# Making real options real

CHARLENE POTTS/SIS LTD

Most applications of option theory have been oversimplified

What happens when options generate more options as well as cashflows?

We offer a second-year course on compound options[*]

Thomas E. Copeland and Philip T. Keenan

ONCE AN OBSCURE MATHEMATICAL TOOL, options and option valuation have entered the mainstream. Long routinely applied on trading floors, the Black–Scholes model for valuing options won the 1997 Nobel Prize for economics. Yet despite the fact that senior managers realize that traditional analytical methods such as net present value (NPV) and economic profit (EP) have been responsible for systematic underinvestment and stagnation, options have penetrated the decision-making processes of large corporations more slowly.

For NPV and EP ignore an important reality: business decisions in many industries and situations can be implemented flexibly through deferral, abandonment, expansion, or in a series of stages that in effect constitute real options (Exhibit 1). Recognizing real options can help decision makers assess the profitability of new projects and understand whether and when to proceed with the later phases of projects that have already been initiated, particularly when they are close

Tom Copeland is a former principal in McKinsey's New York office and *Phil Keenan* is a consultant in the Cleveland office. Copyright © 1998 McKinsey & Company. All rights reserved.

Exhibit 1

Recognizing real options

Situations generating real options, by sector

Aerospace and defense
Valuing options in customer contracts (eg, the cancellation of features)

Assembly and machinery
Timing of investment in new factories

Automotive
Valuing decisions to modify new car designs well into product development cycles

Banking and securities
Valuing real estate leases

Chemicals
Timing of investment in new factories

Consumer/packaged goods
Learning options for new product R&D and test marketing

Electronics
Entry or exit options in the PC assembly business

Energy
Learning options in timing the development of oil and gas fields

Insurance
Valuing such contingency features as loans against value of policies

Media and entertainment
Planning new product launches

Metals
Learning options in timing the development of new mines

Pharmaceuticals and medical products
Establishing priorities among prospective R&D projects

Pulp and paper
Timing of forest harvesting

Retail
Timing of expansion into foreign markets

Telecommunications
M&A programs for geographic expansion

Transportation
The option to extend or contract routing structures (eg, railroad spurs)

to breakeven. Real options are especially valuable for projects that involve both a high level of uncertainty and opportunities to dispel it as new information becomes available.*

Although the overall concept of real options is clear, their specific benefits for individual businesses are not. To date, most attempts to apply real options have been too simplistic to address the complexity of the decisions managers face. In deciding whether to finance a research and development effort, for example, managers must reckon not only with technological uncertainty concerning the outcome of the research but also with market uncertainty about the eventual demand for the resulting product. When a company contemplates developing a mine or a natural gas field, both the market price of the output and the level of output are uncertain.

To bring practicality and detail to thinking about real options, we have worked through a series of cases that capture the business problems and realities they involve.

Compound and learning options

An option depends on an evolving source of uncertainty. An electric utility with modern generating plants may, for example, have the option of bringing a mothballed coal-fired plant back on line. The main source of uncertainty in this case is the volatile price of electricity. Keeping the coal plant dormant gives the utility the option of turning it back on if electricity prices rise sufficiently.

---

* *See* Keith J. Leslie and Max P. Michaels, "The real power of real options," *The McKinsey Quarterly,* 1997 Number 3, pp. 4–22, Lenos Trigeorgis, *Real Options: Managerial flexibility and strategy in resource allocation,* MIT Press, Cambridge, Mass., 1996, and Avinash K. Dixit and Robert S. Pindyck, *Investment under Uncertainty,* Princeton University Press, Princeton, NJ, 1994.

Simple options give their holders the right to buy or sell an item at a particular price on one occasion, at or before a particular expiry date. Even simple options are marked by changing sources of uncertainty, such as the price of electricity or polyester. The holder must decide whether to exercise the option by judging how the uncertain quantity has evolved since the option was acquired.

In this article, we focus at first on compound options: those that when exercised generate another option, as well as a cashflow. In general, compound options involve sequenced or staged investments. Making the first investment gives a company the right but not the obligation to make a second investment, which in turn confers the right to make a third, and so on.

We also look at learning options: those where the holder pays to learn about an uncertain quantity or technology. An oil company may own the rights to a plot of land, say, without knowing exactly how much oil it contains. Rather than trying to predetermine a particular level of production capacity, the company might find it worthwhile to spend money discovering the actual extent of the reserves. Then it can develop the field without wasting resources building facilities designed to process more oil than is actually there – or, conversely, delaying profits by constraining production rates for fear of depleting wells that in reality contain more oil than it realizes. Similarly, companies with R&D programs invest to reduce technological or knowledge-based uncertainty.

Staged investments give managers the option to abandon or scale up projects well into their lifetimes. These compound options can be highly valuable, as in our case below of a chemical company deciding whether to invest in a new plant for manufacturing polyester. Similar factors apply in any manufacturing industry when a new factory or a large capital investment is under consideration. Decisions about expanding into new geographic areas and investing in research and development also involve staged investments of this kind.

Case study: To build or not to build?

In cyclical industries, the key to profitability is often the ability to know at what point in the business cycle to build a new factory. The only trouble is, you can never quite tell where in the cycle you stand. Fortunately, the management of the company whose story we now consider did not have to commit itself outright to a new factory. It had the option of staging the investment over 12 months: $50 million up front for design, $200 million six months later for pre-construction work, and $400 million to complete construction at the end of the year. Even if the company paid out for the design stage, it could walk away after six months with no further investment if profit projections had fallen in the meantime. It could even back out at the end of the initial pre-construction phase and save the $400 million of the final investment.

The factory was designed to convert an input chemical, p-xylene, into an output chemical, polyethylene terephthalic acid (PTA). The price per ton of the input and the output fluctuate. The profitability of the factory would depend on the spread between these prices, adjusted for the number of tons of input needed to create a ton of output. Instead of modeling both prices independently in this analysis, it was possible to simplify the calculation considerably by modeling the spread.

Within a business cycle, the right time to build a factory may vary from industry to industry because of differences in the length of the cycle, the lead times needed to construct facilities, and the lifetimes of the facilities. For the polyester industry, historical data on the spread between the prices of the input and output chemicals suggested that it was mean reverting (Exhibit 2). In other words, it rises and falls but tends to return to a long-term industry average, with high spreads followed by low spreads and *vice versa,* perhaps because high spreads increase competition by encouraging new entrants, while low spreads encourage players to leave the industry.

Exhibit 2

The spread between input and output prices

Price history: p-xylene and PTA, $ per ton

PTA price

p-xylene price

Operating spread = PTA price/ton −0.67 (p-xylene cost/ton)

Received opinion held that a company in this industry should invest in new factories only when the input/output spread was considerably higher than its long-term average. A standard NPV analysis based on the then-current spread found that building the factory would have a net present value of minus $70 million, which would suggest that the project should not be undertaken. On the other hand, the real option valuation, taking into account the value of the option to abandon the project after six months or a year, found that it had a positive value of $350 million, suggesting that the company should invest in the design phase at the very least. The real option approach also clarified the criteria for making decisions later on: namely, the cutoff values for the spread below which it would make sense to abandon the project.

What logic lay behind the dramatic increase in value? The problem with a standard NPV analysis is that it is too sensitive to the spread. Building the

factory costs about $650 million. If the spread declines, future revenues will fall considerably short of that sum, thus causing a net loss; if the spread increases, there will be a net profit. An appropriately weighted average puts the expected revenue at $580 million, yielding the net loss estimate of $70 million.

Unfortunately, a weighted average of outcomes does not capture the flexibility inherent in multistage investments. In reality, phasing in the investment gives the company a chance to cut its losses. At the one-year point, for instance, it has invested only $250 million. If profit projections indicated that revenues would be less than $400 million, the company would choose not to invest the remaining $400 million to complete the factory, confining its loss to the money already spent.

In other words, the NPV analysis assumes that the factory will definitely be built and operated, ignoring management's ability to walk away if revised projections make it advisable to do so. Averaging out large losses and gains, the NPV approach produces the figure of minus $70 million. Because the options approach recognizes management's power to put a floor under the losses, it factors in smaller losses while not losing sight of the potential for large gains. It thus comes up with a much higher value.

The growth staircase as a compound option

The staged investment in the polyester factory exemplifies a growth staircase: a sequence of stages, each contingent on the completion of its predecessor.* Consider a domestic company seeking to expand in foreign markets. It might start by entering a single carefully chosen territory where it could gain experience in international marketing and finance. If it learned that its products had limited appeal there, it could modify them and test-market the changes. When it finally achieved success in that first foreign beachhead, it could expand into similar overseas markets at less risk than it would have faced had it entered several of them simultaneously. Naturally, if its competitors are also likely to enter overseas markets, it must weigh the value of the option to expand cautiously in these countries against the potential cost of coming second in some or all of them.

Another kind of growth staircase is illustrated by the case of a large enterprise acquiring a small one with a proprietary technology that it cannot develop itself for want of resources. The acquirer gains the value not only of the small company's current operations but also of the option to develop and market an improved version of the technology. In such a case, the acquisition may truly create value, since the market valuation of the independent small company may not have reflected the value of the growth option. By taking advantage of projections or historical information about the uncertainty surrounding

---

* *See* Mehrdad Baghai, Stephen C. Coley, and David White with Charles Conn and Robert J. McLean, "Staircases to growth," *The McKinsey Quarterly,* 1996 Number 4, pp. 38–61.

the new technology, real option valuation can be used to quantify the acquisition's potential benefits.

Acquisitions can be viewed as investments in capabilities or infrastructure. A sequence of acquisitions constitutes a staged series of investments. Thus the growth staircase is a compound option. At each stage, it is not clear that an acquisition will provide the anticipated growth platform. If it does, the company that makes the acquisition can proceed with the next one. If it does not, the acquirer has the option of selling the latest acquisition, though not necessarily recouping its full cost. Real option valuation can be used to value all possible contingencies accurately.

Two sources of uncertainty (rainbow options)

Many business decisions are beset by uncertainties. A bank's value, say, may depend on uncertain interest rates and an uncertain number of customers. Although we can use real options to value the bank if it has some degree of flexibility in the decisions it makes, both these uncertainties, which may be connected, increase over time.

Learning options are different because they involve the possibility of reducing some of the uncertainty by making an investment: in other words, paying to learn. This might, for example, involve drilling in a field to learn more about the quantity of natural resources located under it. In this case, though the price of these natural resources becomes more uncertain over time, the quantity becomes less so. For the purpose of illustration, we will assume that the price of the resource in world markets has no bearing on the outcome of the learning option.

Case study: Mining ore

Mining companies have to decide when to develop the properties they own and how much to bid for the right to develop additional properties. Such decisions often involve a combination of options: the option to learn about the quantity of ore present underground and the option to defer development until ore prices are favorable.

These two options frequently conflict. Immediate development sacrifices the deferral option, but provides information on the quantity of ore available. Deferring development allows the owner to wait until the price of ore is high enough to ensure profitability, but provides no information on the quantity of ore. At times, partial or exploratory development, allowing a company to learn something about its holdings while preserving the ability to hold back actual production until prices improve, may be a reasonable compromise.

In general, learning options arise when a company can speed up the arrival of important information by making an investment. As always, an option

arises only if the information can modify future investment decisions. Companies with learning options must balance the value of the option to act on the knowledge they gain against the cost of obtaining that knowledge in the first place.

If a company owns the right to mine land for mineral ore, it must decide whether to develop the property immediately or defer development in the hope that the price of the ore will rise. Imagine that at the current price, the mine is close to breakeven, and the only uncertainty is the price of the ore when sold. Historically, prices have followed a random walk drifting upward over time. If the company knew that this trend would continue, it would decide to defer development until the price rose far enough to make mining clearly profitable. It holds a valuable option: that of deferring the opening of the mine. Real option theory can be used to determine the best time to exercise the option.

If, however, the company is also unsure about the quantity of ore in its mine, it holds a learning option: to pay money to find out. Real option theory can be used to analyze this option too. Note that the two options can conflict: it may not be possible to learn how much ore the mine contains without actually beginning to exploit it. This means that for the company and its stockholders, the combined value of the two options is less than the sum of their individual values.

To give a concrete though somewhat simplistic example, suppose the company thinks that the deposit holds a million tons of ore. It estimates that for an initial cost of $100 million, it can build an infrastructure capable of mining 100,000 tons a year for a decade. Extraction and processing will cost $200 a ton, and current spot and futures market prices indicate that a decade's worth of ore production can be sold for a locked-in rate of $360 a ton. Thus, with a weighted average capital cost of 10 percent, the project has an estimated NPV of minus $2 million – a net loss.

Look at it another way. The price of ore is volatile, and the company does not have to develop the mine immediately; instead, it can wait for a year and see if prices rise. If they go up to $400 a ton, the project will have an NPV of $21 million, even factoring in the year's delay in cashflows. If prices fall to $320 a ton, the company can choose not to open the mine, thus reaping an NPV of zero. Acting on realistic assumptions, we can combine these numbers into an expected NPV of $10.5 million, which means that the company should defer the decision for a year and then reconsider it. Should it be deferred for 10 years, the expected NPV rises even higher, to $30 million, since there is a good chance that ore prices will rise above $400 a ton in the course of a decade.

Now suppose that the quantity of ore in the mine is also uncertain. There is a 50/50 chance of its holding 0.5 million or 1.5 million tons. The expected quantity is thus still a million tons. Given these assumptions, the company

faces an additional element of risk by investing the $100 million in infrastructure. If the deposits contain only 0.5 million tons of ore, an infrastructure costing $50 million would be sufficient to extract it in 10 years. If the deposits contain 1.5 million tons, a larger infrastructure costing $150 million would be needed to handle the ore over a decade. A too small infrastructure delays the extraction of the ore; a too large one is a waste of money, although it accelerates the rate of production.

Exhibit 3

NPV of mine according to ore quantity and infrastructure investment*

* Includes value of deferral option

What should the company do? Even if we factor in the value of the 10-year deferral option, an incorrect guess as to the quantity of ore can be expensive (Exhibit 3). Since the chance of guessing correctly is only 50 percent, the expected NPV of investing $50 million for the small infrastructure is $25.5 million and that of investing $150 million for the large infrastructure $23.5 million. The smaller investment in infrastructure seems to make more sense.

Suppose, however, that there is a learning option. To keep the analysis simple, say that for $1 million the company can give itself a 90 percent chance of learning how much ore the site contains. In this event, learning does not conflict with deferral.

If the company finds out that there is likely to be a lot of ore, it invests in the larger infrastructure for an expected NPV of $39.9 million; if it learns that there is likely to be only a little ore, it builds the smaller infrastructure for an expected NPV of $17.1 million. The overall expected NPV is $28.5 million, or $27.5 million after the $1 million cost of testing has been deducted. Thus it pays to learn. The company can improve the project's overall NPV by spending money up front to find out how much ore there really is.

In this example, the benefit of learning is not very great; if acquiring knowledge cost more than $3 million, the company would be better off investing in the smaller infrastructure without making an effort to learn how much ore it has. The benefit is modest because the smaller infrastructure does not have much of a downside, even if the quantity of ore turns out to be large. After all, it will be possible to unearth all of the ore eventually even though it will take longer.

If there is a time constraint, such as a competitor's expected entry into the market for this kind of ore, long time horizons may become uneconomic. Under such conditions, the level of risk in building a smaller infrastructure for an amount of ore that turns out to be large is higher than it would otherwise be, and paying for additional information about the deposit makes more sense.

Case study: Developing a natural gas field (compound rainbow options)

Combinations of (compound) learning and rainbow options often arise. Consider a company that has to decide how much production capacity to install in an undeveloped natural gas field. Roughly 40 percent of the holding has yet to be explored, so the total amount of gas available is not clear. Nor is the price it would fetch over the many years of extraction. The company has to decide whether to proceed immediately or wait until further exploration has yielded a better estimate of the total size of the reserves. It must also choose between committing to a given level of production capacity or building capacity in a way that allows it to add more in future, a solution that costs more, at least in the early stages.

The company's expectations about the quantity and price of the natural gas can be used to plan the production capacity that would be appropriate if it were to be installed immediately and locked in permanently. But this scenario eliminates all of the embedded options: deferral, further exploration, and expansion. The company therefore creates a decision tree for a real option

Exhibit 4

Decision tree for ROV model of natural gas investment
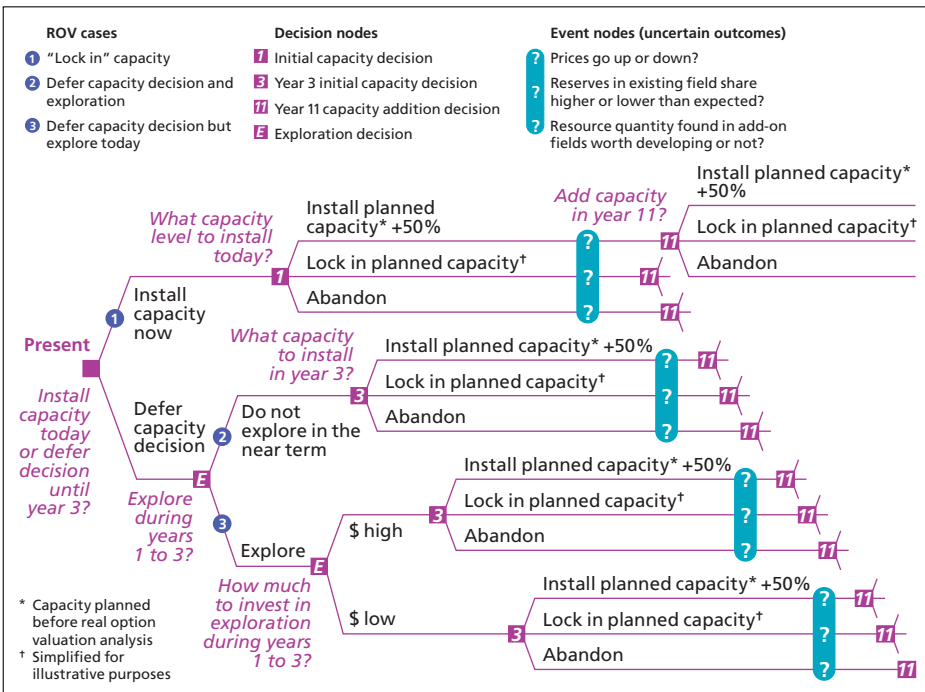
valuation (ROV) model to weigh up the various choices in view of the actual uncertainty as to price and quantity (Exhibit 4).

This model sets out the consequences of installing various levels of production capacity – locking in the level planned immediately or installing enough equipment to extract that amount of gas plus 50 percent extra in reserve – and of deferring the decision for a few years. It also evaluates two different forms of deferral: simply waiting to find out about prices on the one hand, and conducting further exploration to clarify the quantity of reserves on the other. In addition, it compares the consequences of locking in the planned capacity at the end of the deferral point two years later, or of investing extra money at that point to provide for the expansion of capacity in 10 years if the price and quantity information then available seems to justify doing so.

Exhibit 5

Real option valuation model*

ROV case	Initial capacity decision	Exploration decision	Add-on capacity	Total project NPV from valuation Index: Base case = 100
No options base case[†]	Today (planned capacity)	Today (no exploration)	Today (no expansion)	100 + Value of better capacity lock-in decision, exploration option, and expansion option
❶ "Lock in" capacity	Today (planned capacity +50%)	Year 6	Year 11	150 + Value of price information
❷ Defer capacity decision and exploration	Year 3	Year 6	Year 11	200 + Value of reserve size information
❸ Defer capacity decision but explore today	Year 3	Today (explore)	Year 11	225

* For purposes of comparison and because of lack of information, each of the four cases assumes a year 1 exploration cost equal to 0; if "best guess" exploration cost estimate of 40 is used, the NPV for the 3 option cases is 120, 175, and 185, respectively
† Analogous to traditional DCF case (ie, assumes deterministic inputs and no managerial flexibility).

Using information about the volatility of natural gas prices and the uncertain quantity of gas in the field, the real option valuation model estimates the total value of the different courses open to the company (Exhibit 5). Locking in capacity immediately at a level 50 percent higher than originally planned turns out to add 50 percent to the project's total value. Deferring the decision for three years and leaving open the possibility of future expansion (but not exploring in the meantime) doubles the project's original value. The most promising combination, with more than double the value of the project as originally planned, turns out to be deferring the decision about capacity, exploring immediately, and leaving open the possibility of future expansion.

Case study: R&D in pharmaceuticals

Research and development projects combine learning and compound options. Say that a pharmaceutical company wants to rank possible R&D projects in

order of priority. It would probably calculate an NPV for each one and select only those with a positive NPV, or determine its priorities by implementing projects in the order of their NPV rankings, from highest to lowest. Traditional NPV calculations, based as they are on the expected cashflows from undertaking projects fully, can substantially underestimate their true value by ignoring embedded options.

Typical R&D projects face both technological and product market uncertainties: the former from concerns about the ability of researchers to invent safe and effective new products, the latter from the vagaries of future demand. The real option approach recognizes both kinds of uncertainty and the possibility of a flexible response, typically by staging investments. If an initial investment in pharmaceutical research is successful, the company concerned must make further investments in development: that is, various kinds of clinical testing. Should the drug prove safe and effective, the company will then have to make investments to produce and market it. R&D projects are classified as compound rainbow options because they involve a sequence of options, each contingent on those preceding it (compound options) and on multiple sources of uncertainty (rainbow options).

To keep the example simple, let us assume that the drug can be evaluated in a single trial phase (rather than the three required in the United States) and reduce the range of possible outcomes. Once understood, the approach and calculations are easy to extend to as many phases and outcomes as may be appropriate. Suppose that the company is contemplating an R&D project with an initial research phase followed by a development phase. Initial research costs $1.5 million, takes a year, and has three possible outcomes, each with a degree of probability that can be extrapolated from similar research projects in the past. The chance of creating a highly effective drug (or one that will be used widely) is 10 percent; a moderately effective drug (or one with a narrower range of uses) is also 10 percent; and no drug at all is 80 percent.

If the researchers do come up with a new drug, the development phase, also lasting a year, will cost $5 million. There are two possible outcomes at this juncture: that the drug will pass the safety tests (40 percent) or fail them (60 percent). (In a more complex example, the company might have the option of accelerating the test phase, perhaps in response to competitive pressure, by laying out more money.)

Should the drug pass the safety tests, it can be marketed. The company estimates that a highly effective drug would reap annual revenues of $19 million. For the sake of simplicity, assume that this would continue in perpetuity and that the company's weighted average cost of capital is 10 percent, so the project has a present value of $190 million at the start of the marketing phase. A moderately effective drug would have a present value of $140 million. Needless to say, all these revenues lie at least two years in the

future and are thus uncertain. The marketing department suggests that a great product might generate annual revenues ranging from $12 to $30 million, and a mediocre one from $9 to $23 million. Building a factory to produce the drug will cost about $130 million.

Should the company go ahead with the project? Not yet aware of real option valuation, the company uses traditional approaches: NPV and decision trees. Realizing that the problem is too complex for a single-scenario NPV calculation, it computes several scenarios. For a great product, the expected NPV is $43.5 million, since the cashflows (before discounting at the 10 percent cost of capital) are minus $1.5 million in year zero, minus $5 million in year one, and $190 million less $130 million in year two. For a mediocre product, the expected NPV is $2.2 million. Then again, there could be no product at all, which would generate a present value loss of $1.5 or $6 million, depending on whether the company carried the project into the development phase.

Exhibit 6

R&D decision tree
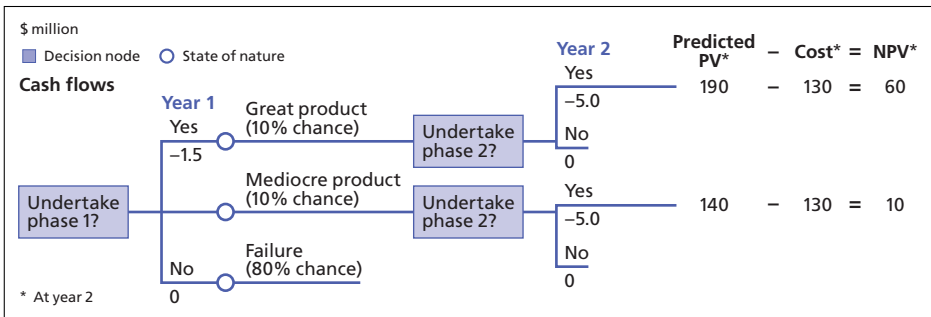
Next, the company combines these scenarios into a single probability-weighted NPV by drawing a decision tree (Exhibit 6). There is a 4 percent chance (10 percent in the research phase times 40 percent in the development phase) of creating a great product with an NPV of $43.5 million, and a 4 percent chance of creating a mediocre product with an NPV of $2.2 million. The chance that the research will be an utter failure costing the company $1.5 million is 80 percent. Finally, the project has a 12 percent chance of running aground during the development phase, for a present value loss of $6 million. Thus, the overall value of the R&D project, calculated with decision trees or a probability-weighted NPV, comes to minus $0.1 million. The company chooses not to go ahead.

By contrast, a real option appraisal of the R&D project puts its value at plus $0.4 million, suggesting that the company should undertake the initial research phase by investing $1.5 million. The real option valuation is so much higher than that arrived at through a decision tree or probability-weighted NPV that it changes the recommendation from "do not invest" to "invest." There are several reasons for this.

First, the real option valuation carefully maps out all of the possibilities available to the company, including those not readily apparent in the decision tree. Although there are two important sources of uncertainty, technological and product market, traditional analysis focuses solely on the former, ignoring the latter by focusing too narrowly on the expected value of cashflows. The real option valuation takes into account an important element of flexibility: the possibility of not marketing the drug, even if it passes the safety testing phase, should the revised product market outlook seem gloomy. The higher the uncertainty surrounding potential cashflows, the higher the real option valuation. Traditional approaches ignore this kind of uncertainty and management's ability to respond to it.

The real option valuation also differs from the one arrived at through a decision tree in another way: by varying the discount rate appropriately throughout the tree instead of using a single rate, such as the weighted average cost of capital, it accounts properly for the relative level of risk that different cashflows involve. Where real options represent substantial leverage, this impact on the discount rate may be enormous: not the difference between 10 percent and 12 percent, but rather between 10 percent and 50 percent (or minus 50 percent, depending on the nature of the option).

Thus the real option valuation identifies the optimal course for the company at each stage in the process. The initial recommendation is to undertake the research at a cost of $1.5 million. If the research fails, there is no development phase. Even if it succeeds, the real option valuation demonstrates that it is sometimes unwise to proceed with development, particularly if the drug in question is a mediocre one and the revised market forecast at the end of a year is less optimistic than originally expected. Should the development phase proceed successfully, it is still not always wise to market a drug, for when forecasts turn pessimistic, even a great drug may not cover its costs.

By factoring in all alternatives appropriately, the real option approach uncovers additional value that can change recommendations for projects at several stages in their evolution.

∿ ∿ ∿

To date, most attempts to apply real options to the formulation of corporate strategy have been woefully simplistic. Few senior managers have tried using them for that reason. Yet to a much greater extent than rival techniques, real options can help companies make their way through the maze of technological and market uncertainties that face them when they make their decisions. By ignoring real options, many companies are undervaluing genuine opportunities for investment – and for growth. Q