Real options are a complement to, not a substitute for, discounted cash flow analysis. To pick the best growth projects, managers need to use the two methods in tandem.

Making Real Options Really Work

by Alexander B. van Putten and Ian C. MacMillan

FOR ALL THEIR THEORETICAL attractiveness as a way to value growth projects, real options have had a difficult time catching on with managers. CFOs tell us that real options overestimate the value of uncertain projects, encouraging companies to overinvest in them. In the worst case, they grant excessively ambitious managers a license to gamble with shareholders' money.

This reluctance stems at least in part from a suspicion that it's risky to apply valuation tools that have been developed for well-defined financial options to complex business projects. The tools oblige managers to make many simplifying assumptions and therefore, the thinking goes, they cannot fully capture a proposal's multifaceted risks and opportunities. These concerns are legitimate, but we believe that abandoning real options as a valuation model is just as bad. Companies that rely on dis-

counted cash flow analysis for valuing their projects fall inevitably into the trap of underestimating the value of their projects and consequently don't invest enough in uncertain but highly promising opportunities.

How can managers escape this dilemma? In exploring their reservations about real-option analysis as a valuation methodology, we have come to the conclusion that much of the problem lies in the unspoken assumption that the real-option and DCF valuation methods are mutually exclusive. We believe this assumption is false. Managers need to integrate the two approaches if they are to make valuations that reflect the reality and complexity of their business's growth projects. Far from being a replacement for discounted cash flow analysis, real options are an essential complement because they allow managers to capture the considerable value

of being able to ruthlessly abandon floundering projects before making major investments.

This is not to say that there aren't serious problems with the way managers calculate the value of real options. There are. For a start, real options, as currently applied, focus almost exclusively on the risks associated with revenues, ignoring the risks associated with a project's costs. It's also true that typical option valuations almost always ignore the fact that the initial investments made in a project, even in one that might eventually be abandoned, often leave the company with an asset it can trade – a benefit, if you like, of failure. These are not, of course, the only difficulties managers encounter using real options, but they are perhaps the most fundamental sources of error, and the integrated approach we present here explicitly addresses them both.

Integrating Options and Discounted Cash Flow

Traditional DCF analysis relies on the straightforward principle that an investment should be funded if the net present value (NPV) of its future cash flows is positive–in other words, if it will create more value than it will cost. This works well if we are projecting future cash flows from some historical context, and we are fairly certain of future trends, but not when our estimates of future cash flows are based on a myriad of assumptions about what the future may hold. In such cases, the odds of accurately forecasting cash flows are pretty slim.

What's more, even supposing we can arrive at a reasonably accurate base estimate for the cash flows, DCF analysis requires them to be discounted at a high rate to reflect the long odds of achieving the projected returns. As a result, all the risks of uncertainty (the possibility that actual cash flows may be much lower than forecast) are captured in the valuation but none of its rewards (the possibility that actual cash flows may be much higher than forecast). This inherent bias can lead managers to reject highly promising, if uncertain, projects.

The challenge, therefore, is to find a way to recapture some of the value lost through the conservative DCF valuation while still protecting against the considerable risks of pursuing highly uncertain projects. This is where options come in. The possibility that the project may deliver on the high end of potential forecasts, so hard for DCF analysis to take into consideration, is the primary driver of option value.

Options provide the right *but not the obligation* to invest in a project. Their value, therefore, is driven by the possibility of achieving a large upside gain combined with the fact that companies can usually abandon their projects before their investment in them has cost too much, thus limiting the downside. The value of an option must therefore increase as the uncertainty (and therefore the potential upside) surrounding the underlying asset increases, whether that asset is financial or "real."

COLIN JOHNSON

Looked at in this way, it seems clear to us that discounted cash flow analysis and real options are complementary and that a project's total value is the sum of their values.[1] The DCF valuation captures a base estimate of value; the option valuation adds in the impact of the positive potential uncertainty. One caveat though. It can hardly be stressed enough that a real-options approach can only be used on projects structured somewhat like options – that is, on projects that can be abandoned before

and DCF value is low – even, conceivably, negative. Now, uncertainty should reduce over time (if it does not, shut down the project!), so we move to the right, and the increasing certainty pushes up DCF value, through lower discount rates. But growing certainty also decreases the option value component of the project.

We do not believe that it is always necessary to calculate both components of a project's value. If the DCF valuation is high, the decision is easy – simply

Adjusting for Cost

That said, there remain two serious problems with option valuations. First, it is hard to find good proxies for the input variables the model requires. Financial options use a volatility measure derived from the easily observed historical prices of the underlying assets. But there are almost by definition no historical numbers that managers can use when trying to derive the option value of an innovative project – even to estimate the net present value of the underlying asset, let alone its volatility. (For a further discussion of the difficulty in finding good proxies, see the sidebar "The Trouble with Financial Option Tools.")

There's an unspoken assumption that the real-option and DCF valuation methods are mutually exclusive. We believe this assumption is false.

you must commit yourself to making major financial outlays if it becomes clear that things will not go well. It would not apply, for instance, to valuing an opportunity that requires you to sink huge sums into building a new factory before you have the first inkling whether the bet will pay off.

Once one accepts the notion that a project's value has both a DCF component and an option component, it also becomes clear that the proportion of a project's total value contributed by each component will vary according to the degree of uncertainty associated with the project. In the early stages of an innovative project, the value of the DCF component will be low because of the need to use a high discount rate to adjust for the uncertain nature of future cash flows. At the same time, the real-option value will most likely be high due to that same uncertainty.

The exhibit "Where the Value Comes From" shows how the relationship between DCF value and real-option value changes as the uncertainty of a project decreases over time. To the left of the diagram, uncertainty is high, so the project value, as measured by the vertical axis, is composed largely of option value,

proceed, since success in the project seems very certain, and it is likely to pay off handsomely. If the DCF valuation produces a strongly negative number and all the value comes from the option, then the project should probably be rejected, unless an investment structure can be created that would allow managers to learn a great deal about the project quickly and for very little cost. This rule of thumb may cause companies occasionally to miss profitable investments, but in our experience most large firms have more projects than they can fund or staff. So even if the option value is high, why waste time on a project that carries a large negative DCF value? It is simply too risky, so move on to something better.

The majority of growth projects, we have found, lie somewhere in the middle. When the DCF value of a project is modestly positive, or somewhat negative, the project lies in what we call the "option zone" – the gray area where managers have usually been forced to rely on their intuition in making the investment decision. It is here that our framework is particularly useful because the option value can provide logic to support or refute that intuition.

Second, even if managers succeed in finding good proxies for the option-model input variables, they remain vulnerable to a major conceptual error. In the current approaches to option valuation, the more variable the profits, the higher the project valuation. The variability of profits, in turn, is derived from estimates of how uncertain both revenues and costs are likely to be. This seems reasonable but leads to an impractical result: Mindless option analysis will value a project with relatively predictable revenues but unpredictable costs more highly than a project with the same predictable revenues but with predictable costs. We think this is wrong. When the uncertainty about potential costs is higher than the uncertainty about potential revenues, cost volatility should decrease, not increase, the value of a project.

Why? Unlike revenues, where volatility can imply as much upside potential as downside, when it comes to costs, the potential for downside is generally much greater. That is, the margin by which costs overrun their estimates is almost always greater than the margin by which they underrun them. Take, for example, a European consortium that set out to build a new fighter aircraft for a projected $20 billion. It's now 15 years behind schedule, and the estimated cost is $45 billion, an overrun of 125%. We do not routinely see cost savings on any-

*Alexander B. van Putten (alexvp@wharton.upenn.edu) is the principal of Triad Consultants and an adjunct professor at the University of Pennsylvania's Wharton School of Business in Philadelphia. **Ian C. MacMillan** is a professor at the Wharton School.*

thing like this same scale. Indeed, with most projects, it's possible to be fairly certain of the minimum cost, which places a floor on how far the costs can fall. But there's no corresponding ceiling for overruns.

What's more, the chance that cost uncertainty will cause overruns increases when companies look for growth opportunities in areas outside their direct expertise. Companies' appreciation of the risks of a project, as the Nobel Prize–winning work of psychologists Daniel Kahneman and Amos Tversky has shown, is systematically too optimistic. The experience of a large industrial company we worked with that was venturing into biotech aptly illustrates how easily the costs of a growth project can spiral out of control when a company is operating in areas far from its expertise and experience.

For millions in R&D dollars, the company had developed a new compound that held great promise as an additive for a number of consumer products. At the time we became involved, project managers had already spent money on toxicity testing and had made other large safety-related expenditures, followed by sophisticated consumer testing, all of which indicated that the compound held considerable potential to command high prices. But the firm had not yet tried to ramp up manufacturing to produce the compound in commercial quantities. Based on long years of experience, management simply guessed that it could be produced for approximately $20 per unit and paid no more attention to the costs of commercial production.

It turned out, though, that the manufacturing process was hugely more difficult than anticipated. The cost to produce the compound would be in the order of hundreds of dollars per unit, which put it outside the range of commercial viability.

Had company managers taken cost volatility into account effectively, they would have managed the project differently. First, they would have realized sooner that the manufacturing process represented the greatest part of the

Where the Value Comes From

The relative amounts that real-option and discounted cash flow valuations contribute to a project's total value vary with the project's uncertainty. The greater the uncertainty, the larger the option component and the smaller the discounted cash flow component. In fact, when the total project value (TPV) of a project is made up almost entirely of option value and the net present value (NPV) of its DCF component is highly negative, the project falls in what we call the flee zone–the zone filled with projects far too risky to consider unless the TPV is truly huge. Conversely, when uncertainty is very low, a project's TPV will be made up almost entirely of DCF value. Then it will be in the deep-in-the-money zone. Its net present value will be so high that it's not even necessary to calculate the project's option value before approving it.

But between the flee zone and the deep-in-the-money zone is what we call the option zone, where the contribution of the option component adds meaningfully to TPV. It is here that traditional DCF valuations usually clash with management intuition, and so it becomes important to compute both the DCF and the option value of a project. In this example, project A (depicted by the solid vertical lines) is squarely in the option zone. As project A progresses, uncertainty should be reduced, so the vertical line should move to the right, as escalating certainty increases the DCF component and decreases the option value component.

The Components of Total Project Value

Option Value Component   DCF Component
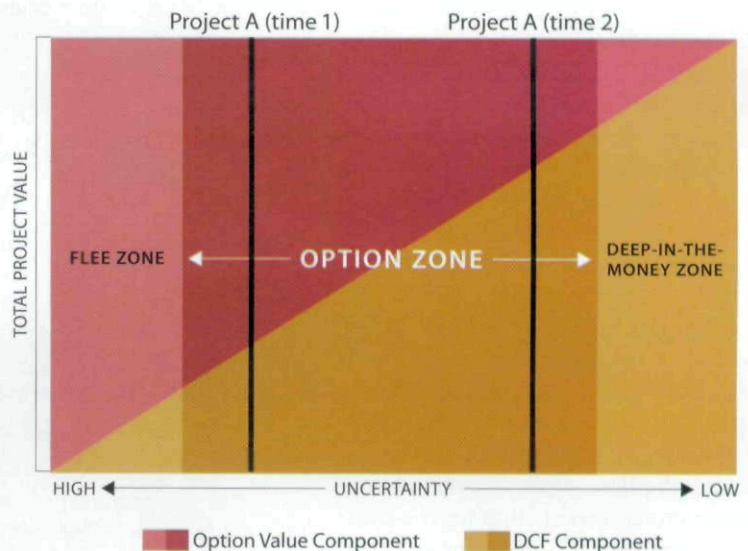
uncertainty surrounding the project. That would have encouraged them to switch the business development effort from product R&D toward process R&D, so that they would first have understood manufacturing feasibility and only afterwards have investigated consumer

Our framework is particularly useful when the discounted cash flow value is modest because the option value can provide evidence to support or refute a manager's intuition.

demand. Second, taking into account cost volatility would also have produced a much smaller total project value, which would have led them to curtail investment in the project at an earlier stage, saving them millions of dollars.

Since costs are volatile in a different way than revenues are, the formula for determining option value needs to be adjusted when cost volatility is greater than revenue volatility. In principle, you could figure an adjusted option value (AOV) that reflects the negative nature of cost uncertainty by separately calculating the option value of the revenues and then subtracting the option "value" of the costs.

In practice, however, there is no need to compute the impact of cost volatility separately from the impact of revenue volatility. There is a simpler approach that is good enough for inferring the AOV of a project, when necessary, and that has the advantage of being simple and quick. Simple and quick is what's needed for most valuations: In any firm with far more projects under consideration than funds or staff to support them, managers need not have a precise value for a specific project; they need only know whether a project is preferable to other projects competing for the company's limited funds and talent. So rather than being concerned with whether a particular valuation is precise, manag-

ers should look at it as a yardstick that allows them to choose the best among competing projects. As long as they feel sure that all the projects applying for funds are being valued in the same way, they can be reasonably confident that they will, on average, select and assign resources to the best ones.

So, keeping it simple, to give costs a truer weight in an option valuation, when cost volatility is greater than revenue volatility, we adjust the volatility of the project as a whole (the volatility number we normally input into an option calculation) to reflect the negative nature of cost volatility. Then we apply that adjusted number to the project's option valuation. The volatility adjustment is made according to the following formula: If cost volatility is greater than revenue volatility: adjusted volatility = project volatility × (revenue volatility ÷ cost volatility).

In other words, if we are more certain about the projected revenues than we are about the projected costs, then the ratio of revenue volatility to cost volatility will be less than one, which will reduce the overall volatility, and that, in turn, will reduce the option value of the project. For instance, for a project with an overall volatility estimate of 45%, a revenue volatility of 40%, and cost volatility of 60%, the adjusted volatility will be 45% × (40% ÷ 60%) = 30%. This adjustment has the effect of discounting

the value of the option due to the higher cost volatility. If revenue volatility is higher than cost volatility, then the project volatility variable in the real-option calculation need not be adjusted.

Adding the Rewards of Failure

Failing to adjust option value to reflect cost risks is not the only source of error. The second option component often missing from managers' calculations that our approach incorporates is the abandonment value (ABV) of a project. In searching for ways to reduce cost volatility, managers often find they can recoup some of the investments they have made, in the event of failure. These opportunities for creating extra value when halting a project can be seen as the equivalent of the put options familiar to financial investors, which serve as a hedge against drops in the price of the underlying asset.

Abandonment value can arise in a number of ways. In some cases, early investments that have to be abandoned can be valuable to another business unit within the same company. Take the example of a large industrial company that had developed a plant-based vitamin precursor. This was novel technology, but it appeared to have little value to the health care industry because it didn't clearly perform any better than existing compounds. Another division

A Formula for Valuing in the Option Zone

$$\text{TPV} = \text{NPV} + \text{AOV} + \text{ABV}$$

| total project value | net present value | adjusted option value | abandonment value |

If cost volatility exceeds revenue volatility, the volatility number for the calculation of AOV should be adjusted as follows:

adjusted volatility = project volatility × (revenue volatility ÷ cost volatility)

If cost volatility does *not* exceed revenue volatility:

adjusted volatility = project volatility

of the company, however, picked up the compound and used it in a joint venture that was developing new food additives for the Asian aquaculture industry, where the compound was shown to accelerate the growth rate of farm-raised shrimp.

In other situations, the early investments may have created an asset that can be traded for cash or equity in another company. GlaxoSmithKline, for example, developed an experimental antibiotic that showed promise in treating drug-resistant staphylococcal infections but was thought unlikely to become the sort of blockbuster drug the company needed to support its growth rate. Rather than consign the intellectual property to its library of interesting compounds, the firm generated abandonment value by trading the patents, technology, and marketing rights to develop this antibiotic for equity in Affinium, a privately held biotech company.

To see how seriously managers take abandonment value, consider what happens in joint-venture negotiations, where the issue of control over the venture's future is usually hotly contested. The cost of owning 1% more of the joint venture than the other party in a two-party agreement is typically far higher than the economic value attached to that additional 1% equity interest. That's because the controlling party can usually force the venture to be liquidated if trouble occurs, and managers understand (at least tacitly) that this privilege has value, which can be explicitly calculated in the course of negotiations.

If the opportunity to create value on exit exists or can be made to exist, then managers should include that factor in their project valuations. This involves another option calculation. Because the exit option is usually a relatively simple real option (a put option), managers can fairly easily apply financial tools like the Black-Scholes-Merton formula. The estimated value of the asset created by the aborted investment is the exercise price. The historical range of prices paid for comparable assets determines the volatility. The date on which the company has to decide whether or not to

The Trouble with Financial Option Tools

Managers should beware of a number of technical difficulties in simplistically applying standard financial option–tool valuations (like Black-Scholes-Merton) to real options. Not the least of them is trying to establish a figure for volatility, for which there are often no historical numbers.

To determine a project's volatility, then, we must first develop a financial model of the business using the *most likely* values for all the factors that drive costs and revenues. We use these to compute the expected total costs and revenues for the DCF component of the project's total value. Then for each factor, we specify the *range of possible* values. These ranges (whose widths reflect their associated uncertainties) are put into a Monte Carlo simulation, from which we extract the means and standard deviations of total profits, total revenues, and total costs. The standard deviations of profits, revenues, and costs are used in the calculation of adjusted volatility described in this article, and this adjusted volatility is then used in the option valuation. The mean of the project value, discounted back at a risk-adjusted rate, becomes the proxy for the current price of the underlying asset.

We would emphasize, however, that if the original projections are flawed (which is very possible with a highly uncertain growth project) or if the discount rate is wrong (even more likely), the volatility and exercise price estimates will also be wrong. Realistically, in fact, with highly uncertain projects, *any* method, no matter how sophisticated, will be wrong. Hence our contention that time spent worrying about the exact option value of a project is time wasted. What valuation can and should do is establish relative values within a portfolio of opportunities, providing a means of ranking the contenders, so that managers can select only the most promising. That way, managers will, in the long run, select better projects than their more timid competitors while keeping risk under control and thus outperform their rivals in both the product and the capital markets.

Another source of error involves the time period used in the calculation, and this is even more difficult to resolve. With a financial option, the more time we have before we commit to buying the underlying asset, the more valuable the option. This makes sense because the stock has more time to increase in value, and if it does not, we need not exercise, so financial options with longer expiration periods have more value than those with shorter lives (all other things being equal). This logic does not extend to the real world, however. Delaying a product launch will not necessarily add value to a project because you end up paying a discount penalty and could even end up missing the market. The relationship between time and value is much less consistent with real options than it is with financial options.

The best way to handle the problem is to formally recognize this competitive reality. What we do is estimate how long it will take before competitors' moves erode net revenues (seldom more than seven years, sometimes as little as three) and use this as the time period for the evaluation. We assume that the project is launched immediately because there is no bonus for delay. If the project is delayed, we actually discount the total project value for the period it is delayed.

continue investing in the project is the time to expiration.

Let's go back to the case of a manager trying to negotiate a controlling stake in a joint venture. Assume that the venture will create an asset in the form of the jointly owned plant and equipment, which the manager expects could be worth approximately $15 million as a going concern. In the past, prices of similar assets, which have been relatively easy to sell, have been worth about $10 million in a liquidation, and price volatility has historically been 45%. The project has three phases, and the assets could be sold if the venture is dissolved at the end of phase one, in about two years. Assuming a risk-free interest rate of 3%, and that current prices for the asset are the same as the forecast price, a Black-Scholes-Merton computation yields an abandonment value of approximately $983,000, which should be added to the total project value and which should also serve as a reference point for the value of control over the liquidation of the venture. In negotiating a 51% stake in the venture, therefore, the manager should be prepared to pay up to, but no more than, $501,330 (51% of $983,000) for the extra control – and that's forgetting for the moment about other benefits that control might offer.

The Option Zone in Real Life

Our integrated approach to investment is not just an exercise in theory. John Hillenbrand and Mary Kay James of DuPont Ventures, working with consultant Hal Bennett and John Ranieri, vice president of DuPont Bio-Based Materials, have for some time been using an expanded concept of total project value that is very similar to the approach set out in this article. DuPont Ventures looks for externally owned new technologies that could be commercialized by a DuPont business unit.

When Ventures finds an interesting technology within an early-stage company seeking financing, the unit will buy into the current round at the same valuation as other investors, on one condition. It must also acquire the right of first refusal to license the other company's technology for specific markets that may interest DuPont but are not primary markets for the target firm. After closing the deal – typically, for between $1 million and $3 million – Ventures assigns the technology to an interested internal business unit, which could then commercialize it using Du-Pont's substantial resources. If no license agreement is completed, Ventures still retains its equity interest in the target company, which may or may not have liquidity in the future.

In making the decision to invest, Ventures uses all the elements of our valuation approach: discounted cash flow, adjusted option value, and abandonment value. When Ventures first considers a new investment, it looks at the target company's projections and makes a DCF calculation as a base case valuation. Of course, the target company's projections do not factor in the benefits derived from being associated with a *Fortune* 100 firm, so they underestimate the value of the technology as Ventures sees it. The next step in the analysis, therefore, is to work with interested business units within DuPont that might possibly commercialize the technology to generate more complete projections and calculate the option value of the investment. In making these projections, Ventures looks closely at the range of costs that DuPont will incur if it were to commercialize the technology, as well as the uncertainty surrounding the yet-to-be negotiated license terms with the tar-

Simple and quick is what's needed for most valuations. Managers need only know whether a project is preferable to others competing for limited funds and talent.

get company. That leads to a cost volatility estimate. The result of this exercise is equivalent to the AOV term in our approach. Finally, Ventures also takes into account the fact that it will retain an equity interest in the target firm, which could potentially be sold whether or not a DuPont business unit invests in the technology. This is equivalent to the investment's ABV and adds to the total project value. The approach has worked well for Ventures, which has developed a robust portfolio of promising opportunities that it would otherwise have missed.

•••

The challenges of growth are forcing companies to evaluate and support increasingly uncertain projects, which in theory require some kind of options framework in order to value them properly. But CFOs and CEOs voice justifiable concerns over the idea of simply replacing the long-trusted DCF model with a real-options calculation. The integrated approach we have presented attends to those concerns and will enable senior managers to make more aggressive investments while meeting their fiduciary responsibilities. We invite managers to test it out on a few pilot projects – ones that their gut feelings tell them deserve funding despite what the DCF numbers suggest or ones with high option values about which they nevertheless have reservations. But remember: Option valuations only make sense when applied to projects that can be terminated early at low cost if things don't go well. And no valuation method will save a company that does not actually pull out quickly, if the project fails to deliver on its initial promise, and redeploy talent and funding elsewhere. If this fundamental option discipline is not baked into every option project, you are not investing, you are gambling.

1. See Robert K. Dixit and Robert S. Pindyck, *Investment Under Uncertainty* (Princeton University Press, 1994) and Lenos Trigeorgis, *Real Options: Managerial Flexibility and Strategy in Resource Allocation* (MIT Press, 1996).

Reprint R0412K
To order, see page 151.