

How much is flexibility worth?



Thomas E. Copeland and Philip T. Keenan

A lot, if uncertainty is high

But discounting cashflows is the wrong way to calculate it

Instead, use options theory to value management's flexibility to act in the future

HAVE YOU EVER BEEN INVOLVED in a capital investment decision where the net present value calculations proved negative, but the management team decided to go ahead anyway? Or been confronted with a positive NPV project where your intuition warned you not to proceed? Often, it is not your intuition that is wrong, but your time-honored NPV decision-making tools.

But there is another way. Managers can use a different tool: real option value. When a situation involves great uncertainty and managers need flexibility to respond, ROV comes into its own. If the decision you face involves low uncertainty, or you have no scope to change course when you acquire new information later on, then NPV works fine. If not, you will want to know more about what real options are and how to value them.

Below, we compare the main decision-making tools and show why traditional techniques such as NPV, economic profit (EP),* and decision trees are incomplete, often misleading, and sometimes dead wrong. We also look at how real options have been used in several practical situations, drawing on simple examples for illustrative purposes rather than going into the mechanics of valuing complicated real options.†

Real options began to be properly understood in 1973, when Fischer Black, Myron Scholes, and Robert Merton devised rigorous “arbitrage-free” solutions to value them. Applications have proliferated, particularly in securities markets, where the theory held up remarkably well when tested against actual prices. However, from our point of view 25 years later, the assumptions of the Black–Scholes model seem somewhat restrictive when applied to real options.

* Defined as the return on invested capital minus the weighted average cost of capital, multiplied by the invested capital; sometimes known as economic value added.

† For a detailed account of this sort, see L. Trigeorgis, *Real Options: Managerial flexibility and strategy in resource allocation*, MIT Press, Cambridge, Mass., 1996.

The authors wish to acknowledge the contributions of Sam Blyakher, Cem Inal, Max Michaels, Yiannos Pierides, and Dan Rosner.

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.

HOW MUCH IS FLEXIBILITY WORTH?

Put simply, “arbitrage-free” means that securities with exactly the same risk/return profiles should be identically priced. If you can describe the payouts on one risky security and then build a portfolio of other securities with exactly the same payout, the price of both must be the same. If the prices were not identical, arbitrage, or buying the underpriced security and selling the other, would be possible. This simple idea is at the heart of option pricing.

As yet, option pricing has not been much used in the evaluation of corporate investments, for three reasons: the idea is relatively new, the mathematics are complex, making the results hard to grasp intuitively, and the original techniques required the source of uncertainty to be a traded world commodity such as oil, natural gas, or gold. Recent McKinsey research has helped to overcome some of these obstacles. We now know how to value several situations involving real options without using complicated mathematics.

What are real options?

Skip this section if you are already familiar with options. It simply describes what real options are and how to recognize them in a managerial rather than securities setting.

Options are rights without obligations

An option is the right, but not the obligation, to buy (or sell) an asset at some point within a predetermined period of time for a predetermined price.

The first account of a real option is found in the writings of Aristotle. He tells of how Thales the Melesian, a sophist philosopher, divined from some tea leaves that there would be a bountiful olive harvest in six months' time. Having a little money, he approached the owners of some olive presses and bought the right to rent their presses at the usual rate. When a record harvest duly arrived and the growers were clamoring for pressing capacity, he rented the presses to them at above the market rate, paid the normal rate to their owners, and kept the difference for himself – proving for all time that sophism is not only an honorable profession, but a profitable one too.

What is the real option in this story? First of all, Thales purchased the right, but not the obligation, to rent the presses. (He purchased a call option, the right to buy or rent. The opposite is a put option, the right to sell.) Had the harvest been poor, he would have chosen not to rent, and lost only his original small investment, the price of the option.

Thales contracted for a predetermined rental price that in option pricing terminology is called the exercise price. If the market price is higher than the exercise price, the call option is said to be “in the money,” and Thales would

exercise it. If the market price is lower than the exercise price, then the call is “out of the money,” and would not be exercised.

The underlying source of uncertainty in the story was the size of the olive harvest, which affected the market rental value of the presses. As the *value of the underlying variable* increases, so does the value of the option. In other words, the greater the harvest of olives to be pressed, the more valuable Thales’ option to rent the presses will be.

The value of the option also increases with the *level of uncertainty* of the underlying variable. The logic is straightforward. If there is no uncertainty over the size of the olive harvest, which is known to be normal, then the market rental value of the presses will also be normal, and Thales’ option will be worthless. But if the size of the harvest is uncertain, there is a chance that his option will finish in the money. The greater the uncertainty, the higher the probability that the option will finish in the money, and the more valuable the option.

So far we have mentioned three of the five variables that affect the value of the option. It increases with the *value of the underlying variable* and with its *uncertainty*, and it decreases as the *exercise price* goes up. The fourth variable is the *time to maturity* of the option. Thales purchased his option six months before the harvest, but it would have been even more valuable two months earlier, because uncertainty increases with time.

To see why, suppose that Thales has agreed to pay 10 drachmas per hour to rent the presses, and the market rental price is also 10 drachmas. With only one second to go before his option expires, it has no value. But with a month to go, there is a good chance that the market value will rise above 10 drachmas and the option will finish in the money. Therefore, the longer the time to maturity, the more valuable an option is.

Finally, the value of the option increases with the *time value of money*, the risk-free rate of interest. This is because the present value of the exercise cost falls as interest rates rise.

Real options can be easy to overlook

One of the problems in learning how to use real options is that we often don’t know how to recognize them in real-life managerial settings. Here are two examples.

In the 1960s, life insurance companies were vying to sign up baby boomers for whole life policies. A feature often included in the policies was the right to borrow against the cash value of the policy at a fixed rate of interest, say 8 percent. At the time, with interest rates of 3 to 4 percent, this feature didn’t seem important. But the insurance companies were unwittingly giving away

HOW MUCH IS FLEXIBILITY WORTH?

a lifelong call option on borrowing that could be exercised at a fixed interest rate of 8 percent.

That option proved extremely valuable when interest rates soared to double digits in the early 1980s. Suddenly, the baby boomers were able to borrow at 8 percent and invest at 12 percent, while the life companies had to borrow at rates higher than 8 percent in order to honor their contracts. Many were threatened with insolvency simply because they had no idea of the value of the options they gave away.

The second example concerns a manufacturer of jet engines. In this highly competitive industry, the secret is to get your engines onto the wings of aircraft; that done, you have locked up the profits from a 30-year stream of spare parts. What the manufacturer's financial officers did was to buy aircraft and lease them, with their own engines on the wings, to airlines. They also extended lease cancellation options that gave the airlines the right to cancel delivery of the aircraft at any time before delivery and a year after, for only a small penalty payment.

The financial officers wondered how much these cancellation options were worth. Analysis revealed that they were worth on average 83 percent of the value of the engine on narrow-body aircraft, and 19 percent on wide-body aircraft (Exhibit 1). The financial officers were horrified.

What should they do?


Their new understanding of real options showed them that the cancellation option was most valuable to airlines that experienced high variability in demand. They stopped

offering lease cancellation options to these airlines. A year or so later, passenger revenue miles fell steeply throughout the industry. Thanks to its change of policy, the company saved tens of millions of dollars.

Options or bets?

Once you understand what real options are, you begin to realize that they are embedded in a whole range of management decisions. Options are everywhere. But it is important to know the difference between a bet and an option.

A situation where there are real options involves uncertainty about two things. One is the future; the other is the ability of management to respond to what it learns as the picture gradually becomes clearer. If management cannot respond in a material manner to new developments, the situation represents a bet, not an option.



Suppose a company must decide whether to invest \$100 million in constructing a new factory in the face of uncertainty about future demand for the product it will produce. If there are no follow-up investments – if it is an “all or nothing” decision – then management faces a bet, not an option. It can put up the money, roll the die, and win or lose.

But suppose the \$100 million investment can be recast as \$10 million spent now on a pilot project and \$110 million invested in a year’s time to build the factory if it still seems like a good idea. Under this scheme, paying the \$10 million immediately gives management the option to proceed with the \$110 million investment a year from now, provided the pilot produces the right results. If it fails, management has no obligation to proceed with the project.

Even though building the factory costs more this way – \$110 million in present value at a 10 percent cost of capital, rather than \$100 million – it may make good economic sense. In particular, if the uncertainty surrounding demand for the product is high, so that the correct decision may well be *not* to go ahead, the option may be worth much more than the cost of creating it.

How real options capture the value of flexibility

Real options capture the value of managerial flexibility in a way that net present value analysis does not. Consider the example described above. Our aim is to illustrate concepts, not describe the methodology, so we will deliberately simplify the calculation.

Suppose there is a 50 percent chance that after investing \$100 million in the new factory, management will be rewarded with strong sales for many years. Revenues exceed costs, and the factory produces an operating income stream with a present value of \$150 million.

On the other hand, suppose there is a 50 percent chance that demand is poor and the present value of the operating income stream is only \$10 million.

A traditional NPV analysis of this bet would put the expected present value of the future operations of the factory at \$80 million (the average of \$150 million and \$10 million). This is not enough to offset the upfront investment of \$100 million, and the project has an NPV of *minus* \$20 million. It will almost certainly be thrown out.

But rather than invest the whole \$100 million up front, what if management elects to buy the *option* to expand, at the price of \$10 million for the pilot?

There is a 50 percent chance that the pilot succeeds; management responds by building the factory (for \$110 million) and reaping profits of \$150 million, as

HOW MUCH IS FLEXIBILITY WORTH?

before. However, both building the factory and obtaining the profits are delayed a year because of the pilot, so we discount them both back one year at 10 percent to a \$100 million investment and a \$135 million profit, or a net \$35 million gain.

There is also a 50 percent chance that the pilot fails, in which case management halts the project with no further outlay.

The overall value of the project (ignoring for simplicity's sake subtle points about investor risk sensitivity and the degree of correlation between the project and the market) is thus \$7.5 million (the average of \$35 million and zero, minus the upfront \$10 million). Management should indeed proceed with the pilot.

What we have so far is a classic decision tree analysis. So how does real option valuation differ from decision tree analysis, and which is best?

Decision trees and real option valuation are closely related: if you can implement the first, it is not much work to implement the second. Decision tree analysis involves building a tree representing all possible situations and the decisions management can make in response to them. To value a decision tree, one calculates expected cashflows based on their objective probability and then discounts them at some chosen rate – usually the weighted average cost of capital.

Option valuation differs from decision tree analysis in calculating values in accordance with the “no arbitrage” principle, or law of one price. This states that two different investment opportunities that produce the same (equally uncertain) payoffs must be worth the same amount; otherwise, arbitrageurs would buy the undervalued investment and sell the overvalued investment, making a risk-free profit in the process.

The option approach can be interpreted in the decision tree context as modifying the discount rate to reflect the actual riskiness of the cashflows. A call option, for instance, corresponds to a leveraged position in the underlying asset, and is therefore by definition riskier than the asset.* As a result, the appropriate discount rate is considerably higher than the weighted average cost of capital. Moreover, it changes throughout the decision tree depending on how far the option is in or out of the money.

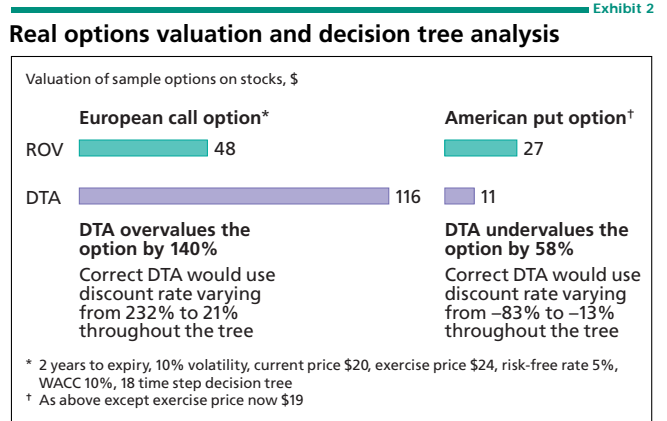
Decision tree methodology gives no guidance on how to choose the discount rate or adjust it for risk or leverage. Traditional decision tree analysis using the

* Imagine a given stock price is \$20 and the exercise price of the call option is \$15. The call option will be worth roughly the difference between the two, namely \$5. If the stock price goes down by \$1, a 5 percent change, the option value will fall by \$1, a 20 percent change. Therefore the option is riskier than the stock.

HOW MUCH IS FLEXIBILITY WORTH?

weighted average cost of capital as the discount rate can thus lead to false results.

To compare decision trees and options on a level playing field, we applied the two approaches to the pricing of financial options – a call and a put on a stock – so that questions of investor risk preferences and how to quantify uncertainty and managerial flexibility would not cloud the comparison. Although in some situations the approaches give similar results, decision trees can be wrong by a factor of two, as Exhibit 2 shows.



Although in some situations the approaches give similar results, decision trees can be wrong by a factor of two, as Exhibit 2 shows.

Why traditional tools are inadequate

Managerial corporate finance theory has been struggling for years with the question of how to evaluate investments under uncertainty. The “obvious” approach of comparing the costs and benefits of an investment is actually highly complex, since costs are incurred today with certainty while benefits are uncertain and reaped in the future. In practice, managers use a range of methodologies including earnings per share or per share growth, economic profit, decision trees, discounted cashflow, and real options.

Until recently, discounted cashflow and economic profit were the most popular approaches. The DCF method forecasts future cashflows, discounts them at a risk-adjusted rate (the weighted average cost of capital), and subtracts the current investment cost to estimate the net present value of a project. Projects with positive NPV are said to create value and are accepted; negative NPV projects are thrown out.

Advocates of this approach point out that it fulfills important criteria: it is cashflow based, risk adjusted, and multi-period or forward looking. However, as we saw earlier, it does not capture flexibility. Exhibit 3 compares the methodologies using these criteria.

DCF techniques were originally developed in order to value investments such as stocks and bonds, and

Exhibit 3

Key criteria for decision-making tools

	Cashflow based	Risk adjusted	Multi-period	Captures flexibility
Real option value	✓	✓	✓	✓
NPV/DCF	✓	✓	✓	✗
Decision trees	✓	✗	✓	✓
Economic profit	✓	✓	✗	✗
Earnings growth	✗	✗	✗	✗

HOW MUCH IS FLEXIBILITY WORTH?

assume that companies hold investments passively. They overlook management’s flexibility to alter the course of a project in response to changing market conditions. In effect, they assume that management makes an irrevocable decision based on its view of the future, and then does not deviate from its plan no matter how things actually shape up. The life of the project is assumed to be fixed, and the possibility of abandoning it in the face of adverse circumstances or, conversely, expanding it in response to unanticipated demand is not even considered. Such rigid assumptions are rather like planning a drive from New York to Los Angeles with only fragments of a map, and then sticking firmly to your original route even as you see highway signs and find out what traffic and road conditions are like.*

Both real options and decision trees capture the mechanics of flexibility. However, only options adjust for risk.

Horses for courses

Real option valuation is most important in situations of high uncertainty where management can respond flexibly to new information, and where the project value without flexibility is near breakeven. (If the NPV is very high, the project will go full steam ahead, and flexibility is unlikely to be exercised. And if the NPV is strongly negative, no amount of flexibility will help.) Optionality is of greatest value for the tough decisions – the close calls where the traditional NPV is close to zero (Exhibit 4).

Consider two investments: a new brewery and a pharmaceutical R&D program. The brewery is a one-off investment in a fairly stable environment in which demand can be forecast reasonably confidently. If the brewery’s operating margins are high, its NPV will be high. The only alternative to going ahead is deferral, which is unlikely since there is little uncertainty about when new capacity will be needed or what the operating margins will be. DCF methods will work well in this setting because their implicit assumptions are valid.

The pharmaceutical R&D program is another matter. Investment is needed at several stages, with substantial outlays on basic research, developmental testing, clinical testing, and product rollout. At each stage, management can choose to abandon the project, defer it, go ahead as planned, or spend more to accelerate

Exhibit 4

When managerial flexibility is valuable

		Uncertainty Likelihood of receiving new information	
		Low	High
Room for managerial flexibility Ability to respond	High	Moderate flexibility value	High flexibility value
	Low	Low flexibility value	Moderate flexibility value

Flexibility value greatest given

1. High uncertainty about the future (very likely to receive new information over time)
2. Much room for managerial flexibility (allows management to respond appropriately to this new information)
3. NPV without flexibility is near zero (if a project is neither obviously good nor obviously bad, flexibility to change course is more likely to be used and therefore more valuable)

Under these conditions, the difference between ROV and other decision tools is substantial

* We thank Sam Blyakher for this example.

it. Flexibility is high, as is uncertainty, and the value of the project without flexibility may be marginal. In this setting, most of the assumptions of DCF do not hold. Option valuation will give much better results.

In 1995–96, the PC assembly business was in turmoil. Gateway was one of the few players making money. Exhibit 5 shows the estimated spread between return on invested capital (ROIC) and weighted average cost of capital (WACC) for a number of major players. Amid tremendous uncertainty about how the industry would shake out, managers had to decide whether to exit or stay as their businesses bled cash. The standard decision tools suggested they should exit immediately. Those businesses with gross operating margins of 11 percent or below have negative DCF and EP values (Exhibit 6).

Exhibit 5
Profitability in PC assembly





Exhibit 6
Different methods, different values



Suppose we focus on a business with an operating margin of 11 percent. Its EP is minus \$90 million, and DCF minus \$590 million. Yet its real option value is \$790 million. Why the huge difference?

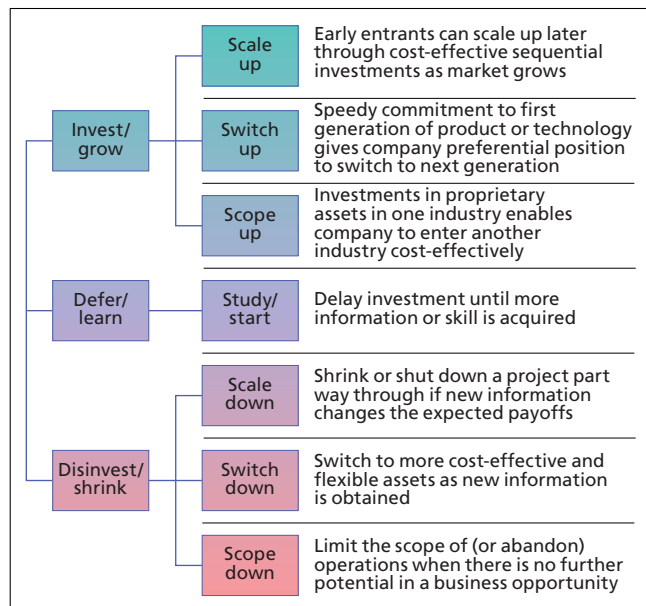
DCF overlooks flexibility. In particular, it ignores the possibility of exiting and reentering the industry. If the costs of doing so were trivial, the best strategy would be to exit immediately and reenter if conditions improve. But exit and reentry costs are high. Managers will thus decide to stay in an industry despite a negative DCF because there is tremendous uncertainty about what operating margins will be like in the future, and because exit and reentry costs are high. ROV not only captures the value of flexibility, but also indicates how long a company should stay in a business before exiting, and when to reenter.

Classifying real options

Individual real options can be classified into growth options (scaling up, switching up, or scoping up a project), deferral/learning options, and abandonment options (scaling down, switching down, or scoping down a project) (Exhibit 7).

Exhibit 7

7S framework: Grow, defer, or quit



The seven basic real options can also occur in combinations, as compound options. A company that invests in an R&D project, say, may be buying both the option to commercialize the resulting product and the option to engage in subsequent R&D projects to develop future generations of related products. These subsequent R&D projects themselves contain options for commercialization and further development, leading to a type of compound option called a growth staircase.

Real options can also depend on more than one source of uncertainty. The value of an option to commercialize an R&D project, for instance, depends on at least two: technological uncertainty (will the scientists succeed in inventing the new product?), and market uncertainty (what will the demand for this new product be?). Options that depend on multiple sources of uncertainty are often called rainbow options.

A deferral option (natural resource development)

Real options played an important role in a decision made by one coal-mining company. It needed to work out how much to bid for the lease of a plot of land that could be developed into a coal mine. Using the current price of coal and extrapolating its growth into the future, and forecasting extraction costs, taxes, and the estimated quantity of coal in the mine, the company calculated the NPV of the cashflows involved in developing the mine and selling the coal, and concluded that the lease was worth \$59 million – not very much.

However, the company knew that the price of coal could fluctuate substantially. As its current price was close to the project’s breakeven point, the revenue projections were highly sensitive to future price changes.

The company realized that acquiring the lease would give it an option: to defer opening the mine until such time as the price of coal rose far enough to make the project’s economics reasonable. This option turned out to be worth \$57 million – almost as much as opening the mine immediately. When the value of the option was factored in, the lease was actually worth \$116 million.

The company successfully bid \$72 million, waited until the price of coal rose, and made a tidy profit from the mine.

Learning options (natural resource development and R&D)

Learning options arise when a company is able both to spend money to speed up its acquisition of important information (for example, to reduce technological uncertainty in R&D or to learn more about the quantity of resources in the ground in exploration and development projects) and to use what it has learned about the market demand for the project output to modify future investment decisions. It must balance the value of the option to act on the information learned against the cost of acquiring that information.

Take the company considering a site for a mine. It must weigh the value of deferring development of the mine against the value of the information it could gain about the quantity of ore in the mine as a result of full or partial development. This is an example of a rainbow option. The sources of uncertainty are the price of coal and the quantity of coal in the mine.

Deferral is valuable because of the uncertainty surrounding the price of the coal the mine will eventually sell. If the economics are presently near breakeven, waiting gives management the chance to react to price shifts. However, partial development is also valuable because it reduces uncertainty about the quantity of coal in the mine, while preserving management's ability to adjust future investment according to what is learned. Partial development thus represents a learning option that is in conflict with the deferral option; the company cannot exercise both.

Multi-stage R&D projects generally contain a series of embedded options based on technological and market uncertainty. They too are learning options. Undertaking an R&D project gives management the right, but not the obligation, to commercialize a product if and when the R&D effort is successful and the economics of producing and marketing the product are attractive. Although an R&D project viewed in isolation may have a negative NPV, the option to commercialize the result is often extremely valuable – enough to determine that the project be undertaken anyway.



Making irreversible investment decisions in the face of uncertainty is risky. Being able to change a decision as new information becomes available helps reduce the risk. Traditional decision-making tools such as NPV, EVA, and earnings per share neglect the value of such flexibility. Real options, on the other hand, provide a theoretically sound tool for valuing management's strategic scope. Recent advances in theory have made ROV techniques applicable to a multitude of real-world situations. At the same time, advances in technology have enabled option pricing capability to move out of Wall Street and into mainstream corporations. 