusion of the three-stage project described earlier and for a one-shot project in which the staging options are ignored. The results show that the risk profile for the three-stage project has many outcomes with positive value, while the one-shot project has greater exposure to losses. Staging clearly reduces risk. Also, staging made the project more palatable to senior managers who weren't comfortable giving the CIO carte blanche but wanted milestones and contingent approval of future budgets.

The real options analysis highlighted the potential learning of a pilot project, and it was successfully completed. The demand for mortgages fell, however, as interest rates rose, and NMT decided to delay Stage 2 until market conditions improved.

This case study has shown how the real options approach valued an infrastructure investment by a combination of "thinking wide," tracing through the impact of infrastructure on business competencies, and "thinking far," linking the value of the investment to external as well as internal factors. The next case study evaluates the option to delay, the option actually used by NMT.

Chapter **15**

Valuing Vacant Land

The city of Metropolis has been expanding for years, coming out to the peach orchards. Laga Realty, a developer, owned two orchards: one was overgrown long ago, but the other continues to produce a harvest, generating profits of $10,000 per year. It was estimated that, at current values, the properties can be developed immediately for a net profit. Should the property be developed immediately or should development be postponed?

Forty years ago, the valley surrounding Metropolis had been filled with fruit orchards. Now technology companies that need room to expand drive the local economy. Laga Realty owns two 40-acre orchards: one has been overgrown for 10 years and is no longer harvested; the other generates $10,000 per year from its harvest.

Immediate development would cost $2 million, and the developed property would be worth $2.2 million. However, as the technology sector of the economy experiences boom and bust cycles, so does the value of developed property. These fluctuations are well tracked by several regional real estate investment trusts (REITs), portfolios of real estate properties that have shares listed on a stock exchange.[1]

The Questions

Should the properties be developed immediately or should the developer wait? What is the effect of the harvest on the development timing decision?

Parallels to Other Applications

Managers often intuitively use the option to wait, comparing the value of an investment today with the uncertain value of the same investment later. Investing in either period could have positive value. The problem is to choose the strategy with the highest value. This leads to a decision rule to invest now if:

Value of Immediate Investment *less* Cost of Investment
is greater than
Value of Option to Wait

or

Standard NPV *is greater than* Value of Option to Wait

The option to wait has many applications, including waiting to expand a product line, waiting to enter a new market, waiting to destroy an old hotel and build a new one, and so on. Abandonment options can also be viewed as options to wait: abandon today for an immediate salvage value or wait and abandon later for an uncertain salvage value. The common factors across the applications are uncertainty about the underlying asset and irreversibility—the option to wait avoids unpleasant surprises.

The Application Frame

Current development is somewhat attractive, but there is tremendous uncertainty about the value of the developed property. The developer is worried that the market for developed property will fall, making it impossible for him to recoup the development expenditure. In other words, investing today might lead to regret tomorrow.

However, the investment opportunity does not disappear by waiting. There is clearly some value to the option to "wait and see before investing." By investing today, the developer not only gives up the $2 million, but also kills a valuable option. The prudent investment decision must compare the value of developing today to the value of waiting and developing at later points in time. If postponed, however, the cost of development will rise with the local construction inflation index. The decision today requires looking into the uncertain future and then working backward in time.

Figure 15.1 is a simple characterization of how the option to wait is valued. Begin in the far future at time T, a date further in the future than the true time horizon to ensure that the arbitrary starting point does affect the valuations.[2] At time T, decide whether or not to immediately develop the land. Then move back in time to $T - 1$: Should you develop the land immediately or wait and possibly develop at time T? The value of the land at time $T - 1$ includes the value of the option to wait and depends on the current value of developed property. The fold back continues to the present, and the value of the land includes the option to develop it at any future date.

Figure 15.1 Solution Method for the Option to Wait
The solution begins at time T, far in the future, and folds back to today, time t. In each period, the optimal strategy is found by comparing the value of immediate development ($V - C$) with the value of waiting to develop (*option to wait to develop* – M).

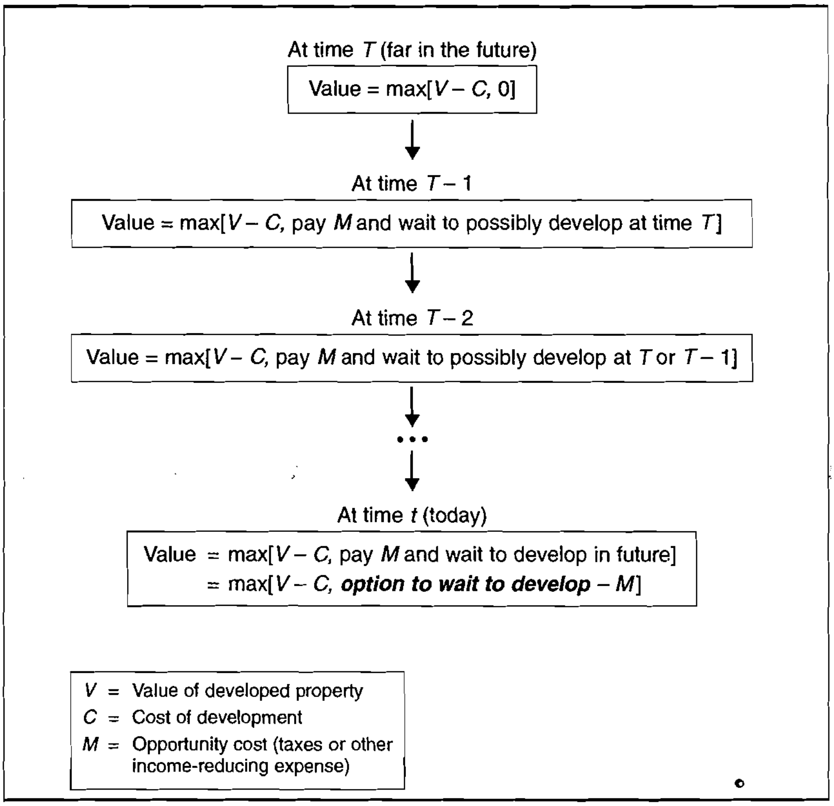

The value of the developed property can be tracked, for instance, by the value of several regional REITs, which are also used to obtain an initial estimate of the value of the developed parcel as 40 times the share price of the REIT. The volatility of the developed property, 50%, was estimated from several regional REITs.

The Results

At current values for developed properties, both properties should be held without development. The value of the postponement option for the idle land is $1.8 million. Hence, immediate investment costs the developer $2 million in development expenditure plus the $1.8 million postponement option that is killed. A second result showed that the developed property value must be nearly twice the initial investment to justify immediate investment.

Figure 15.2 shows the value of the postponement option plotted against possible values of developed property. The optimal exercise of the option is at the point of tangency between the option to wait and the value of immediate development. Higher levels of uncertainty increases the option value and pushes this point of tangency to the right.

As intuition would suggest, the harvest revenues will further delay development because killing the option to wait forgoes the profits from the peach harvest. The effect will be tempered by the carrying costs (property taxes) of the undeveloped land. (This situation is analogous to the effect of operating idle parcels of land in city centers such as parking lots.)

Cash flows accruing to the developed property create the opposite incentives to those accruing to undeveloped land. Because postponing development also postpones the commencement of rental income, rents act as an incentive to invest early, thus lowering the threshold value needed to justify immediate investment.

The developer can increase the value of development by staging the investment, for example, by starting development and then halting it if the economy sours. A staged investment allows for mid-course corrections and decreases the value of waiting, leading to earlier development.

Let's add a few other wrinkles. Suppose development restrictions are introduced, reducing the uncertainty about the value of developed land.[3] What would happen to the peach orchards? The reduction in uncertainty reduces the value of waiting and increases the incentive to exercise the option. In Figure 15.2, the two curved lines representing the

Figure 15.2 Value of the Option to Wait

The option to postpone the development of idle land increases with the value of the developed property. The value of the option to wait to develop largely depends on the volatility of the value of developed property and the opportunity cost of waiting. The optimal time to develop is when the option to wait is just tangent to the value of immediate development.

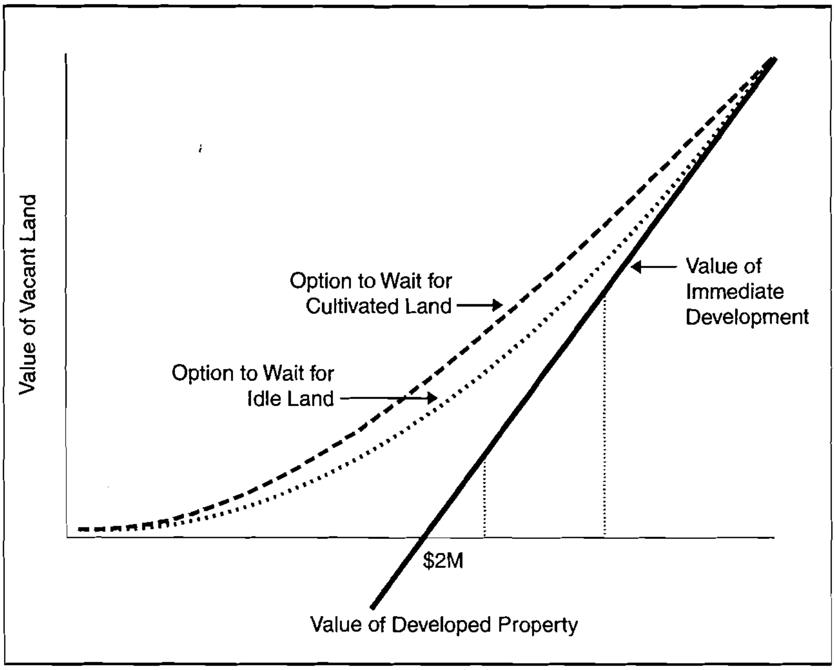

value of the option to wait would shift down, and the points of tangency would move to the left. This change alone might be enough to trigger immediate development. The value of waiting is driven by uncertainty. Other realistic wrinkles include rental income that fluctuates with the value of developed property—which tends to accelerate investment when the value is high and delay investment when the value is low—and the uncertain interest rate on construction loans. This is an additional, independent source of uncertainty and will increase the value of the option to wait, shifting up the curves in Figure 15.2, and hence raising the trigger point of development.

This case study has shown how to value the option to wait, an option that can be "killed" only once. In contrast, the value of the option to switch, the subject of the next case study, depends on its reversibility, the option to switch back.

Chapter **16**

Buying Flexibility

MidAmerica Manufacturing is choosing between three industrial boilers to generate steam. The first boiler burns natural gas, the second burns No. 2 fuel oil, and the third can switch between the two inputs. The price of the third boiler is $2,000 more than the first boiler and $5,000 more than the second, a premium for the built-in flexibility. Which boiler should MidAmerica buy?

To compare the industrial boilers, MidAmerica Manufacturing put the fuel costs on an efficiency-adjusted basis. An efficiency-adjusted price is the spot price multiplied by a factor that accounts for the thermal efficiency of the boiler. The efficiency-adjusted prices quoted below are for the particular boilers in question; different models and makes will have adjustments of a different magnitude.

Here are some more details on the boilers MidAmerica is evaluating:

- *Natural gas.* This boiler costs $63,500 and its efficiency-adjusted price of fuel is 1469 P_{GAS}
- *No. 2 fuel oil.* The cost of this boiler is $66,600, and its efficiency-adjusted price of fuel is 1408 P_{OIL}
- *Dual-fuel.* This boiler costs $68,700. It has the same thermal efficiencies as the single-fuel boilers

MidAmerica also gathered historical data on fuel prices. The recent history of the ratio of the efficiency-adjusted prices for the boilers under evaluation is given in Figure 16.1. When the price ratio equals 1.04 (1469/1408), the efficiency-adjusted prices are equal. The figure shows that the price ratio crosses this level quite frequently, and it

Figure 16.1 Recent History of the Price Ratio

The price of natural gas equals the price of No. 2 fuel oil, on an efficiency-adjusted basis, when the price ratio equals 1.04. The option to switch fuel inputs is more valuable if the efficiency-adjusted price ratio frequently crosses the point of equality.

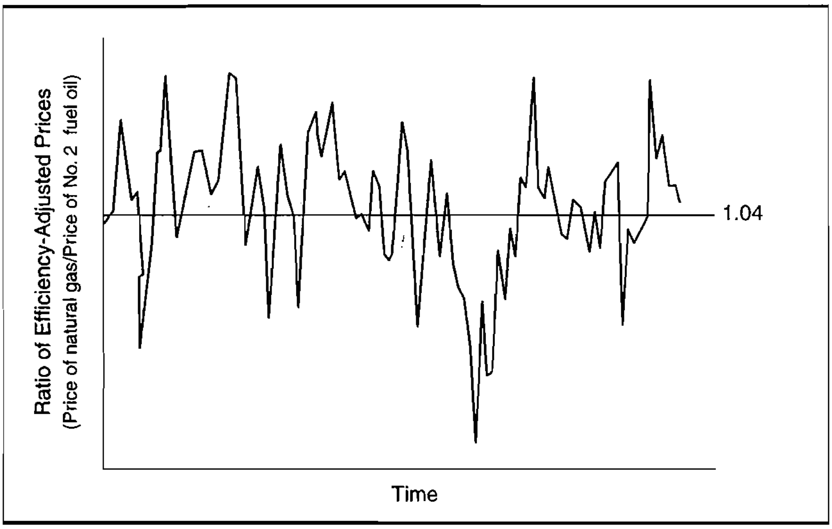

appears that the ratio is mean-reverting. There is an economic story behind this feature. Because the two fuels are substitutes for each other, the efficiency-adjusted price of gas will tend to converge to the efficiency-adjusted price of oil in the long run. The long-run mean is at the point where the cost of fuel is the same for the marginal purchasers of oil and gas.

After examining this price history, MidAmerica wondered how often they would actually switch fuels. If it cost nothing to switch, the cost-minimizing decision rule is switch to the cheaper fuel at any time. With switching costs (S), the cost-minimizing strategy depends on the current fuel being used. If gas is currently being used, there is a cost to switch to oil fuel. If gas is currently being used, the cost for the next month is:

$$\text{Fuel cost in next month} = \min(1.04P_{GAS}, P_{OIL} + S)$$

Similarly, if fuel oil is currently being used the cost for the next month is:

$$\text{Fuel cost in next month} = \min(1.04P_{GAS} + S, P_{OIL})$$

The subscripts on the cost C indicate which fuel is currently being used.

Parallels to Other Applications

Managers are often faced with decisions requiring the valuation of flexibility. Flexibility can be purchased through special features in capital equipment, or it can be obtained by investing in training, routines, or flexible contracts. Flexibility in capital equipment is common and includes flexible manufacturing systems, options to change the product mix in oil refineries, and the opportunity to temporarily shut down and restart equipment. Flexibility also includes options to switch production across locations, depending on local labor conditions, demand, and currency fluctuations. Some contracts contain flexibility, such as a contract to deliver natural gas or electricity each month, whichever is cheaper.

In many cases, the need for flexibility is largely predictable, but with some uncertainty. For example, some electricity-generating equipment is used just during hours of peak demand, certain times of day, and certain seasons. It is purchased for the predictable need for flexibility, but the exact price at which each hour of generated electricity will be sold is still uncertain, and in fact, the realized price may be so low that no power will be generated. Predictable flexibility can be handled by a discounted cash flow analysis, but the uncertain component requires a real options analysis.

The Questions

What is the value of the flexibility of the dual-fuel boiler? How should the flexible boiler be operated? What is the role of the switching costs?

The Application Frame

The method to value the dual-fuel boiler is very similar to that used to value the option to wait in Chapter 15 and is shown in Figure 16.2. Valuation starts at time *T*, at the end of the physical life of the boiler. All three boilers were expected to have a 10-year life. The currently lower-priced fuel, gas, is selected as the initial fuel; if the time horizon is sufficiently long (if *T* is far enough away from the present), the valuation results will not be sensitive to the initial fuel choice.[1] The initial monthly profits are calculated at time *T* as revenues,

Figure 16.2 Solution Method for the Option to Switch Fuels
As the flow of this figure shows, the option to switch has the feature that the options available and the cost of switching depend on the fuel currently being used and the fuel expected to be used in the next period. Dynamic programming folds back the possible price paths of two fuels, the expected fuel choices and the current fuel choice. In this illustration, it is assumed that gas is used at time T in the far future. The notation has been simplified and omits discounting and expectations operators.

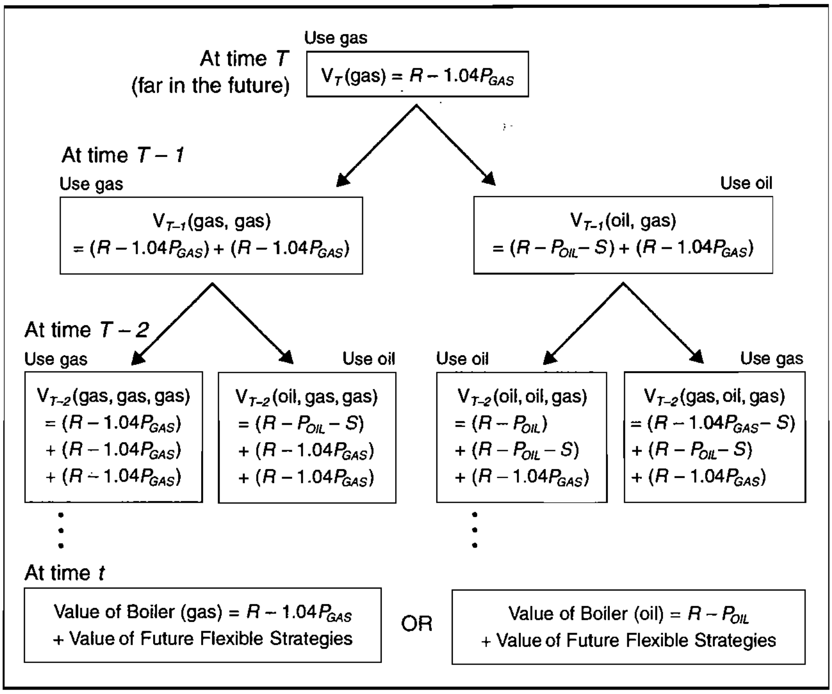

R, minus the fuel cost, $1.04P_{GAS}$ or P_{OIL}. The next month, $T - 1$, the optimal fuel is chosen based on current fuel prices and, because of switching costs, is based on which fuel was used at T.

In the same manner, the valuation continues to fold back. At $T - 2$, four values of the dual-fuel boiler are possible, based on the optimal fuel (given current prices) and the strategies that would be optimal in the subsequent periods. Because these strategies include switching, the value of the boiler obtained by folding back includes the value of flexibility. As the figure shows, the current value of the boiler depends on the current fuel choice and future strategies. This valuation model can be solved using a stochastic dynamic programming solution method.

The application has two sources of uncertainty, the price of oil and the price of gas, but the solution is easier if the analysis is framed in terms of a single source of uncertainty, the ratio of efficiency-adjusted fuel prices. The contingent operating decisions depend on this single source of uncertainty. Let Q denote the ratio of the fuel prices, $Q = 1.04\ P_{GAS} / P_{OIL}$. For example, the cost function, given that oil was used last month, is:

$$C_{OIL} = \min(1.04Q + S, 1)$$

The decision rule does not change, but now we need only a single estimate of volatility, that of the ratio of the efficiency-adjusted fuel prices.

The recent history of the price ratio shown in Figure 16.1 suggests that the price ratio is mean-reverting, and this assumption was made in estimating the volatility. A regression analysis was done to estimate how quickly the price ratio moves to the mean after accounting for the expected price trend. Volatility was calculated as the deviation from that trend. Thus, the volatility estimate depends on the assumption of mean-reversion and the empirical estimates of the speed of mean-reversion.[2] The estimated volatility is used to roll out a tree showing the range of possible values for the price ratio, and then the dynamic program described in Figure 16.2 was used to fold back the valuation.

Mean-reverting prices can arise from a mean-reverting convenience yield or from mean-reversion in the risk premium, the additional returns required by investors for bearing risk. A mean-reverting convenience yield requires an adjustment to the tracking portfolio; a mean-reverting risk premium is already captured in the fuel prices. In some cases, the source of mean-reversion can be identified with empirical tests of the price data; in other cases, the source is imposed by assumption.

The Results

MidAmerica is willing to pay the extra cost of the dual-fuel boiler if its flexibility adds value. Table 16.1 summarizes the valuation results for the three boilers. For consistency, the value of a single-fuel boiler was obtained by setting the switching cost S to such a large number that switching was never the optimal strategy. (The real options valuation method gives the same answer as discounted cash flow when the option is "hardwired" out of the valuation.)

As Table 16.1 shows, MidAmerica actually did two separate calculations. The first was based on 10 years of fuel prices ending in 1977, and the second on 10 years of fuel prices ending in 1997. MidAmerica estimates showed that the volatility of the price ratio, after accounting for mean-reversion, was higher in the 1970s than in the 1990s.

Table 16.1 The Declining Value of Fuel-Switching Flexibility

In 1977, the dual-fuel boiler was the optimal investment, after accounting for all fuel-switching options. During the subsequent 20 years many firms and utilities purchased dual-fuel boilers, reducing the volatility of relative fuel prices. In this example, the natural gas boiler emerged as the optimal investment in 1997.

Type of Boiler	Purchase Price (constant dollars)	Value to MidAmerica at Current Fuel Prices	NPV
1977 Results			
Natural gas	$63,500	$65,000	$2,500
No. 2 fuel oil	$66,600	$62,000	*approx.* $0
Dual-fuel	$68,700	$98,000	$26,000
1997 Results			
Natural gas	$63,500	$68,000	$3,000
No. 2 fuel oil	$66,600	$62,000	*approx.* $0
Dual-fuel	$68,700	$69,500	$800

The dual-fuel boiler was the best choice in 1977. The value of its flexibility to MidAmerica exceeded the purchase price premium. By 1997, however, the value of the dual-fuel boiler flexibility had decreased dramatically, and the single-fuel natural gas boiler was the best choice. Why?

During the 20-year interval, many industrial manufacturers and utilities added dual-fuel boilers, increasing the speed of mean-reversion and reducing the volatility of the price ratio. As these companies switched between fuels, decreasing demand for one and raising demand for the other, their actions caused the spot price of gas to rapidly converge toward the spot price of fuel oil. The long-run forces reflecting the efficiency of the least-efficient dual-fuel boiler in the market and its switching costs. In recent years financial contracts have been introduced that synthetically create a dual-fuel boiler at a lower switching cost than many boilers in place, further reducing the volatility of relative fuel prices.

This case study has shown how to value flexibility. Electric utilities face the same flexibility choices for generation capacity, but at a scale 10 times larger than shown here. When important capital equipment has built-in input flexibility, it can actually change the dynamics of commodity markets. The next case study illustrates an additional force at work: Financial market contracts can also be used to obtain flexibility, further disciplining the optimal operating decision.

Chapter 17

Combining Real and Financial Flexibility

New Jersey Refining had been watching the new financial products available on the crack spread, the gross profit margin for oil refiners. Currently it operated its plant using a dynamic programming algorithm, temporarily shutting down whenever the crack spread was too low, after accounting for shutdown and start-up costs. New Jersey Refining thought that its competitors were using the new financial contracts to avoid temporary shutdowns and wondered whether it should adopt this approach as well.

New Jersey Refining was a typical oil refiner, using a process called "cracking" to break down crude oil into a mix of refined products, including gasoline, heating oil, kerosene, jet fuels, solvents, and so on.[1] The crack ratio summarizes the refining process. New Jersey Refining's ratio is 5:3:2, meaning that five barrels of crude oil yield three barrels of gasoline and two barrels of heating oil. Crude oil is about 85% of the total operating cost, while gasoline and heating oil are about 80% of New Jersey Refining's revenues, so the crack ratio is an approximate indicator of its gross margin.

There had been times in the past when New Jersey Refining had temporarily shut down, incurring substantial costs. These included the fuel costs of warming up the refinery furnace and distillation tower and from "wear and tear" on the equipment. In addition, while shut down, New Jersey had ongoing maintenance and safety costs. Despite

This chapter is based on research done by Nalin Kulatilaka and Changku Yi at Boston University.

these costs, sometimes the crack spread was so low that shutdown was the best strategy.

Managers at New Jersey Refining had noticed that futures and options contracts were trading on several exchanges based on the 5:3:2 crack ratio and wondered if these contracts could be used to smooth cash flow. The contracts had a constructed underlying asset, the crack spread, which is a weighted average of the appropriate spot prices. Typically there were two months between the date a refiner purchased the crude oil and the date it sold its output, and the constructed underlying asset included this timing difference in spot market prices.

In addition, a number of investment bankers and commodity brokers had approached New Jersey Refining about creating options and futures contracts tailored to its own gross profit margin. The crack spread in the traded contracts imperfectly tracked New Jersey's gross margin because New Jersey's sale of output lagged its oil purchases by only one month. Although a tailored contract would eliminate the basis risk of a traded contract, it would also have a higher cost, because it would include a fee to the investment bank for designing and providing the contract.

The Questions

How can financial market information help New Jersey Refining determine its optimal operating strategy, including the optimal rule for temporary shutdown? Can New Jersey Refining use the traded crack spread contracts to smooth cash flows, or should it buy contracts from the investment bankers?

The Application Frame

Under its current operating strategy, based on a dynamic programming model, New Jersey Refining shuts down when its gross margin falls below a critical value and restarts operations when the gross margin rises above a second critical value. Using a real options analysis, the refinery uses financial market information—the convenience yield and the volatility of crack spreads—to discipline its operating strategy, changing the critical values for shutdown and restarting operations. The real options analysis does not require any forecasts of future crack spreads, another difference with standard dynamic programming. The application frame for the shutdown option is similar to the frame for the switching option discussed in Chapter 16, because a shutdown op-

Parallels to Other Applications

What if a new competitor came along with lower costs of delivering product? In some product markets, financial contracts are available that mimic operating options. The presence of financial options expands the opportunities to acquire flexibility and undertake risk management. The list of financial contracts available to track gross profit margin is growing, and it includes:

- *The crack spread in oil refining,* the difference between the weighted average of prices for refinery outputs and the price of crude oil
- *The crush spread in soybean processing,* the difference between the price of crushed soybean oil and raw soybeans
- *The spark spread in electricity generation,* the difference between the price of electricity and the price of gas adjusted for a standard thermal efficiency

tion is really an option to switch between two modes, refining and shutdown.

The shutdown decision rule depends on the current refinery status (shutdown or operating), the shutdown cost, d; the start-up cost, s; the maintenance costs during shutdown, f; and the crack spread, k. For example, if the refinery is currently operating, it should continue operating for the next period only if:

Value when in operation + k
+ Value of option to shut down in the future
is greater than or equal to
Value when shutdown $- f - d$
+ Value of option to start up in the future

Each side of the inequality captures current cash flow (k or $- f - d$) and the future value of the refinery under the operating mode. The decision rule for a refinery that is currently shut down is similar. The real options analysis is required to correctly value the options.

The value of the shutdown and start-up options depends on the cost of each action (*d* and *s*), the current crack spread (*k*), the volatility of the crack spread, and the convenience yield on the crack spread. The parameters *s*, *d*, and *f* depend on the technology, organizational structure, and contractual arrangements of New Jersey Refining. The crack spread parameters are obtained from spot and futures market data.

The Results

Figure 17.1 shows the optimal shutdown and start-up rules, the output of a real options analysis. The figure shows an area of inertia, known as a hysteresis band. The band ranges between $2.75 per-barrel crack spread (the critical value for shutdown) and $4.25 per-barrel crack spread (the critical value for start-up). Inside the hysteresis band the refinery has negative cash flow; the refinery will continue operating, although it is incurring losses and will remain closed although it could be earning profits. The refinery incurs these short-run losses because the costs of switching operating modes make the strategy

Figure 17.1 Critical Values for the Refinery's Shutdown Option
The refinery's operating strategy for a particular path of the crack spread is shown. When the crack spread falls below the critical value for shutdown, the refinery pays $*d* to temporarily close and then $*f* per month while idle. The spread must exceed the critical value for start-up before the refinery is willing to pay $*s* to reopen. In this band of inertia or hysteresis the refinery remains in operation although it would generate negative cash flow and stays idle although operations would generate positive cash flow.

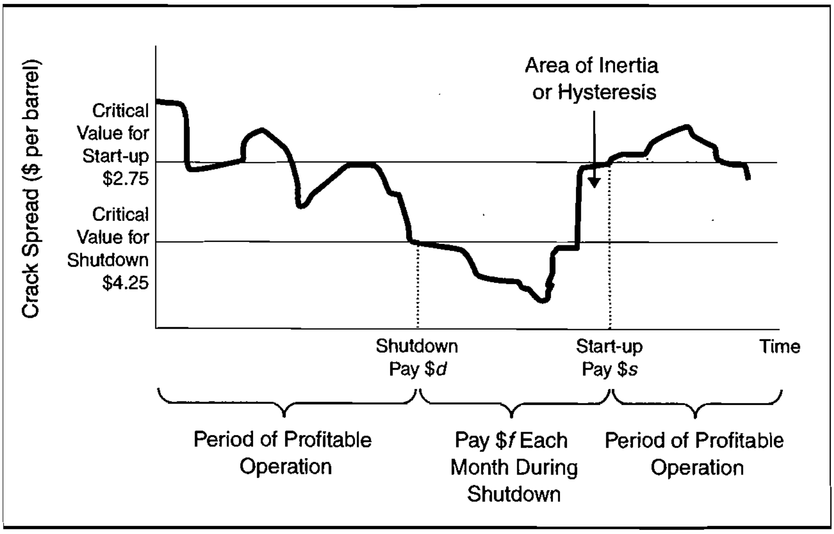

optimal in the long run. Note that the critical values of crack spread at which the refinery should be shut down and started up do not depend on a private forecast of crack spreads; the critical values depend on objective information, the cost parameters and market information about the crack spread.[2]

In addition to the critical values, the real options analysis produces another output, the risk profile of future cash flows that would result from using these operating rules. This profile may not be acceptable to New Jersey Refining. For example, senior management may feel that the company is too exposed to oil price fluctuations and would prefer to smooth cash flows. Financial hedges based on the crack spread can then be used to modify this risk profile. Because these financial contracts are based on the cost structure of the marginal producer in the industry, they benchmark New Jersey's performance over a range of outcomes and provide a reference point on the cost of modifying the risk profile.

Does the existence of financial contracts change the actual operating strategy of the refinery? It may if the market-priced insurance against crack spread fluctuations is cheaper than the cost of shutting down the refinery. In this case, New Jersey Refining would never shut down. Instead it would purchase insurance to compensate for the losses incurred by continued operations below the critical shutdown level of the crack spread. Cheap insurance can change the operating strategy. Note that even if New Jersey Refining does not buy the crack spread hedge, it should use the price information to improve its operating strategy.

Note that for a new entrant, the financial options become an entry barrier—don't enter unless your physical operating options are more valuable than the financial options. Industry capacity expansion will take place through an increase in superefficient physical capacity and the expansion of virtual refineries through contracts.

Financial contracts are useful for shedding commodity price risk but may introduce other risks, including:

- *Counterparty risk.* Generally, tailored contracts are available through investment banks, which are the only source of pricing information as well. Credit risk can also be a concern when other companies are the counterparties.
- *Legal risk.* Specialized unproven contracts may have ambiguous terms of sale or undisclosed risks—matters that sometimes incur legal costs to resolve.

- *Basis risk.* The refinery's gross margin can differ in composition from the market-traded crack spread.
- *Liquidity and execution risk.* Some contracts are illiquid and/or have high transaction costs. These costs must be paid whether the owner of the physical asset or the bank runs the tracking portfolio (although the bank is likely to have cheaper transaction costs).

This case study has demonstrated the power of using the synthetic flexibility of financial assets in the operation of a real asset, a strategy that has important implications for competition in the marketplace. The firms in the previous case studies have been "pricetakers"; the next case study shows how the applications involving "pricemakers" are treated in the real options approach.

Chapter **18**

Investing to Preempt Competitors

GoodTech was nearing the end of product development for its latest electronic product, a personal digital assistant (PDA). GoodTech usually introduced its new products once a year in October, just before the big Christmas selling season. The size of the market for PDAs was very uncertain, and the PDA product introduction would be very expensive. Senior management at GoodTech was divided over which strategy to take: hold back the product one year and wait to see how this season developed or introduce the product this year, using it to set industry standards on PDA features.

Everyone knew that the market for PDAs was very uncertain. Other companies had introduced products, but these had been more consumer novelties than products ready for the large-scale consumer market that GoodTech usually sold into. This season it might be different.

GoodTech's product management team has a dilemma. If GoodTech launches its PDA this year, other competitors will be discouraged from entering the market; if competitors did enter, GoodTech would be well positioned against them. If the market for PDAs turns out to be weak, however, GoodTech will not be able to recoup its investment, even with a strong market position. If GoodTech waits, it will not have an advantage over potential competitors, but it will be saved from the regret of making an investment in a weak market.

Parallels to Other Applications

GoodTech faces a trade-off common to many new markets: Uncertainty increases the value of waiting, but early investment may actually change the payoffs to investment. Certainly many companies believe that this sort of trade-off is at the heart of their market. Drug companies believe that early entry locks up market share. Technology companies believe early entry sets standards. Internet companies believe early entry is needed to be at the center of an increasing returns spiral. There are also numerous examples of this type of thinking in recent corporate spending. Western consumer goods companies have rapidly acquired most brand names in Eastern Europe, though uncertainty over these markets is still enormous. Pharmaceutical firms routinely purchase the rights to new compounds developed by biotech companies well before their commercial viability is established. Cable companies have been committing large amounts to the media industry at a stage when uncertainty over future regulation, technology, and consumer tastes is still large. This case study shows what must be behind the "gold rush" mentality to justify giving up the option to wait. As strategy guru Gary Hamel summarizes: "We live in a discontinuous world. . . . Who will capture the new wealth? . . . In an increasingly nonlinear world, only nonlinear strategies will create substantial new wealth."[1]

The Question

Should GoodTech retain the waiting-to-invest option or should they take the plunge in order to gain competitive advantage?

The Application Frame

Framing an application in which today's investment changes the rules of the market or tilts a market outcome toward one competitor requires blending insights from game-theoretic models of strategic market interactions with real options models of investment in uncertain markets.[2] This is quite hard to do in a way that provides a general way of thinking about the trade-offs. Often the results are specific to features of the model used.

This case study has a second difference from the others in this book. In the other case studies, the firms use strategic investments to shelter themselves from adverse outcomes while capitalizing on the good outcomes. Here the firm also uses strategic investments to change the market structure itself, enhancing upside benefits. Increasing returns investments are currently being generated in many similar new markets.

For GoodTech, product introduction expenses include advertising, branding, and development of distribution channels. GoodTech's early sales would attract other companies to write programs, design accessories, and create plug-in components, such as talking address book software, cellular phone adapters, and personalized carrying cases for the basic GoodTech PDA. This would establish the GoodTech PDA as the industry standard.[3]

GoodTech's preemptive investment opportunity can be viewed as the purchase of a "strategic growth option." This option not only allows GoodTech to capture the upside of a potentially good market but also has the strategic effect of influencing competitor behavior. One way to model the benefits of the growth option is as the capture of a greater market share. Investment creates capabilities to possibly take advantage of future growth opportunities. The competitive value of these capabilities may deter potential new entrants and restrict the output of those who do enter. This stylization captures the benefits and risks of a preemptive investment in a simple, tangible way, and is entirely consistent with the richer story of the real world.

Figure 18.1 summarizes how the strategy choice can be modeled as the choice between two options—the option to wait and the growth option—that are triggered by the same underlying asset, the value of the PDA market. The payoff to the growth option is very sensitive to the underlying asset because early entry results in larger market share and higher profits. The option to wait has a flatter payoff and allows GoodTech to avoid investment altogether if the value of the PDA market is unattractive.

Now turn to the uncertainty. The option to enter is triggered by the value of a newly emerging market. The players in a new market are often large, diversified companies (the ones that can more easily finance platform investments). What data might we draw on to understand the volatility of a new market? Much of the early industry experience is buried in the data of large companies, and traded companies that would serve as "pure plays" (companies facing the risks of just one industry) are unlikely to be available. The traded securities that

Figure 18.1 Invest Immediately or Wait?

In this example the payoff to immediate entry has a steeper slope than the payoff to the option to wait because the firm influences the market structure, creating greater up side exposure. By investing early the firm buys a growth option with a very steep payoff.

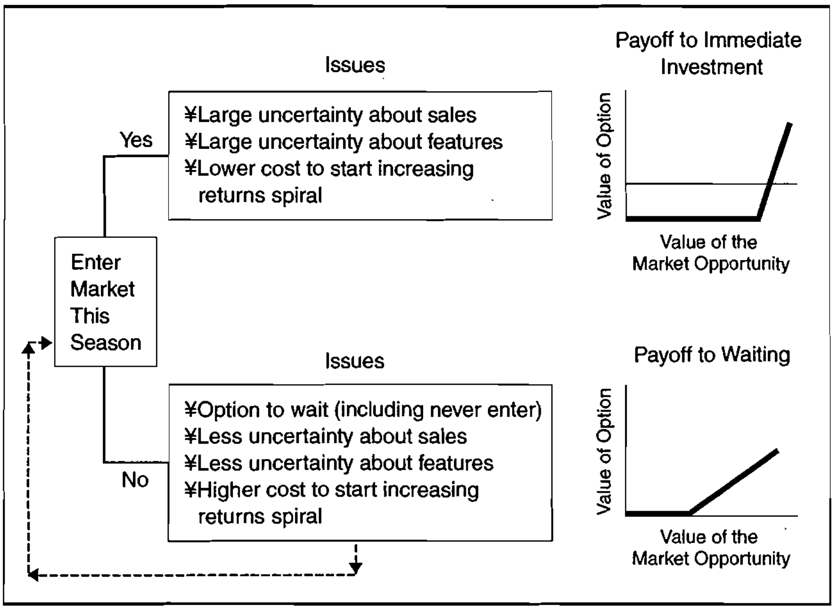

provide the best tracking are likely to be recent IPOs selling into new markets, but significant tracking error will remain.

The Results

Figure 18.2 summarizes an important result. Figure 18.2(a) shows the waiting and growth options for a market with low uncertainty, and Figure 18.2(b) shows the two options for a market with high uncertainty. The investment threshold is defined by the intersection of the two curves. The firm should commit the investment and acquire the strategic growth option only if the expected value of the market is higher than the threshold.

The best strategy will be determined by the relative values of three variables: the cost to enter (lower now, higher later); the impact of the investment on the market structure, which determines the slope of the payoff function (steeper now, flatter later); and the level of uncertainty (higher now, flatter later). In many applications, the growth option is more sensitive to uncertainty than the option to wait, making preemp-

Figure 18.2 Greater Uncertainty Makes Entry More Likely

The firm enters the market at the investment threshold, when the value of immediate investment equals the value of waiting. (b) shows that as volatility increases, the increase in the value of the preemptive growth option (and hence the value of immediate investment) is larger than the increase in the option to wait. An increase in uncertainty moves the investment threshold to the left, making entry more likely.

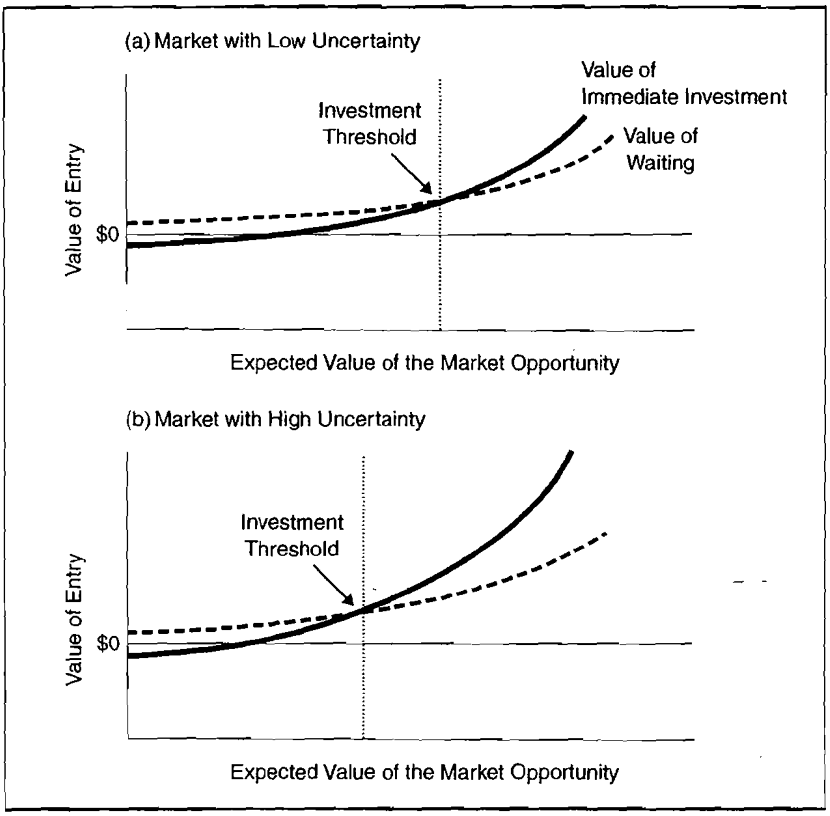

tion the best strategy in markets with high levels of uncertainty.[4] This is shown in Figure 18.2(b). (A note of caution: Think through whether this result can be generalized to your application or if your market has some strategic interaction not captured in the model just described.)

In some industries, winning the time-to-market game changes the industry rules, and the preemptive model of this chapter will capture the right trade-offs. James Morgan, CEO of Applied Materials, recounts one success story: "In 1997 we lost money for two quarters so that we could introduce our . . . architecture six months ahead of schedule. But we broke even for the year, and now we own the architecture of choice."[5]

In other industries, time-to-market pressures are different from the "price-making" investments of this chapter. If market arrival time does not confer any special market power, "rushing to the market" consumes investment resources but does not increase profits to the firm. It just takes additional money to get to the starting gate.

In some cases there is trade-off between preemptive investment and the option to wait, but at high levels of uncertainty preemptive investment is likely to be the dominant strategy. The next case study looks at another kind of opportunity to shape outcomes: setting the terms of technology license contracts.

Chapter **19**

Writing a License

Bonzyme is a "virtual" biotech firm, one constructed by licensing arrangements. It has contracts with academic researchers, supplying part of their funding in exchange for rights to discoveries. It contracts with firms specializing in running clinical trials to get the drugs through testing. In the late stages of clinical tests, Bonzyme gets its payoff: BG Pharma, a large pharmaceutical company, may acquire a license from Bonzyme giving BG Pharma the rights to market the drug. Bonzyme must negotiate a license knowing that the terms of the contract give options to BG Pharma and that they may even influence BG Pharma's decisions regarding the launch and marketing of the drug.

Most managers recognize that licenses are hard to value with traditional valuation tools because future sales are uncertain. One common solution is to establish payments as royalty rate, sharing the risk of the uncertain market between the two parties of the license. If the license changes how the real options are exercised, the license valuation can be complex. Licenses also have explicit options written into them; the license may contain options regarding an asset that itself has many implicit real options.

Bonzyme owns a drug that contains the option to launch and a series of options to abandon. When it grants BG Pharma a license, Bonzyme may transfer the options, it may influence the exercise of these options with the payment terms, or it may effectively destroy these options with onerous cancellation penalties. To negotiate the license, Bonzyme must know how to value the alternative contract terms. A real options model must be used to account for how the terms of the license affect BG Pharma's optimal exercise strategy for the drug's options.

Parallels to Other Applications

In the new economy, you don't have to buy a firm to get its technology. Some firms have competitive capabilities in the production of intellectual capital, and other, often larger firms, have complementary capabilities in turning the intellectual capital into a product. In the biotech industry, drug companies offer financial assistance during the lengthy development process, distribution channels, and large marketing budgets. In the semiconductor industry, some companies specialize in aggregating and integrating blocks of code from others and then combining the blocks into "a system on a chip." Many high-tech firms started with a technology developed at a university, including Open Market, Silicon Graphics, Yahoo!, and Lycos.

The Questions

How do the terms of a license affect BG Pharma's decisions to launch the drug? How can the license, with its contractual options, be valued?

The Application Frame

Although actual licensing terms can be more detailed and complex, the basic arguments for how to structure a license are captured by tracing through the effects of two different forms of payment, an annuity of fixed payments over time that are triggered by the market launch and a royalty rate (royalties are computed as a percentage of sales). In practice, a license may have both terms. To keep our arguments clear, we model BG Pharma as having a single market launch option.

Let's consider the effect of each form of payment on BG Pharma's launch decision. First, to consider the most simple case, think of the annuity as a perpetuity, a schedule of fixed payments lasting forever. Let the value of this perpetuity be equal to L. The value of the market opportunity to BG Pharma has fallen by the amount L; its option to launch has an underlying asset equal to the value of the market opportunity minus L. Equivalently, the cost of the launch to BG Pharma is increased by L. This is shown in Figure 19.1.

Figure 19.1 A License Changes the Payoff to Market Launch
The option to launch is modified by a license granting Bonzyme a perpetual annuity of fixed payments upon launch. The present value of the payments equals *L*, and can be treated as a cost of launch in the option analysis.

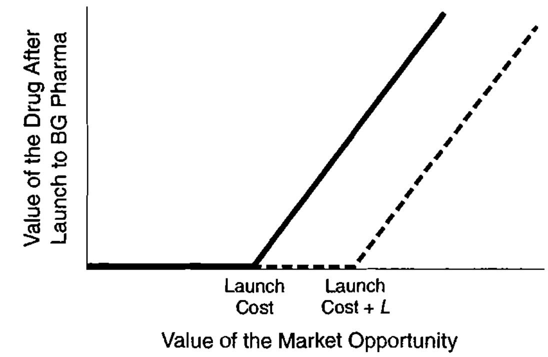

If BG Pharma had developed the drug in-house (or if the license required the fixed payments regardless of the launch), the development costs would be considered sunk costs and would not affect the launch decision. This is the payoff with the solid line in Figure 19.1. When the fixed payments to Bonzyme are included in the option valuation, the launch occurs only at a higher value of the market, shown by the dotted line. The fixed payment terms will delay the launch, all else being equal.

To show the effect of a fixed payment in its most simple form, we assumed that the payments were made in perpetuity. In practice, they are likely to be made for a limited time. As time goes by, less is owed to Bonzyme, and the difference between the two payoff functions in Figure 19.1 will shrink. BG Pharma holds an American-style option to launch; thus, the effect of the fixed payments on the option solution method is complex. (See Figure 9.3.)

Now consider the case when the payment is a royalty on sales. At each date a launch decision is evaluated, the present value of expected royalties, L^*, can be computed from the sales forecast. A sales forecast will be needed because the launch decision will depend on the value of the market opportunity. In most firms, there is not a constant relationship between expected sales and value of the market opportunity—expected sales must be computed outside the real options framework.

As before, the expected payment can be added to the launch costs for valuation of the option to launch.

The effect of the license terms on the option to launch will depend on the specifics of the application, including the magnitude of the payment terms, the remaining term of the license, the time to maturity of the option to launch, whether the option to launch is American or European, and the level of market uncertainty. These factors must be taken into account in the design and evaluation of the contract. The dynamic programming solution method can be used to value the license. The foldback will simultaneously value the real options, the effect of the license terms on these options, and the value of the license (including any options specified in the license contract).

In Chapter 13, we suggested that the value of the market opportunity for a drug could be tracked by a stock index of pharmaceutical companies. This index does not represent the present value of life-cycle profits of the drug but more generally the value of being in the business. It is appropriate for tracking the underlying asset of the option to launch. (It may not be appropriate, however, for tracking the value of the remaining market opportunity when the drug goes off-patent.)

The Results

Three broad-ranging results emerge from this case. First, license valuation requires that the contractual terms be simultaneously valued with the real options contained in the asset. The real options approach can consistently value the internal real options of BG Pharma and the value of the license to Bonzyme.

Second, the value of a license depends on the magnitude of uncertainty in the product market. How BG Pharma exercises its real options will influence the value of the license, and exercise decisions depend on uncertainty. Even when the license is a one-time up-front fee, license terms will affect the real options, and the effect of the modified real options must be included in the license value.

Third, terms of the license affect the way the licensee, in this case, BG Pharma, develops the market for the drug. For example, a fixed fee per unit sold would create an incentive for BG Pharma to discontinue selling the drug if prices fall sufficiently low, thus truncating the stream of license fees received by Bonzyme. BG Pharma can credibly threaten to discontinue sales, and it might demand new negotiations to lower the licensing fee. It might be possible to prevent this outcome by imposing a penalty for renegotiations in the original contract, but it is not possible to completely contract against all perverse outcomes. Depending on the terms, the license may be subject to the "holdup problem," the threat of negotiations for new terms at a particularly disadvantageous moment. The holdup problem occurs when one party holds all the decision rights. The interplay between issues of incomplete contracting and valuation is crucial in the design of licensing contracts and is an area of current academic research.

This chapter has shown how the structure of license terms can affect the value of the business opportunity. Licenses are an increasingly important way of transferring and combining assets and an area where real options converges with "financial engineering"—the design of contingent contracts.