Published in:: G. Succi, M. Marchesi, L. Williams, D. Wells: Extreme Programming Perspectives, Addison Wesley, 2002

Chapter 43

Keep Your Options Open: Extreme Programming and the Economics of Flexibility

Hakan Erdogmus and John Favaro

Financial evaluation and strategic analysis have long been considered two distinct approaches to evaluating new capital initiatives. An emerging valuation approach, known as real options, attempts to align finance and strategy through a new perspective: The value of an asset lies not only in the amount of direct revenues that it is expected to generate, but also in the options that it creates for flexible decision making in the future. In general, the more uncertain the future is, the higher the value of flexibility embedded in an asset, whether financial or real. This perspective has significant implications for the economics of flexible processes. Applied to software development, it could imply that a lightweight process that is well positioned to respond to change and future opportunities creates more value than a heavy-duty process that tends to freeze development decisions early. Thus, the feasibility of Extreme Programming (XP) can be supported by the option value of flexibility inherent in it. What is the theory that underlies this statement? How does it relate to the fundamental assumptions of XP? How does it impact the value of an XP project? What are the implications of such value propositions for project decisions? If you are curious, read on …

Introduction

Change: Ally or Enemy?

Kent Beck, during a workshop on Extreme Programming (XP) for capitalists, provoked the audience by proposing that XP could create ten times more value than a heavyweight process. How can this ever be possible? Consider the two fundamental premises of XP.

A. Change is inevitable. Just about the only thing you can predict with some certainty is that *change will happen*. In *Extreme Programming Explained*, Beck emphasizes this point.

> *Everything in software changes. The requirements change. The design changes. The technology changes. The team changes. The team members change. The problem isn't change per se, because change is going to happen; the problem, rather, is the inability to cope with change when it comes.* [Beck1999]

B. Change is easy. The cost of change does not rise exponentially as the system grows. Contrary to popular belief, the rise in the cost of change gradually diminishes.

We don't question premise A. We take it as a given.

> *XP is a lightweight methodology for small to medium teams developing software in the face of vague or rapidly changing conditions.* [Beck1999]

Premise B is more controversial. We don't know whether it's true or whether it is universally true. We don't know whether it is a consequence of the 12 XP practices or of the advancements in software practice and technology in general. Thus, we will condition our conclusions and insight on the truth (or falsity) of premise B. For the time being, let's take it as a given as well.

Consider the following XP principles and practices:

1. Embracing change
2. Simple design
3. Small initial investment
4. Incremental change
5. Small releases
6. Continuous refactoring

How does one get from the premises A and B to the principles and practices 1 through 6? At a gut level, if change is inevitable, naturally the best way to manage it would be to embrace it. It all seems to make sense, but the cause-effect relationships between the premises and the resulting principles and practices of XP, as well as among the principles and practices themselves, are more subtle and complex than they first appear. True, a flattened cost curve would make 1 through 6 possible. But why would it also ultimately make XP more profitable? True, investing in a complex design would not make sense under highly volatile and vague requirements. But wouldn't this argument hold under an exponential cost-of-change curve even more strongly than it does under a flat cost-of-change curve? So again, how does XP create more value here?

The answer lies behind a crucial characteristic of XP: *flexibility*. Change is driven by uncertainty. At the heart of any process designed to cope with uncertainty is flexibility. Embracing change means treating uncertainty as an ally, rather than viewing it as an enemy. Embracing change means embracing flexibility. Most of the principles and practices, indeed most things that are fundamental to XP, can be one way or another traced back to flexibility. And flexibility creates value under uncertainty. The more uncertainty there is, the more value it creates.

Published in:: G. Succi, M. Marchesi, L. Williams, D. Wells: Extreme Programming Perspectives, Addison Wesley, 2002

What Is It With Flexibility, Anyway?

> *Flexibility can be viewed as an option.*
>
> —Nobel Prize Lecture in Economics, 1997

This simple yet provocative statement, made during the conferral of the most prestigious prize in economics, forms the point of departure for the discussion of the economics of XP.

Let's begin by considering a simple example of flexibility: a fully refundable plane ticket. Such a ticket gives its holder the flexibility to recover the full cost of the ticket in case of an unexpected event. In other words, the customer has the *option* to exchange the ticket for its cost on or before the travel date should such an event occur. The flexibility provided by the option is desirable if the future is uncertain—the more uncertain the future is, the more desirable the flexibility is. The customer can think of the ticket as a risky asset: If such an event occurs, without the flexibility, the ticket will be worthless; if everything goes well, the ticket will preserve its value.

A refundable ticket costs more than a nonrefundable ticket. Why? Because customers are willing to pay for the additional flexibility, which protects them in case of a negative development, or if they simply change their mind. The airline company demands a premium for this option over the price of a nonrefundable ticket because by offering a full refund, it risks flying with an empty seat, and incurring a loss as a result. Customers, by agreeing to pay for the additional flexibility provided by the refundable ticket, implicitly believe that the value of the option, with respect to the amount of uncertainty they are facing, is comparable or superior to the premium demanded by the airline.

Options, Options Everywhere

Options arise everywhere in the business world. Here are some more concrete examples from software development.

- A pioneering Internet security project with a follow-on opportunity in the growing e-business market: Undertaking the pilot creates the option to be a player in an emerging market. This is an example of a *growth option* [Benaroch+2000; Favaro+1999; Taudes1998].
- Development of a framework for a future product line: The infrastructure investment enables efficient generation of a multitude of closely related applications without committing to a particular one. This is an example of a *platform option* [Erdogmus2001B; Favaro+1998B].
- Abandoning a staged migration project midstream when budget overruns overtake the expected benefits: The ability to stop adds value proportional to

the losses that would be incurred with continuing. This is an example of an *exit option* [Erdogmus+1999; Favaro+1998B].

- Developing a prototype before the full application to resolve technical and user uncertainty: Investing first a small amount to learn reveals the feasibility of the larger investment. This is an example of a *learning option* [Sullivan1996].
- Waiting to see whether the Java technology gains acceptance before migrating a stable application to Java. Waiting before committing may be a cheap way to learn. This is an example of a *delay*, or *timing*, *option* [Benaroch+1999].

These strategic options, both technical and business-driven, commonly arise in the general software industry, but the topic of the discussion is Extreme Programming. What is the relationship of XP to the concept of such options?

XP as an Options-Driven Process

> *We need to make our software development economically more valuable by spending money more slowly, earning revenue more quickly, and increasing the probable productive lifespan of our project. But most of all, we need to increase the options for business decisions.* [Beck1999]

Beck's declaration makes it clear that the founders of XP also believe in the importance of creating business options—and believe that XP is capable of creating them. Some examples that should be familiar to the XP practitioner are the following:

- Checkpoints after every iteration, where the customer can make midcourse decisions
- Talented, trained personnel able to switch course rapidly with new or modified stories
- The ability to modify project at a small cost through enabling technologies and best practices
- Waiting to see whether the customer really wants a feature before implementing it

The 12 practices and four values of XP are also a fertile source of business options.

- *Small releases* introduce decision points and opportunities to change course. At the end of a release, the customer has the option to continue, modify the course of the project, or stop based on what has been learned from the previous releases. This flexibility increases the value of the project while reducing its risk.
- *Refactoring* makes future options to modify the system more valuable by keeping the cost of change at bay.
- *Collective ownership* increases the chances of an option to improve the system to be exercised in a timely manner, which in turn increases its value.

- *Continuous integration* preserves business value. Anytime, you can stop and still have a working system that can be delivered to the customer with some inherent value. The option to exit is more valuable for the customer because of this salvage value.
- *Simplicity* creates options to modify the system. Complex code and design ossify the system. The simpler the code, the easier to modify it, and the higher the resulting option value.
- *Communication*, *feedback*, *pair programming*, *on-site customer*, and *testing* all help reveal information and resolve uncertainty. When uncertainty is not resolved, options cannot be exercised in a rational and timely manner, destroying the value of flexibility.
- *Courage* is required to exercise the options created. Without it, the options are worth nothing. Without courage, options virtually don't exist. Conversely, courage is also required to let go when it is revealed that an existing option is no longer likely to create business value.

All of these points make a convincing argument that XP is a powerful options-driven process, capable of generating significant value. However, we also need an economic foundation for analyzing *why* and *how much* value is created by options. This brings us to one of the central activities of finance: valuation.

Valuation Basics: How to Quantify the Value of Capital Investments

Valuation is the process of estimating how much an asset is worth. An XP project is subject to the same fundamental principles of valuation as any other real asset, as summarized this way:

> *By adding up the **cash flows** in and out of the project, we can simply analyze what makes a software project valuable. By taking into account the effect of **interest rates**, we can calculate the **net present value** of the cash flows. We can further refine our analysis by multiplying the **discounted cash flows** by the **probability** that the project will survive to pay or earn those cash flows.* [Beck1999]

That single paragraph contains references to most of the fundamental principles of valuation (shown in boldface), so let's pick it apart now. Comprehensive coverage of the subject is beyond the scope of this chapter. In what follows, we provide an overview of only the most basic concepts as they relate to the current discussion. Suggestions for further reading are provided at the end of the chapter.

In finance, the costs and benefits associated with an investment are called *cash flows*. Investments are compared only on the basis of their cash flows. Cash flows are often represented in tabular form according to chosen time periods—for example, in years, quarters, or months. Usually, there is an original investment, C_0, represented as a negative number. Subsequent cash flows are denoted as C_1, … C_n, spanning the time horizon in which the investment incurs costs and generates benefits.

Discounted Cash Flow and Net Present Value

The *present value* (PV) of a future cash flow is the value of the cash flow as though it were received today. How does one calculate the present value of a future cash flow?

Moving forward from present to future, an investment is expected to grow at a certain rate of return. Now turn it around: Moving backward from future to present, an investment shrinks with the same rate of return.

When moving back in time, the rate of backward adjustment is called the *discount rate*. The process of backward adjustment itself is called *discounting*. The general technique of valuing a capital investment project by summing its discounted future cash flows is known as *discounted cash flow* (DCF).

The DCF calculation doesn't usually include the initial investment C_0. When that initial investment is included (represented as a negative cash flow), the *net present value* (NPV) is obtained: the benefits minus the costs. All of this is expressed in the following simple formula:

$$\text{NPV} = C_0 + \frac{C_1}{(1+k)} + \frac{C_2}{(1+k)^2} + \ldots$$

Here, the C's represent the cash flows, the subscripts represent the periods in which the cash flows occur, and k is the per-period discount rate. The NPV formula tells us whether the investment is worth more than it costs. The rule is that if NPV is positive, the investment is worth undertaking—it generates more value than it costs. If it is negative, it should be forgone—it generates less value than it costs. If it is zero, we are indifferent between undertaking and forgoing it.

A Valuation Example

We illustrate DCF and NPV in action by considering the development scenario illustrated in Figure 43.1. The horizontal line represents the time horizon extending to five years out. The outgoing arrows represent negative cash flows, or expected costs including the initial investment (development cost) and the subsequent investments (maintenance costs). The incoming arrows represent positive cash flows, or expected benefits from sales revenues. The discount rate is given as 7% annual.

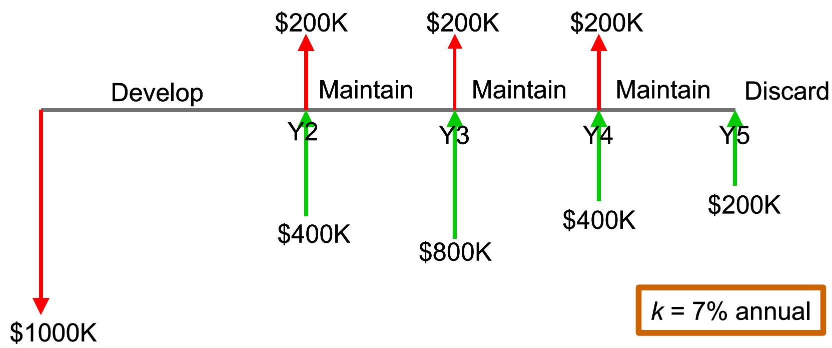

Figure 43.1: Cash flows of a software development project

The straight net value of the investment is calculated simply by summing all the cash flows. In thousands of dollars, the calculation is as follows:

Straight Net Value = –1000 – (3 × 200) + 400 + 800 + 400 + 200 = **200**

According to this result, the net value is positive, and the project should be undertaken. However, the DCF approach yields a very different conclusion. The net present value, calculated using DCF, is as follows:

$$\text{NPV} = -1000 - \left(\sum_{t=2..4} \frac{200}{(1+0.07)^t} \right) + \frac{400}{(1+0.07)^2} + \frac{800}{(1+0.07)^3} + \frac{400}{(1+0.07)^4} + \frac{200}{(1+0.07)^5}$$

$$= \mathbf{-40}$$

The negative result tells us that the project is not worth undertaking. The cost of the investments exceeds the return on investment expected from the project. Therefore, if undertaken, the project would destroy value rather than create it.

It's All about Risk

Risk management is taken very seriously in XP. Clearly, then, risk has to be taken into consideration in any economic valuation of an XP project.

Software engineers have an intuitive view of risk that is related more to project management, even to sociology or psychology, than to finance. XP is no exception. Usually, risk is characterized by what can go wrong in the project, and the strategies for dealing with this problem have been limited to implementing the riskiest artifacts first. From the financial point of view, however, risk has a much more precise, well-defined meaning. Financial risk refers to the variability in the returns of an asset [Ross+1996]. It has two components:

- Systematic component—Market risk, a.k.a. systematic risk or nondiversifiable risk
- Unique component—Private risk, a.k.a. unsystematic risk or diversifiable risk

Private risk corresponds to the traditional software engineering view of risk. However, no business works in a vacuum—all businesses participate in a market and are affected by systematic risks that permeate the system in general and the sector in which they operate in particular. These systematic risks range from the overnight bank loan rate determined by the Federal Reserve Bank (in the United States) to the outbreak of war.

Table 43.1 contrasts the well-known risks identified for an *individual* XP project with the market risks that affect *many* projects. Market risks are often easier to tackle because they are priced by financial markets. Those are the risks that well-diversified investors are mainly worried about because diversification can minimize, if not completely eliminate, private risk.

Table 43.1: Private Risk Versus Market Risk in XP Projects

Private	**Market**
• Project canceled • System goes sour • Business misunderstood • Business changes • False feature-rich • Schedule slips • Staff turnover • Defect rate • Technology	• How much will the clients be willing to pay? • How much will skilled programmers cost? • How uncertain are fixed costs? Overhead? • How well is the economy doing? • Where are the short-term interest rates heading?

Both market and private risk can figure into the simple NPV equation. When cash flows are estimated, effectively private risk must be taken into account. If things go well for the project, more will be earned. If things go badly, the cash flows will be smaller. So the private risk is accounted for in the *unbiased* estimates of cash flows in the *numerator* of a DCF term in the NPV equation. An unbiased estimate of a cash flow is calculated as a statistical expectation by considering as many scenarios as is feasible and the respective likelihood of these scenarios.

In contrast, market risk is accounted for in the *denominator* of a DCF term, by adjusting the discount rate. The higher the market risk, the higher the discount rate. Figure 43.2 illustrates how the NPV equation accounts for private and market risk.

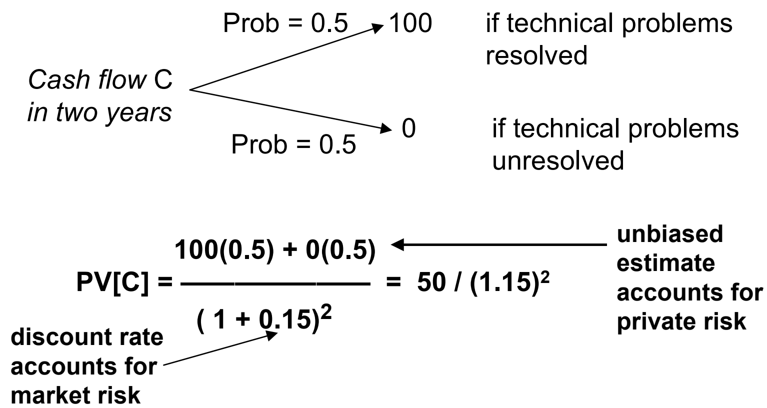

Figure 43.2: Accounting for private and market risk in the NPV equation

Corporations have developed practices to determine discount rates based on the returns of past projects and grouping of like projects into risk categories. In addition, many organizations specialize in the estimation of market risk, examining the historical returns of companies and deducing the amount of market risk borne to develop projections.

From Traditional Valuation to Real Options

Discounted cash flow is the foundation of modern valuation. It provides a method for capturing the value over time of operational benefits and costs associated with any investment so long as those benefits and costs can be cast in currency terms. We have seen that DCF techniques can deal with both project-specific, or private, risk (through unbiased expected cash flow forecasts) and systematic, or market, risk (through a suitably adjusted discount rate) associated with these operational costs and benefits.

DCF alone, however, is not sufficient for capturing *all* value inherent in a project. DCF can be used to evaluate the operational benefits from *business as usual*, often the case in a stable environment with well-understood and measurable costs and benefits, but it has little to offer to capture additional business value due to flexibility under uncertainty, such as strategic opportunities, learning, and the ability to respond to changing conditions.

This orthogonal dimension of value generation requires techniques that can explicitly model *active management*. Although DCF works well only for deterministic projects with a linear timeline, projects that can be represented by a linear stream of expected cash flows, it does not work well for projects with future decisions that depend on how uncertainty resolves—for example, XP projects. For this purpose, we must turn to the intuition and more powerful techniques offered by the theory of option pricing.

Option Basics

In its most general form, an option refers to a future discretionary action. Financial options have been traded for centuries. They date back to seventeenth-century Holland, where tulip options were common. Investors bought options to buy and sell yet-to-be-developed tulip varieties.[1]

Options are a form of *derivative* [Hull1997]. The value of an option—that is, the price to be paid to acquire the option or the value it adds to an existing portfolio of assets—depends on the value of an underlying asset. For financial options, the underlying asset can be a stock price, an exchange rate, or a commodity spot price. For *real options*, the underlying asset is a real asset, typically a stream of future cash flows.

A large body of jargon is associated with the options trading industry. Fortunately, we need only the most basic terminology in this chapter. A *call option* refers to the right, without a symmetric obligation, to buy a risky asset at a preset price—called the *strike price* (a.k.a. *exercise price* or *exercise cost*)—on or before a future date, called the *expiration date* (a.k.a. *maturity date*) of the option.[2]

Figure 43.3 illustrates how an option works with a simple example. Consider a call option on a stock whose *current price* is $50, with an *expiration date* after six months, at a *strike price* of $45. Now let's consider two cases, where the stock price either goes up, to $75, or down, to $35, in six months. If the stock price goes up, the holder of the option exercises the option by buying the stock for $50 and selling it at its market value of $75, making a profit of $30. Otherwise, the holder of the option does nothing, and the option expires worthless. Thus, the option is worth either $30 or nothing at maturity.

---

1. Before long, the practice led to enormous speculation and a spectacular crash. To this day, speculative bubbles, such as the market crash in dot-com stocks in 2001, are commonly referred to as Tulip Mania.

2. The opposite is a *put option*, which refers to the right to sell an asset at a preset price on or before a future date.

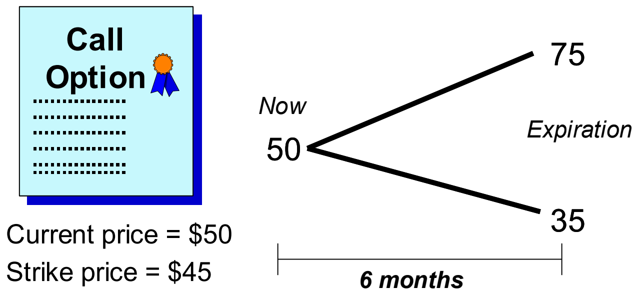

Payoffs at expiration	If stock rises to $75	If stock falls to $35
(Stock price) – (Strike price)	$75 – $45 = $30	$35 – $45 = –$10
Option value at expiration	**$30.00**	**$0.00**

Option pricing: How much should I pay to acquire this option now?

Figure 43.3: Call option example

Five parameters determine the value of a call option, as shown in Figure 43.4. The arrows next to each parameter indicate whether a higher value of that parameter *increases* or *decreases* the value of the option.

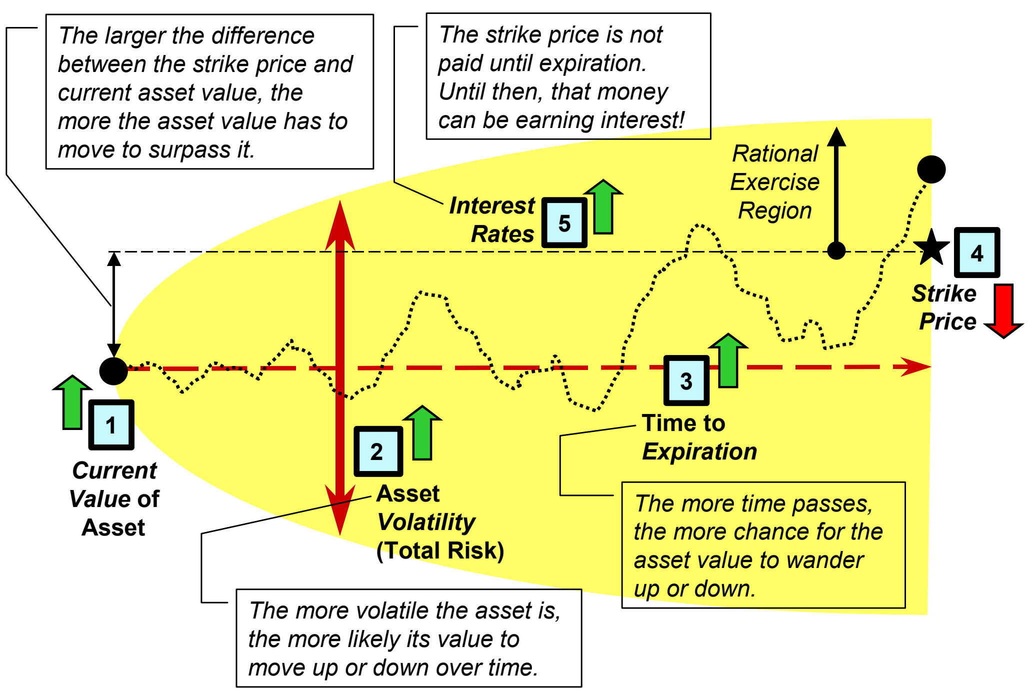

Figure 43.4: Five parameters determining the value of a call option

Rational Exercise

The holder of a call option exercises the option only if the price of the underlying asset is above the strike price (the upper straight line in Figure 43.4), to avoid a loss. This practice constitutes a fundamental assumption of option pricing, called *rational exercise*.

The rational exercise assumption is behind the behavior of an option's value in response to changes in volatility and the expiration date. As the volatility (total risk) of an asset and time horizon increases, the tendency of the asset's value to move away from its initial value also increases. Rational exercise prevents such an increased tendency to affect the maximum loss, thereby limiting downside risk, but without a symmetrical restriction on the size of the payoff in the case of a positive development.

From Financial Options to Real Options

So far the discussion has focused on the *financial* world of stocks and options; but the main interest of this chapter is in the *real* world of projects—and processes that drive them. How do we make the leap from one to the other?

The term *real options* was coined in 1977 by Stewart C. Myers of the Massachusetts Institute of Technology (MIT), who first realized that financial option pricing techniques could be applied to the evaluation of projects. The essence of his insight is illustrated in Figure 43.5. The figure maps the five parameters affecting a financial option's value shown in Figure 43.4 to the analogous factors in real-world projects.

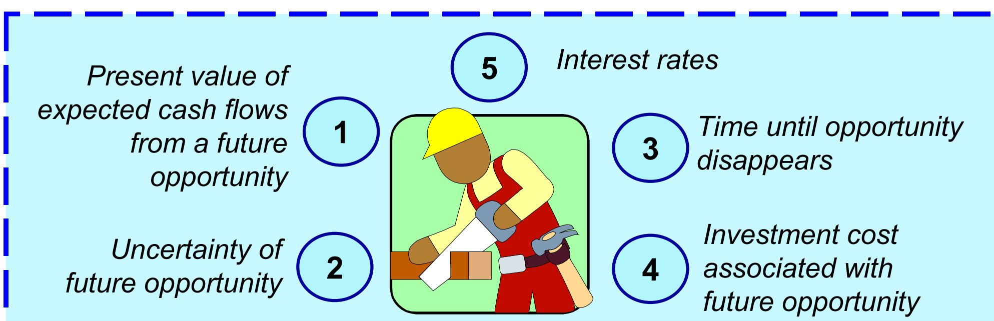

Figure 43.5: Analogy between financial and real options

The analogy is not quite as simple as Figure 43.5 suggests. In the financial world, parameters such as the current price of a stock (parameter 1 in Figure 43.4) and its uncertainty (parameter 2 in Figure 43.4) are determined by the markets. In the world of real assets, however, it is usually necessary to estimate them through other means—often without much information to work from.

A full treatment of the relationship between financial options and real options is outside the scope of our discussion, but the main differences are summarized in Table 43.4 at the end of the chapter. For the purposes of this chapter, we take the analogy for granted and move on to the application of options to XP.

XP and Options

A basic understanding of how options work helps us understand how some of the basic XP value propositions can be justified using the fundamental tenets of XP. In this section, we examine two of these value propositions.

- *Proposition 1:* Delaying the implementation of a fuzzy feature creates more value than implementing the feature now.
- *Proposition 2:* Small investments and frequent releases create more value than large investments and mega-releases.

We begin with the technical premise of XP and its relation to the first proposition. Then we tackle the second proposition in the context of staged investments and learning. The option pricing models used to analyze each scenario are introduced just in time along the way.

The Technical Premise of XP

> *The software development community has spent enormous resources in recent decades trying to* ***reduce the cost of change*** *– better languages, better database technology, better programming practices, better environments and tools, new notations … It is* the *technical premise of XP.* [Beck1999]

XP challenges one of the traditional assumptions of software engineering: that the cost of changing a program rises exponentially over time, as illustrated inFigure 43.6.

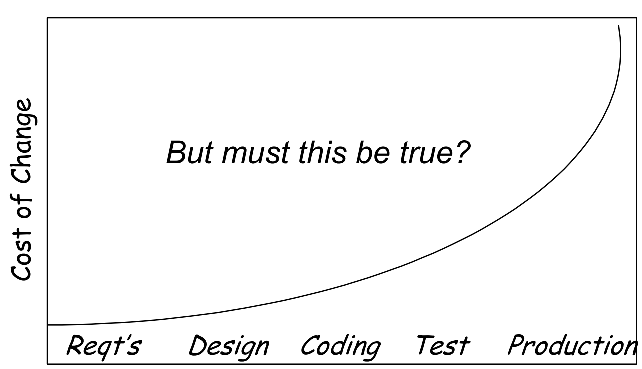

Figure 43.6: The traditional assumption about the cost of change

The technical premise of XP is that this pathological behavior is no longer valid. Better technologies, languages, practices, environments, and tools—objects, database technologies, pair programming, testing, and integrated development environments come to mind—all help keep software pliable. The result is a cost-of-change function that resembles the dampened curve in Figure 43.7.

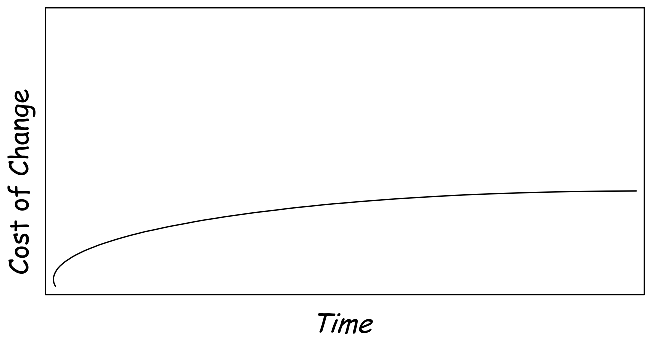

Figure 43.7: The technical premise of XP

Why is a flattened cost curve important for an options-driven process? A flattened cost curve amplifies the impact of flexibility on value. It does so by creating new options that would not have existed under an exponential cost function, and by reducing the exercise cost, and therefore increasing the value, of existing options.

You Aren't Going to Need It: Now or Later?

> *We are traditionally told to plan for the future … Instead, XP says to do a good job … of solving today's job today … and add complexity in the future where you need it. The economics of software as options favor this approach.* [Beck1999]

One of the most widely publicized principles of XP is the You Aren't Going to Need It (YAGNI) principle. The YAGNI principle highlights the value of delaying an investment decision in the face of uncertainty about the return on the investment. In the context of XP, this implies delaying the implementation of fuzzy features until uncertainty about their value is resolved. YAGNI is a typical example of *option to delay*, an all too common type of a real option.

Extreme Programming Explained provides an example of the application of options theory to YAGNI.

> *Suppose you're programming merrily along and you see that you could add a feature that would cost you $10. You figure the return on this feature (its present value) is somewhere around $15. So the net present value of adding*

> *this feature [now] is $5. Suppose you knew in your heart that it wasn't clear at all how much this new feature would be worth—it was just your guess, not something you really knew was worth $15 to the customer. In fact, you figure that its value to the customer could vary as much as 100% from your estimate. Suppose further that it would still cost you about $10 to add that feature one year from now. What would be the value of the strategy of just waiting, of not implementing the feature now? … Well, at the usual interest rates of about 5%, the options theory calculator cranks out a value of $7.87.* [Beck1999]

The scenario is illustrated in Figure 43.8.

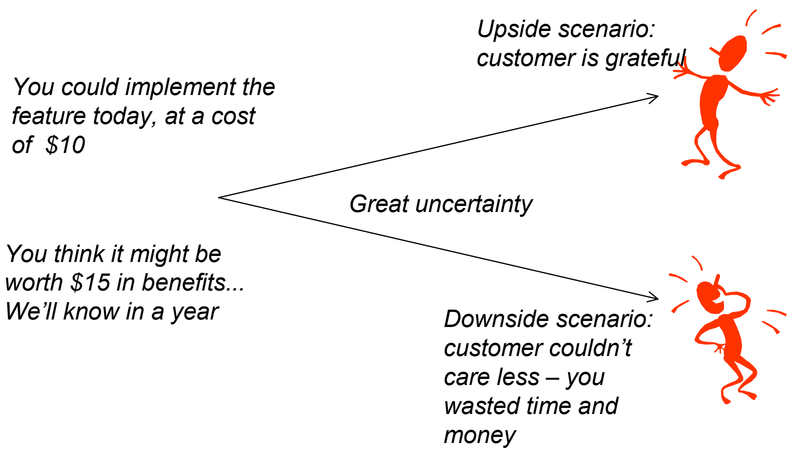

Figure X.1 YAGNI Scenario

The delay option underlying the YAGNI scenario is much akin to a financialcall option, an option to acquire a risky asset on a future date. We will analyze the scenario using the famed Black-Scholes formula for calculating the value of a call option on an uncertain asset—the same formula used by Beck in *Extreme Programming Explained*. To understand in what way the cost of change affects the value proposition underlying the YAGNI scenario, we need to dig a little deeper into the option pricing theory.

Option Pricing 101

Three financial economists, Fisher Black, Myron Scholes, and Robert Merton, undertook the groundbreaking work on option pricing in the early '70s. Their efforts won them a Nobel Prize in economics in 1997. The equation published in a seminal paper in 1973 on the pricing of derivatives and corporate liabilities became known as the Black-Scholes formula [Black+1973]. The Black-Scholes formula revolutionized the financial options trading industry. Both the theory and the resulting formula in various forms are now widely used.

The Black-Scholes equation is illustrated in Figure 43.9. In the equation, C denotes the value of a call option on a non-dividend-paying asset with a strike price of L. M is the current value of the underlying asset, the asset on which the option is written. The option expires at time t. The risk-free interest rate is denoted by r_f, expressed in the same unit as t. The risk-free rate is the current interest rate on the risk-free asset, such as a short-term Treasury bill or government bond. Its value can simply be looked up in the business section of a daily newspaper. $N(.)$ is the cumulative normal probability distribution function, and "exp" denotes exponential function.

$$C = N(d_1) \times M - N(d_2) \times L \exp(-r_f t)$$

$$d_1 = \frac{\ln\left(\frac{M}{L}\right) + t\left(r_f + \frac{\sigma^2}{2}\right)}{\sigma\sqrt{t}}$$

$$d_2 = d_1 - \sigma\sqrt{t}$$

- C Value of call option on risky asset
- L Strike (exercise) price
- t Time to expiration
- r_f Risk-free interest rate
- M Current price of asset
- σ Volatility (standard deviation of asset's rate of return)
- $N(.)$ Cumulative normal probability distribution function

Figure 43.9: The Black-Scholes formula for the value of a call option

The parameter σ denotes the *volatility* of the underlying asset. Volatility is a measure of total risk, which subsumes both market and private risk. It is given by the standard deviation of the continuous rate of return on the asset's price (value) over time. Usually, this parameter is estimated using historical data. For a stock option, volatility can be estimated by calculating the standard deviation of the stock's past returns over small intervals spanning a representative period—for example, using weekly returns over the past 12 months. For real assets, estimation of volatility is much more tricky, but sometimes market data can still be used. An example from software development is provided in [Erdogmus2001B].

The parameters M (the current value or price of the underlying asset), σ (the volatility of the underlying asset), t (the option's time to expiration), L (the option's strike price), and r_f (the risk-free interest rate) correspond to the five standard parameters of option pricing illustrated in Figure 43.4.

How did Black, Scholes, and Merton invent this magic equation? All earlier attempts at solving the option pricing problem involved calculating the net payoff of the option at expiration under the rational exercise assumption and then discounting this payoff back to the present to determine its current value. This approach required identifying the proper discount rate for the uncertain payoff. Essentially, the risk of an option is different from, and often much higher than, the risk of its underlying asset. Even if the discount rate for the underlying asset is known, choosing the proper discount rate for all possible payoffs of the option under different exercise scenarios is inherently problematic. Black, Scholes, and Merton succeeded not by *solving* the discount rate problem, but by *avoiding* it. Their solution is based on two key concepts:

- Replicating portfolio
- The law of one price, also known as no arbitrage

The first concept, *replicating portfolio*, states that the behavior of an option can be replicated by a portfolio consisting of a certain amount of the underlying asset and a risk-free loan to partially finance the purchase of the underlying asset. Thus, it is not necessary to buy options—one can create a *do-it-yourself* or synthetic option through a combination of the underlying asset and a loan. The option is then effectively equivalent to a *levered* position in the underlying asset. Indeed, the idea of financial leveraging has been known for a long time: Buying on margin has been widely practiced, especially during the stock market boom of the ’90s.

The second concept, the *law of one price*, or *no arbitrage*, states that an efficient market lacks money machines. If one can replicate the behavior of an option exactly by a corresponding portfolio, the portfolio and the option are interchangeable for all practical purposes and thus must be worth the same. The two assets—the option and the replicating portfolio, with exactly the same payoffs under the same conditions—must have the same price. If the exact composition of the replicating portfolio, and therefore how much it is worth in the present, can be determined, then how much the option is worth in the present will also be known. Option pricing problem solved!

The original derivation of the Black-Scholes equation is based on solving a specific stochastic differential equation in continuous time. Cox, Ross, and Rubinstein provide a much simpler derivation originating from a discrete model [Cox+1979], which we will also take advantage of later in the chapter. In the YAGNI example, we will stick with the Black-Scholes model.

Evaluation of the YAGNI Scenario

Table 43.2 illustrates the application of the Black-Scholes formula to the YAGNI scenario.

Table 43.2: Calculation of the Option Value of the YAGNI Scenario

B-S Variable	**Value**	**Explanation**
M	15.00	B-S: Current price of underlying asset *YAGNI: PV of benefits from proposed, deferrable feature implementation*
L	10.00	B-S: Strike (exercise) price of the call option *YAGNI: Cost of implementing proposed, deferrable feature*
r_f	0.05	B-S: The risk-free rate of return *YAGNI: The opportunity cost of implementation; the return that the implementation cost would earn if invested in a risk-free security*
t	1.00	B-S: Years until expiration of the option *YAGNI: Date on which feature implementation decision must be taken*
σ	1.00	B-S: Volatility of the underlying asset (standard deviation of the asset's rate of return) *YAGNI: Volatility of the feature's benefits (the standard deviation of the return of feature's benefits)*
C	**7.87**	B-S: Value of Black-Scholes call option *YAGNI: Value of waiting one year before implementing the feature*

The NPV of implementing the feature now is $5 (the $15 present value of expected benefits, minus the $10 cost of implementation). If the implementation is deferred one year, at a volatility of 100%, the Black-Scholes model yields an option value of $7.87, provided that the cost of implementation stays the same. Because deferring implementation incurs no initial cost, the option value equals the NPV of waiting a year before deciding whether to implement the feature. This value takes into account the possibility that the planned feature may be worthless in one year, which would force its implementation to be forgone, as well as the possibility that the actual benefit of the feature may very well exceed today's estimate of $15 (because of uncertainty), which would make the feature a much more profitable investment. The flexibility of deferring the decision increases the value created, because it helps limit the downside risk of the investment without a symmetric limitation on its upside potential.

Uncertainty is a key factor in this example. Figure 43.10 illustrates how the value created by waiting in the YAGNI scenario varies in response to the level of uncertainty, everything else being equal. Uncertainty is captured by the volatility of the feature's benefit. As the volatility increases, the option value of waiting also increases.

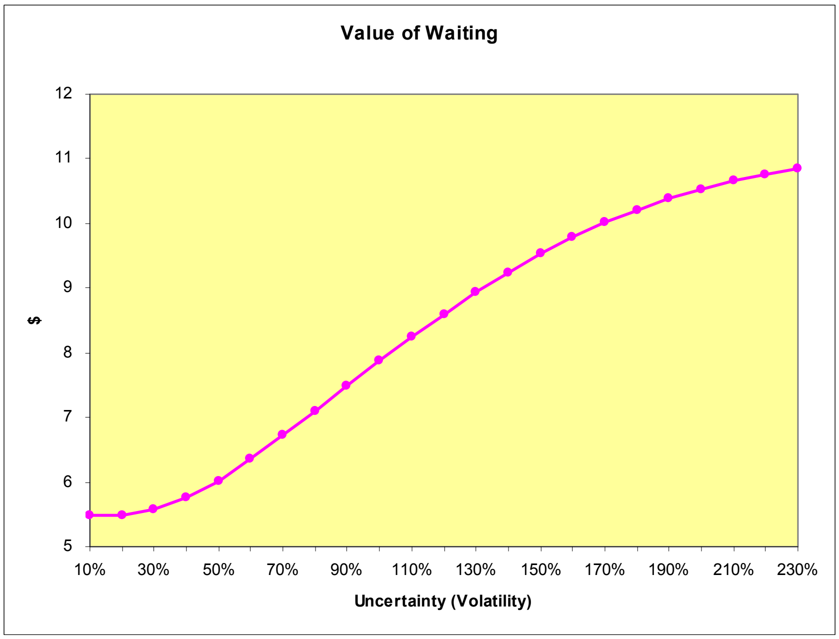

Figure 43.10: Sensitivity of the value of the YAGNI scenario to uncertainty

A Deeper Look at YAGNI

The YAGNI example discussed in the previous section assumes that the cost of change is constant over time. More insight can be gained through a closer look at the value of the YAGNI delay option under other cost functions. Consider the following two cost curves:

- A *traditional* cost curve, where the cost of change exponentially increases over time
- A *flattened* cost curve, where the cost of change gradually increases over time at a diminishing rate

An example of each type of cost curve is plotted in Figure 43.11. To see how the shape of the cost curves and waiting time affect the value created, we reevaluate the YAGNI scenario, using these sample curves.

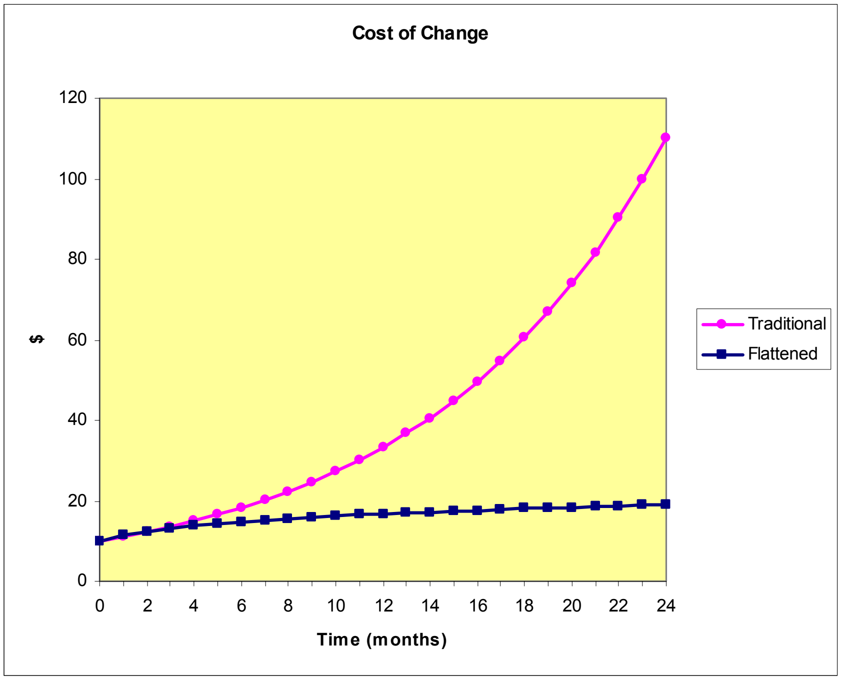

Figure 43.11: Sample cost curves: traditional versus flattened cost of change

Assume that the volatility of the feature's benefit is constant at 100% per year. Because this is per-period volatility, as waiting time (or the expiration date of the option) increases, cumulative volatility—total uncertainty around the benefit—also increases. The longer one waits, the more likely it is for the actual benefit to wander up and down and deviate from its expected present value of $15.

Figure 43.12 shows the result of reevaluating of the YAGNI option under the two cost curves. The option value, the value of waiting before implementing the feature, is shown for different waiting times for each curve. The dashed line represents the benchmark NPV of $5—the value of implementing the feature now, without any delay.

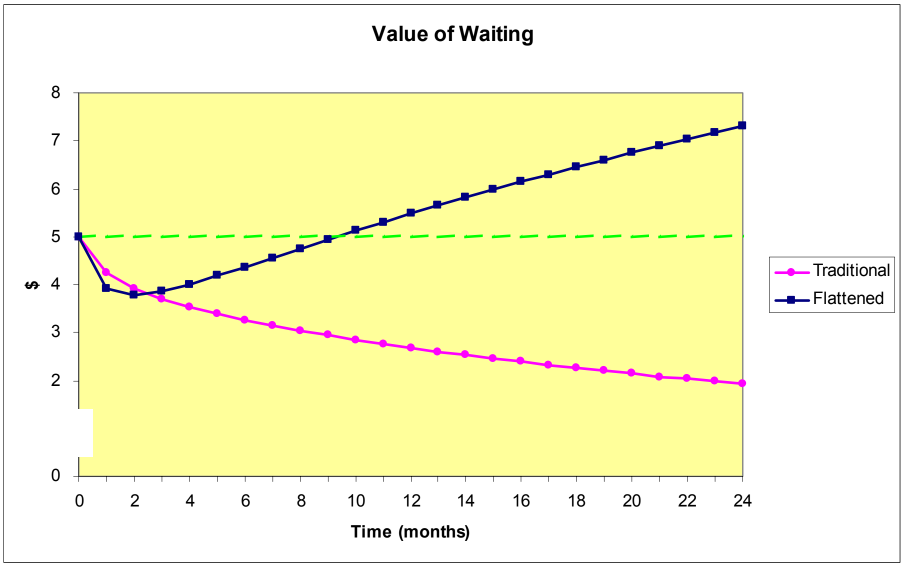

Figure 43.12: Option value of waiting under traditional and flattened cost curves

The bottom curve in Figure 43.12 reveals that under the traditional cost curve, waiting does not make much economic sense. Delaying the implementation decision destroys value because the increase in the cost of change overtakes the benefit of the flexibility to make the implementation decision later. As a result, the longer we wait, the less value we create.

Under the flattened cost curve (the top curve in Figure 43.12), however, the behavior is drastically different. If the uncertainty is expected to be resolved within a threshold waiting time, waiting is not profitable because of the initial ramp-up in the cost of change. After this initial, rapid ramp-up, the cost curve flattens, and waiting becomes increasingly profitable. The option value crosses over the $5 benchmark at approximately ten months, the threshold waiting time. Beyond this point, delaying the implementation decision creates more value than the immediate implementation of the feature.

In summary, the option pricing theory confirms that under the traditional cost model of change, decisions about system features should be committed to as soon as possible: Waiting is not desirable in this situation. However, under a flat cost curve, the timing of commitment depends on the level of uncertainty and when uncertainty about the benefits of the features is expected to be resolved. If uncertainty is high or it is expected to be resolved over the long term, decisions about system features should be committed to as late as is feasible; otherwise, they should be committed to now. Finally, under a constant cost function, commitment should always be made later rather than sooner. Figure 43.13 summarizes these conclusions.

Type of Cost Function

Level of Uncertainty / Time Horizon	Traditional	Flattened	Constant
Low / Short	Now	Now	Later
High / Long	Now	Later	Later

Figure 43.13: YAGNI scenario and the cost of change: implement now or implement later.

Why Small Investment and Frequent Releases?

Another important principle of XP is to start with a small initial investment. How can XP afford to start a project with few rather than many resources? What is the rationale behind this principle? Consider this statement from the CEO of an international consulting firm, made during a discussion of the strategy of a start-up venture in Silicon Valley.

> *I'm convinced that successful new ventures—successful new anythings—come from* thinking big, but starting small. *Most big failures come from* thinking big and starting big *and getting into trouble financially or strategically because there hasn't been enough learning to translate the big idea into a workable idea before overcommitting the amount of money or how the big idea is implemented. Iridium—the Motorola satellite-based mobile phone venture—comes to mind as an example. Note how [the president of the start-up being discussed] is gradually building up his capital base through a series of small financing rounds rather than a big-bang financing that, had he been successful in getting it, probably would have led to poor use of the money because he hadn't learned enough about how to translate his big idea into a workable one.*
>
> —K. Favaro, CEO, Marakon Associates

In XP, the rapid feedback supplied by tight iterations resolves uncertainty, whether technical or business-related, and permits the results of the learning process to be incorporated into subsequent iterations. Tight implementation cycles and frequent releases allow for decision points where the information that has been revealed can be taken advantage of to modify the course of the project. If the project is going badly, it can always be stopped. If it's going well, there is an option to continue with the next cycle. This process of continuous learning and acting based on the information revealed improves flexibility and minimizes risk. The cost of learning is limited to the small investment required to complete a small cycle, and its impact is therefore proportionately small. Taking proper action after learning increases value if the cost of learning is relatively small

In the remainder of this section, we illustrate exactly how small investments and frequent releases increase the value created.

A Black Hole: Large Investment, No Learning

First consider the complete opposite of small investments and frequent releases: a scenario involving a large initial commitment, but no learning, no decision points. Essentially, this is a single-stage project with a large investment in the beginning and a mega-release at the end.

Figure 43.14 illustrates the scenario. The only decision in the scenario is that of go/no-go in the beginning. Alas, whether the large investment will pay off is not known *a priori*. Uncertainty about the success of the project is resolved only once the release goes out the door, at the end. The probability of the project ending up worthless may be substantial because the course of the project could not be modified in the face of new information. There are no opportunities to take corrective action.

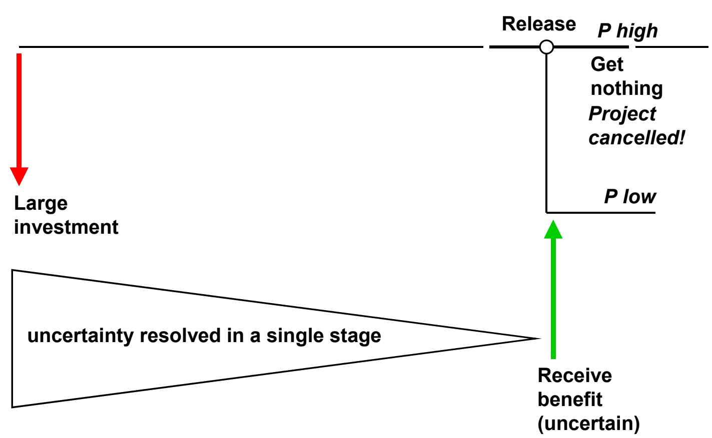

Figure 43.14: A single-stage project with no intermediate learning

Now to lighten up things a little, assume that the probability of complete failure is zero. Throwing in a few numbers will make things more concrete.

- The large investment will cost $110 in present value terms.
- The total duration of the project is four months.
- The expected benefit of the whole project, again in present value terms, is $100.

- The benefit is subject to a monthly volatility of 40%.

Where multiple sources of uncertainty are present, the volatility measure collapses the different factors involved into a single factor. Each of these factors may have both a private and a market component. In this case, let's suppose that changing customer requirements are the sole source of uncertainty as with the YAGNI scenario, which again may be affected by both external and internal developments. What does the 40% figure imply? If the product were ready now, the customer would expect an immediate benefit of $100 in present value terms. Think of the 40% volatility as the standard deviation of the monthly percentage change in this expectation based on past experiences.

NPV in the Black Hole

The NPV of the single-stage project is simply the present value (PV) of the expected benefit, net of the PV of the large investment. Because all figures are expressed in PV terms, they have already been discounted. Thus, the NPV is calculated as follows:

$$NPV = 100 - 110 = -10$$

A negative NPV! The project does not look attractive. According to the NPV rule, it should not be undertaken, because it is expected to destroy rather than create value.

Remarkably, here we did not use the volatility of the benefit in the calculation of the NPV. This is because the benefit was already specified in *expected* PV terms—that is, the risk of the benefit is factored into the $100 estimate. Alas, such is not always the case. As we will see, the volatility plays a crucial role when the project involves decision points in the middle.

Light at the End of the Tunnel: Small Investments with Learning

Having established a benchmark for comparison with the single-stage project, let's now consider the alternative scenario, which is the real focus of the current discussion. This time, the same project is undertaken in multiple stages, each stage requiring a relatively small investment and resulting in a new release. This new scenario is illustrated in Figure 43.15.

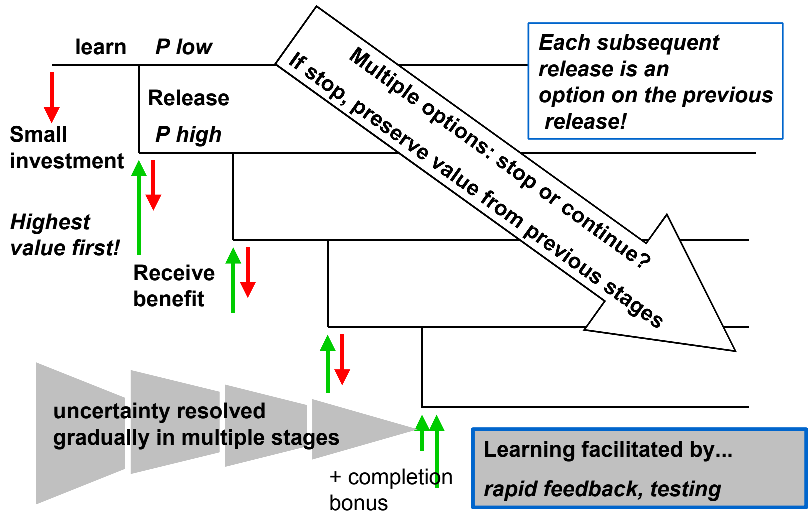

Figure 43.15: Staged project with small investments

Here are some characteristics of the new scenario. The releases progress in small increments. The stages can be ordered to implement the higher-value features first so that the present value of the total value realized is maximized (*earn early*, *spend late*). Moreover, each stage provides a learning opportunity. The customer can revise the estimates of future benefits and make an informed decision on whether to stop, continue as is, or modify the course of the project. The development team can similarly learn and steer technical choices and manage customer expectations according to the revised estimates. Uncertainty is gradually resolved in multiple steps.

Most remarkably though, each stage effectively creates a *real option* to undertake a subsequent stage. If the project is abandoned midstream, the value created during previous stages can at least be partially preserved: Only the investment associated with the last release will be completely lost. The additional value created by staging over the benchmarked single-stage scenario may be substantial. The more uncertain the expected benefits are, the higher this difference will be.

A Project with Two Stages

To see why, consider a seemingly small improvement over the simple single-stage scenario discussed in the previous subsection: a two-stage version of the same project with a single, midpoint decision.

Each stage covers half the original scope, takes half the total time, yields half the expected benefit, and incurs half the total cost of the single-stage project. Learning is incorporated into the scenario as follows. At the end of the first stage, the customer will revise the estimate of the remaining benefit, the expected benefit of the second stage, and decide whether to continue. Therefore, the second stage is conditional on the outcome of the first stage. Initially, the project benefits are subject to the same uncertainty as the benchmarked single-stage project, at a volatility of 40% per month. Though, unlike in the single-stage project example, this time the volatility will have a serious effect on value. Table 43.3 summarizes the set up of the two-stage project

Table 43.3: Setup of the Two-Stage Project

	Stage 1	**Stage 2**	**Overall**
Flexibility: Purpose: Uncertainty:	Mandatory Learning More uncertain	Optional Completion Less uncertain	
Cost	55	55	110
Benefit	50	Stage 1 outcome	?
Volatility (per month)	40%	?	?
Duration (months)	2	2	4
Risk-free rate (per month)	0.41%	0.41%	0.41%

The costs and benefits in each column of Table 43.3 are stated in PV terms relative to the beginning of the period covered by the column. The risk-free rate is assumed to be a constant 5% per year, or 0.41% per month. The overall cost of 110 is the sum of the first-stage cost and the second-stage cost, but the latter is first discounted at the risk-free rate back two months from the beginning of the second stage.

The correct way to calculate the NPV of this scenario is by viewing the second stage as an option that will be exercised only if its expected benefits (estimated at the end of the first stage) exceed its expected cost of 52.2. This contrasts with the DCF approach, which would view undertaking the second stage as a given.

To value the option underlying the two-stage project, we need a model that is richer and more accommodating than that of Black-Scholes. We will employ a closely related, but more general model, of which the Black-Scholes model is a special case. Figure 43.16 demonstrates how to calculate the *expanded* NPV of the two-stage project—the NPV including the option value—using this model and an accompanying technique called *risk-neutral valuation*. The details of the calculation are given next.

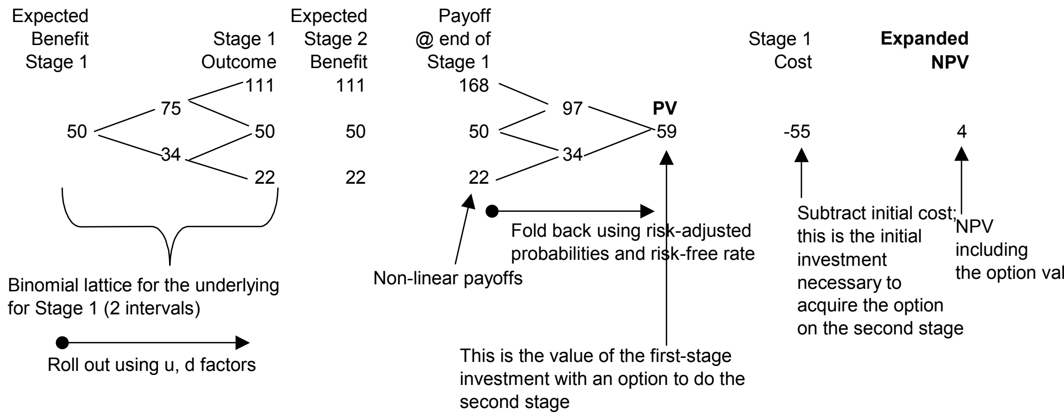

Figure 43.16: Valuation of the two-stage project

Uncertainty in a Staged Project: The Binomial Model

The first step is to determine how to model uncertainty. The *binomial model* [Sundaram1997] is frequently used in option pricing to model uncertainty for solving problems with more complex structures than standard option pricing formulas can accommodate.

In the binomial model, the underlying asset of an option is modeled using a two-state, discrete-time random walk process. Starting from an initial value, the asset moves either up or down in a fixed interval. The process is then repeated for successive intervals such that two consecutive opposite moves always take the asset to its previous value, generating a *binomial lattice*. The resulting structure represents the possible evolution of the asset in discrete time, starting with an initial value. It is essentially a binary tree with merging upward and downward branches.

In the two-stage scenario, the underlying, uncertain asset is the benefit of the first stage. The value of the overall scenario depends on the behavior of this asset. On the left side of Figure 43.16, a binomial lattice is shown for this asset. The root of the lattice is represented by the value 50, which is the specified expected present value of the first-stage benefit. Recall that the total duration of the first stage is two months. Suppose that at the end of the first month, enough information will exist to revise this estimate. Thus, we divide the duration of the project's first stage into two equal intervals, resulting in an interval size of one month.

The values of the subsequent nodes of the binomial lattice are determined using the volatility estimate of 40% per month. From the volatility, first we calculate an upward factor, u, that is greater than unity and a downward factor, d, that is smaller than unity. Over each interval, the value of the asset either increases by a factor of u or decreases by a factor of d. The upward and downward factors are chosen to be consistent with the volatility estimate, the standard deviation of the rate of percentage change in the asset's value. If the volatility is σ, u and d can be chosen as follows [Cox+1979]:

$$u = \exp(\sigma\sqrt{\tau}) \text{ and } d = 1/u$$

where τ is the chosen interval size, expressed in the same unit as σ, and "exp" denotes the exponential function. In the current example, the volatility is 40% per month and the selected interval size is one month. These choices yield the upward factor u = 1.49 and the downward factor d = 0.67. Before proceeding, we need to verify that the monthly risk-free rate of 0.41% + 1 = 1.0041 is greater than d and smaller than u, a condition that must be satisfied so that we can apply the principles of replicating portfolio and law of one price to the scenario.

Treating Nonstandard Payoffs

As shown in Figure 43.16, the PV of the stage 1 benefit is 50, which constitutes the root node of the binomial lattice. The lattice is rolled out beginning with this initial value and multiplying it repeatedly with the upward and downward factors to cover two intervals, which takes us to the end of the first stage. This process yields three terminal nodes—111, 50, and 22—each representing a possible stage 1 outcome. For each of these states, the expected stage 2 benefit equals the stage 1 outcome, as was stipulated in Table 43.3. This yields the estimate of the stage 2 benefit, conditional on the actual benefit of stage 1. Stage 2 will be undertaken only if its estimated benefit at the end of stage 1 exceeds its estimated cost of 55. Thus, applying the rational exercise assumption at the end of the first stage yields the following for each terminal node of the binomial lattice:

(Net Value of Stage 2) =
max{0, (Conditional Benefit of Stage 2) – (Cost of Stage 2)}

The overall net value, or payoff, at the end of the first stage therefore equals the folowing:

(Outcome of Stage 1) + (Net Value of Stage 2)

From top to bottom, the payoffs are calculated as 168, 50, and 22 for the three terminal nodes. Note that stage 2 will be undertaken only for the top node, the one with a payoff of 168. For the remaining nodes, the payoff simply equals the stage 1 outcome because the subsequent option on stage 2 is not exercised in those states.

Now comes the tricky part: recursively folding back the lattice to obtain the present value of the calculated payoffs. We perform this by invoking the same two concepts that underlie the Black-Scholes option pricing model: replicating portfolio and law of one price. Note that the Black-Scholes formula couldn't be used directly here, because the payoff function is not exactly the same as that of a standard call option: It does *not* simply equal the greater of zero or the maturity value of the asset net of an exercise price. We develop the general technique on the fly using the current example.

Calculating the Present Value of the Payoffs

Consider the top two terminal nodes of the binomial lattice in Figure 43.16 with the corresponding benefits of 111 and 50 and payoffs of 168 and 50. The terminal benefits of 111 and 50 are derived from the benefit at the parent node using the upward and downward factors $111 = 75u$ and $50 = 75d$. What is the expected discounted payoff at the beginning of the preceding interval? We can always attach probabilities to the upward and downward branches, calculate the expected payoff using these probabilities, and then discount the result back one interval using a proper discount rate. This procedure would have worked, except that (a) we don't know what those probabilities are, and (b) we don't know what the proper discount rate is. Besides, even if the probabilities were given, we would have to figure different discount rates for different branches, because the risk of the project changes after the option has been exercised. For large lattices, this procedure is simply impractical.

Instead, we appeal to the concept of replicating portfolio. According to this concept, the payoffs of 168 and 50 at the terminal states can also be realized artificially by forming a portfolio composed of a *twin security* and a fixed-interest loan. Assume now that there exists such a security—one whose movement parallels that of the benefit. The absolute value of the twin security is not important, but it must be subject to the same upward and downward factors. When the benefit moves up or down, the twin security also moves up or down by the same factor. Assume the value of the twin security at the beginning of an interval is M.

The replicating portfolio is formed at the beginning of the interval this way.

- Buy n units of the twin security. This represents the position of the replicating portfolio in the underlying asset.

- Take out a loan in the amount of *B* at the risk-free rate of interest to partly finance this purchase. This represents the position of the replicating portfolio in the risk-free asset.

The worth of the replicating portfolio at the beginning of the interval then equals $nM - B$. If we can determine the value of *n* and *B*, we can calculate the exact value of the replicating portfolio (as we will see, we don't need to know the value of *M*). This is the right point to apply the law of one price: The value of the replicating portfolio must equal the expected value of the terminal payoff at the beginning of the interval, the price one would have to pay at that time to acquire the option to continue with the second stage at the end of the interval.

Now let's consider the possible values of the portfolio at the end of the interval. After one interval, the loan must be paid back with interest to receive the payoff. Regardless of what happens to the price of the twin security, the amount of the loan will be $B(1 + r_f)$, including the principle and the interest accrued. Here r_f is the risk-free rate, the total interest rate on the loan over one interval.

On the one hand, if the price of the twin security moves up to uM, the portfolio will then be worth $uMn - B(1 + r_f)$. For the portfolio to replicate the payoff, this amount should equal 168, the payoff after the upward movement. On the other hand, if the price of the twin security falls to dM, the portfolio will be worth $dMn - B(1 + r_f)$, which must equal 50, the payoff after the downward movement. Thus the law of one price provides us with two equations.

If the price moves up:

$$168 = \text{(Terminal payoff)} = \text{(Terminal value of replicating portfolio)} = uMn - B(1 + r_f)$$

If the price moves down:

$$50 = \text{(Terminal payoff)} = \text{(Terminal value of replicating portfolio)} = dMn - B(1 + r_f)$$

Because r_f, u, and d are all known, we can solve these two equations for *B* and *n* as a function of *M*, and then calculate the portfolio value at the beginning of the interval by plugging the solution into the expression $nM - B$. Fortunately, the unknown *M* is eliminated during this process, yielding a value of 97. This amount is precisely how much the option to continue with the second stage would be worth at the node labeled 75 in the binomial lattice. We can repeat the same procedure for the middle and bottom terminal nodes to obtain a value of 34, and then once again with the two computed values 97 and 34, regarding them as new payoffs, to reach the root of the lattice. In the end, we obtain a final root value of 59. This amount is precisely how much the option to continue with the second stage would be worth at the beginning of the project.

A Simple Procedure: Risk-Neutral Valuation

The procedure described in the previous subsection may seem somewhat cumbersome. Fortunately, there is an easier way. Solving a system of simultaneous equations to obtain the portfolio value at the beginning of an interval is equivalent to computing the expected value of the payoffs at the end of the interval using an artificial probability measure, and then discounting back this expected value at the risk-free rate by one interval. Figure 43.17 illustrates this simple technique.

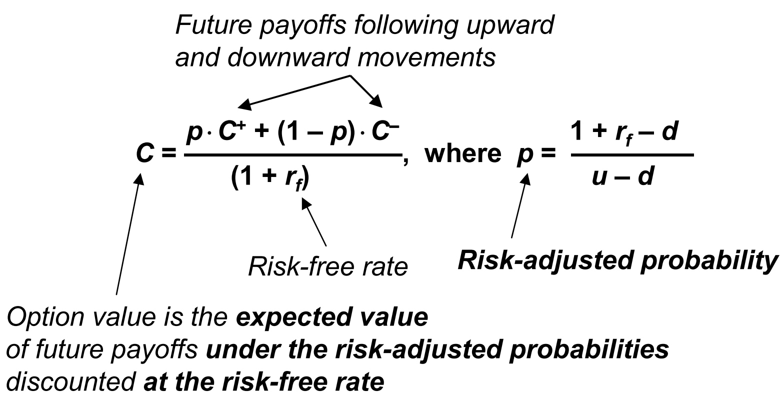

Figure 43.17: Risk-neutral valuation in the binomial model

In the middle portion of Figure 43.16, the portfolio values at the intermediary nodes and at the root of the binomial lattice are computed using the simplified procedure as follows. Starting with the terminal payoffs and recursively moving back in time:

$$97 = \frac{168 \cdot p + 50 \cdot (1-p)}{1+r_f} = \frac{168(0.41)+50(0.59)}{1.0041}$$

$$34 = \frac{50 \cdot p + 22 \cdot (1-p)}{1+r_f} = \frac{50(0.41)+22(0.59)}{1.0041}$$

$$59 = \frac{97 \cdot p + 34 \cdot (1-p)}{1+r_f} = \frac{97(0.41)+34(0.59)}{1.0041}$$

where:

$$p = \frac{1+r_f-d}{u-d} = 0.41 \quad \text{and} \quad 1-p = 0.59$$

The quantities p and 1 – *p* here and in Figure 43.17 are referred to as *risk-adjusted*, or *risk-neutral*, probabilities. They are not the actual probabilities of the upward and downward movements of the underlying asset, yet they are used to compute an expected value (in Figure 43.17, the numerator in the equation on the left). The expected value is simply discounted back at the risk-free rate r_f. The artificial probabilities *p* and 1 – *p* depend on the spread between *u* and *d*, the upward and downward movement factors of the twin security. In a way then, *p* and 1 – *p* capture the variation—or the total risk—of the underlying asset relative to the risk-free asset.

The general, recursive process of computing the present value of an asset based on replication and law of one price (no arbitrage) principles is referred to as risk-neutral valuation.

A number of features are remarkable about this technique. First, the value calculated does not require the actual probability distribution of the underlying price movement. Second, it does not require a discount rate, given the initial value of the underlying asset. Third, the procedure is independent of how the future payoffs are calculated. Because the rules used to calculate the payoffs don't matter, the process is the same for any payoff function.

Two-Stage Project: NPV with Option Value

The root value of 67 obtained in the previous subsection represents the PV of stage 1 and stage 2 combined, viewing stage 2 as an option on stage 1. This amount, however, does not account for the initial cost, the cost of stage 1, or the investment necessary to create the option on stage 2 in the first place. If we subtract this cost of 55 (which is already given in PV terms) from the calculated value of 59, we obtain an expanded NPV of 4, as shown on the right side of Figure 43.16. This value is an *expanded* NPV in the sense that it subsumes the value of the staging option.

Remarkably, the new NPV is not only positive, but also significantly higher than the benchmark NPV of the single-stage project, which was calculated to be –10. The difference of 22 is sizable compared with the total expected benefit of the single-stage project. Although they incur the same cost in PV terms, the two-stage project with learning creates a lot more value at the given level of volatility.

Impact of Uncertainty on the Option Value of Staging

The uncertainty of the expected benefit has a great impact on how much value is created when learning and additional flexibility are incorporated into the scenario. In the previous subsection, we calculated the expanded NPV using a volatility of 40% per month. This volatility captures the uncertainty of the benefit in terms of the variation in percentage changes in the estimate of the benefit from the start to the end of the first stage. What happens when this volatility increases or decreases?

Figure 43.18 plots the expanded NPV of the two-stage project as a function of volatility. As the volatility increases, the project value increases as well: The more uncertainty there is, the more important it is to have flexibility in the project. Remarkably, this effect was not observed in the single-stage project: As long as the present value of the benefit does not change, NPV remains constant. Although uncertainty also exists in the single-stage project, it was not accompanied by a discretionary, midproject action that depended on the uncertainty. Consequently, volatility has no further impact on the project value as long as it has already been accounted for in the PV estimate of the benefit.

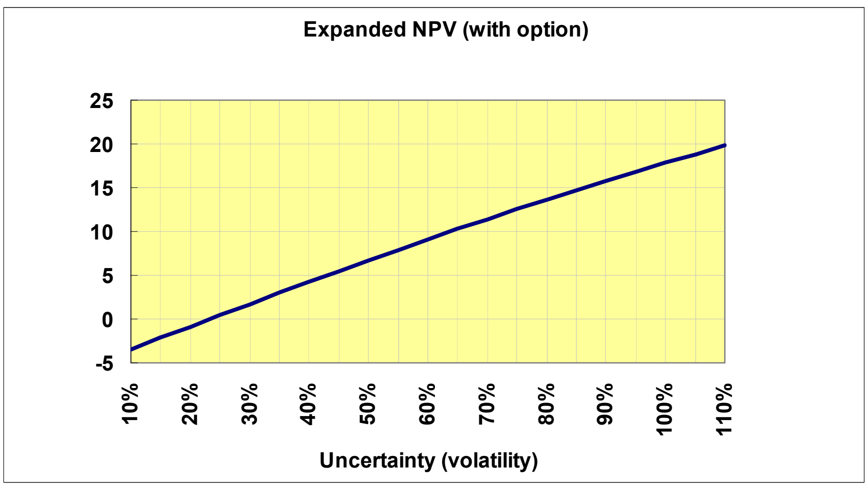

Figure 43.18: Effect of uncertainty on the value of the two-stage project

Implications of Real Options Thinking

The flexible principles and practices of XP create several kinds of real options. These options allow us to view XP as a process that maximizes value. The insight gained can be expanded to other contexts as well, from risk and contract management to compensation and incentive systems.

A Different Attitude toward Risk Management

> *More careful, risk-averse companies get left in the dust.*
>
> —Dan Scheinemann, Cisco

We've seen that although project uncertainty may endanger the value of a project, uncertainty in the environment *increases* the value of real options embedded in a project. Risk-averse companies get left in the dust because they are not able to take advantage of external opportunities that accompany risk.

The most important implication of this argument is the need for an *expanded* view of risk management. The real options approach refutes the notion that all risk is bad, that all risk has to be reduced, and especially, that all risk reduces value. This observation shifts the emphasis from the project-level idea of *contingency plans*, which have the connotation of something going wrong, to the idea of *contingent investments*, where it is economically justifiable to move toward risk, knowing that value is best maximized through active management [Favaro2002]. In other words, in an environment of high uncertainty, the emphasis shifts to *managing* risk from *reducing* risk. In fact, options theory goes one step further: It maintains that it is *total risk*, not only project-level or market-level risk, that needs to be managed.

Contractual Innovation

One of the most recent uses of options is in the design of innovative contracts with nonlinear payoffs. For example, a contract with a ceiling price or a floor price can be synthesized through a combination of buying and selling call and put options. In general, any set of contingent payoffs—payoffs that depend on the value of an underlying asset—can be priced as a combination of options on that asset. Thus derivative concepts can be used to engineer contracts using a mix of financial and real options, combined appropriately to manage risk. In Silicon Valley, these concepts are being used regularly to work out financing for venture capital start-ups and for licensing agreements. For an example of the use of contractual options in license agreements, see [Erdogmus2001A]

The options perspective can also be applied to XP contracts. In a white paper, Beck and Cleal [Beck+1999] note that XP contracts have characteristics of options, in the sense that the features to be implemented are optional—they don't necessarily have to be implemented. That is, the scope of contracts is not predetermined.

XP contracts are not *fire-and-forget*. The customer and the development team have a set of decision points, which give them the ability to actively manage the contract. The customer exercises its options by asking the team to implement a set of features. As the team completes new features, the customer has to decide which new options to exercise next. These may be options that were in the original scope, options that were under consideration originally but not in the original scope, or newly discovered options. Thus, as successive iterations resolve uncertainty, the customer has decision points in which to intervene and maximize new opportunities while minimizing downside effects.

One surprising insight that can be gained from this perspective is that XP contracts add the most value when feature benefits are least certain. When a feature is *deeply in-the-money* (that is, the value of its immediate implementation is high), the feature should be implemented. Conversely, when a feature is *deeply out-of-the-money* (that is, the value of its immediate implementation is highly negative), it should not be implemented. However, when the feature is *at-the-money*, the NPV of immediate implementation is around zero, and the customer is uncertain of its benefits. In this latter case, the option to delay implementation may add a great deal of value, because the future is likely to resolve that uncertainty. This is precisely where the YAGNI principle makes the most economic sense.

Similarly, additional value can be created when the customer has a choice to implement the best of a set of alternative features. Such a choice is represented by a *best-of,* or *rainbow,* option. In an option on the *best of two assets*, the investor holds two options but is only allowed to exercise one of them, effectively creating a hedge. Options theory can demonstrate that best-of options have the most value when the underlying assets are negatively correlated—that is, if one asset goes up, the other goes down. Intuitively, this makes sense because it gives the holder of the option a real choice rather than a hypothetical one. In this light, XP contracts can also be seen as a portfolio of best-of options, where the customer is offered a set of suggested features plus the alternatives. A choice between contrasting alternatives is more valuable to the customer than one between related alternatives.

The Role of Discipline

The nonlinear thinking that underlies real options also has profound implications for information technology management in an increasingly uncertain environment.

An organization embracing this thinking systematically identifies leverage points where flexibility would be desirable, analyzes these leverage points, and structures projects with options that take advantage of them. Projects are continually refined to embed in them further options that increase their value.

An organization embracing the real options approach becomes less averse to total risk. It recognizes that opportunity and risk go hand in hand. Consequently, it encourages and supports learning investments that explore new opportunities.

Paradoxically, an option derives most of its value from rational exercise. Therefore, if the organization is to move responsibly toward valid contingent investments with many embedded options, more rather than fewer projects will be started (although many of these will not be taken to completion due to rational exercise). To create and maximize value, decision makers need the discipline to abandon projects when their option value no longer justifies further investment. This attitude has serious repercussions in terms of how compensation and incentive systems should be redesigned. It will be important for management not to penalize teams for killing projects that aren't working out, moving from a philosophy of *killing careers* to a philosophy of *killing projects*. Project abandonment needs to have a positive connotation if done for the right reasons, not the negative connotation that it has in a myopic perspective of risk management.

An interesting anecdote for this line of thinking comes from the entertainment industry. A screenwriter is highly rewarded for the successful completion of a screenplay, even if a film is never made from it. The reason is the recognition of movies as high-risk, high-payoff projects. Simply the fact of creating the option to make a film (by developing a finished, professional-quality screenplay) is correctly recognized as having high economic value. The possibility that the option to make the film may not be exercised (or that the project may be abandoned in midcourse) is fully accounted for. By the same token, venture capital firms invest in a portfolio of highly risky projects. Most fail, but the few that succeed justify the investment in the portfolio as a whole.

It may seem that such an approach is destined to remain up at the relatively abstract levels of strategic planning and never be seen down in the trenches, where development is carried out. This is not at all the case. Recall that XP explicitly encourages technical experiments—miniprojects within projects, or *spikes* in XP terminology, –- that quickly test new ideas. These experiments ultimately are either incorporated into the main project if successful or abandoned otherwise. Developers are rewarded for this kind of creative exploration even if only a few of those experiments succeed.

In an options-oriented management system, the incentive and compensation package should be aligned with the concept of rational exercise—one of the cornerstones of option pricing. Such alignment encourages teams and individuals to undertake the kind of experimentation that will create valuable options for the project's or the firm's future—and not fear for their jobs if some of those options are not realized. One of the four values of XP, courage, is fundamental here: To impact creation of economic value, courage should imply not only the courage to create and exercise options, but also the courage to practice rational exercise in a disciplined way.

Summary

Software development takes place in an inherently uncertain environment, one that is constantly changing. In such an environment, continuous formulation and implementation of options enhances the capability to deal with the vast uncertainty through increased flexibility, exploitation of opportunities, and avoidance of pitfalls. By its very design, XP is well positioned to take advantage of this mind-set. The language of real options and the financial theory behind it demonstrate how the principles and practices of XP can guide a project toward value creation.

Final Remarks

Academics and practitioners frequently debate the suitability of the financial options analogy, and the techniques developed to price these options, to the valuation of real options scenarios. The main differences between financial and real options are summarized in Table 43.4.

Table 43.4: Financial Options Versus Real Options

Financial Options	**Real Options**
Complete markets. Payoff structure can be emulated by a replicating portfolio.	*Incomplete markets.* Payoff structure often cannot be practically emulated by a replicating portfolio.
Traded asset. The underlying asset is traded in the financial markets.	*Twin security.* The underlying asset is not traded; instead, a proxy, or twin security whose value is correlated with the underlying asset, must be assumed. This also applies in DCF and NPV.
Observed current price. The current price of the underlying asset is observed.	*Lack of an observed current price.* The current price of the underlying asset is not observed. It must be estimated as a present value from a stream of future cash flows.
No discount rate. A discount rate is not needed to value the option because of the existence of an observed price and the use of replication and no arbitrage assumptions.	*Discount rate needed.* A discount rate is often needed to calculate the PV of a stream of future cash flows as a substitute for the current price of the underlying asset.
No interaction. Financial options are self-contained contracts. They don't interact.	*Extensive interaction.* There are often complex interactions among different real options within a project or even across different projects. The behavior of one option affects the value of the other.
Sources of uncertainty constrained. Financial options involve one or two uncertain underlying assets.	*Multiple sources of uncertainty.* Real options often involve multiple underlying assets or assets with multiple sources of uncertainty.
Single ownership. Financial options have defined ownership.	*Shared ownership.* Real options are often shared among competitors. A company's exercise of a real option may kill or significantly undermine the same real option for a competitor and vice versa.
Value leakage. The holder of a financial option may be subject to the loss of benefits while waiting to exercise the option because of dividend payments or convenience yield that are available to the holders of the underlying asset, but not to the holders of an option. The rate and pattern of this can be estimated using historical data or using industry conventions.	*Competition, partnerships, and sharing.* The holder of a real option may be subject to the loss or amplification of benefits while waiting to exercise the option because of the actions of competitors and partners and shared ownership, all of which may be very difficult to quantify.

Two points are important to keep in mind here.

First, the existence of trading markets and assumptions about their efficiency, completeness, and liquidity are critical for financial option pricing techniques. These techniques are designed to treat risks that can be priced in such markets. The suitability of a specific option pricing technique to value a real options scenario may thus depend on the nature of the uncertainty being treated and *how close to the market* the underlying assets are. In software development projects, where private risk is an important factor, the financial options analogy may be weak.

In cases where the financial option–real option analogy is weak, the option values yielded should thus be treated as *idealized values* computed under the assumption of the existence of a market equivalent that reasonably closely tracks the risk being tackled. Fortunately, sometimes risks that seem purely private at first can be market-priced, thanks to an expanding and vibrant securities market in the technology sector. For example, it was difficult to model the market risk associated with the growth option that Netscape had when it introduced the first browser, simply because it was the first of its kind and there weren't yet any other Internet companies on the market. This situation has changed. Netscape was rapidly joined by other Internet companies, with the result that there are now several Internet stock indexes that track the market risk and volatility associated with Internet investments. Another example is provided in [Erdogmus2001B]. For further discussion of the analogy between financial and real options and the limitations of this analogy, see the last sidebar in [Amram+1999].

Second, like all forecasts, the numbers used in options calculations will be more or less precise. The final numbers obtained are as good as the estimates used in their calculation. Where these estimates are unreliable, option values can still provide much insight if they are used in an informed manner, especially along with comparative, sensitivity, and scenario analyses.

Further Reading

Introductory corporate finance texts provide more comprehensive discussions of basic valuation concepts—in particular, capital budgeting, discounted cash flow techniques, net present value, and the relationship between risk and return. Recommended texts are Ross et al. [Ross+1996] and Brealey and Myers [Brealey+1987].

Hull's book [Hull1997] provides an undergraduate-level overview of derivative securities (including options), the general techniques for their pricing, and derivative markets. Pindyck and Dixit [Pindyck+1992] offer a deeper and more theoretical exposition of option pricing theory together with the econometric foundations of the Black-Scholes and related models.

The seminal paper on option pricing is by Black and Scholes [Black+1973], which explains the original derivation of their and Merton's Nobel Prize–winning model. Cox, Ross, and Rubinstein [Cox+1979] provide a much simpler derivation of the same model using the binomial model and the risk-neutral approach. The Black-Scholes model has many variations that have similar analytic solutions. The most relevant ones are discussed by Margrabe [Margrabe1978] and Carr [Carr1988]. Margrage derives a formula for the option to exchange two risky assets, of which Black-Scholes is a special case. This formula can be used when the exercise cost of an option is also uncertain [Erdogmus2001A]. Carr provides a comprehensive discussion of Margrabe's formula and other, more complex variations, including compound options. Kumar [Kumar1996] provides a compact discussion of the impact of volatility on option value.

Sundaram [Sundaram1997] gives the best exposition of the binomial model and risk-neutral valuation. Smith and Nau [Smith+1995] explain the relationship between option pricing and decision trees, and demonstrate how the two models together can account for both market and private risk. These two articles are highly recommended for those interested in the practical application of option pricing to real assets.

Many excellent, high-level articles exist that discuss the use of option pricing theory in valuing options on real assets [Amram+1999; Luehrman1998; Myers1984]. Myers, who originally coined the term real options, explains how the real options approach links strategy and finance. Texts that focus on applications of real options include Copeland and Antikarov [Copeland+2001] and Amram and Kulatilaka [Amram+1999]. Further applications can be found in two books by Trigeorgis [Trigeorgis1999; Trigeorgis1994]. The older of these is a self-contained textbook, and the more recent is an edited collection focusing mostly on applications.

Applications of real options to information technology in general and software development in particular have been addressed in many articles. Those addressing investment decisions in software development include [Erdogmus2001A; Erdogmus2001B; Favaro+1999; Boehm+2000; Erdogmus+1999; Favaro1996]. Sullivan et al. focus on applications to software design [Sullivan+1999], and the book by Clark and Baldwin focuses on applications to modularity in general in the context of hardware design [Baldwin+1999]. General applications to information technology investments are also available [Benaroch+1999; Benaroch+2000; DosSantos1991; Taudes+2000].

References

[Amram+1999] M. Amram, N. Kulatilaka. “Disciplined Decisions: Aligning Strategy with the Financial Markets.” *Harvard Business Review*, 1999.

[Baldwin+1999] C. Baldwin, K. Clark. *Design Rules: The Power of Modularity*. MIT Press, 1999.

[Beck1999] K. Beck. *Extreme Programming Explained: Embrace Change.* Addison-Wesley, 1999.

[Beck+1999] K. Beck, D. Cleal. Optional Scope Contracts. White paper. 1999.

[Benaroch+1999] M. Benaroch, R. Kauffman. “A Case for Using Option Pricing Analysis to Evaluate Information Technology Project Investments.” *Information Systems Research*, Volume 10, 1999.

[Benaroch+2000] M. Benaroch, R. Kauffman. “Justifying Electronic Banking Network Expansion Using Real Options Analysis.” *MIS Quarterly*, Volume 24, 2000.

[Black+1973] F. Black, M. Scholes. “The Pricing of Options and Corporate Liabilities.” *Journal of Political Economy*, Volume 81, 1973.

[Boehm+2000] B. Boehm, K. Sullivan. “Software Engineering Economics: A Roadmap.” In *The Future of Software Engineering.* A. Finkelstein, ed. ACM Press, 2000.

[Brealey+1987] R. Brealey, S. Myers. *Principles of Corporate Finance*. McGraw-Hill, 1987.

[Carr1988] P. Carr. “The Valuation of Sequential Exchange Opportunities.” *Journal of Finance*, Volume 18, 1988.

[Copeland+2001] T. Copeland, V. Antikarov. *Real Options: A Practitioner’s Guide*. Texere, 2001.

[Cox+1979] J. Cox, S. Ross, M. Rubinstein. “Option Pricing: A Simplified Approach.” *Journal of Financial Economics*, Volume 7, 1979.

[DosSantos1991] B. Dos Santos. “Justifying Investments in New Information Technologies.” *Journal of Management Information Systems*, Volume 7, 1991.

[Erdogmus2001A] H. Erdogmus. Management of License Cost Uncertainty in Software Development: A Real Options Approach. Presented at the Fifth

Annual Conference on Real Options: Theory Meets Practice. UCLA, Los Angeles, California. 2001.

[Erdogmus2001B] H. Erdogmus. "Value of Commercial Software Development under Technology Risk." *The Financier*, Volume 7, 2001.

[Erdogmus+1999] H. Erdogmus, J. Vandergraaf. Quantitative Approaches for Assessing the Value of COTS-centric Development. Presented at the Sixth International Software Metrics Symposium. Boca Raton, Florida. 1999.

[Favaro1996] J. Favaro. A Comparison of Approaches to Reuse Investment Analysis. Presented at the Fourth International Conference on Software Reuse. 1996.

[Favaro2002] J. Favaro. "Managing Requirements for Value." *IEEE Software*, 2002.

[Favaro+1998A] J. Favaro, S. Pfleeger. "Making Software Development Investment Decisions." *Software Engineering Notes*, Volume 25, 1998.

[Favaro+1998B] J. Favaro, K. Favaro, P. Favaro. "Value Based Software Reuse Investment." *Annals of Software Engineering*, Volume 5, 1998.

[Favaro+1999] J. Favaro, K. Favaro. "Strategic Analysis of Application Framework Investments." In *Building Application Frameworks: Object-Oriented Foundations of Framework Design*. M. Fayad, R. Johnson, eds. John Wiley & Sons, 1999.

[Hull1997] J. Hull. *Options, Futures, and Derivatives*. Prentice-Hall, 1997.

[Kumar1996] L. Kumar. "A Note on Project Risk and Option Values of Investments in Information Technologies." *Journal of Management Information Systems*, Volume 13, 1996.

[Luehrman1998] T. Luehrman. "Investment Opportunities as a Portfolio of Real Options." *Harvard Business Review*, 1998.

[Margrabe1978] W. Margrabe. "The Value of an Option to Exchange One Asset for Another." *Journal of Finance*, Volume 33, 1978.

[Myers1984] S. Myers. "Finance Theory and Financial Strategy." *Interfaces*, Volume 14, 1984.

[Pindyck+1992] R. Pindyck, A. Dixit. *Investment under Uncertainty.* Princeton University Press, 1992.

[Ross+1996] S. Ross, R. Westerfield, B. Jordan, G. Roberts. *Fundamentals of Corporate Finance.* Times Mirror Professional Publishing, 1996.

[Smith+1995] J. Smith, R. Nau. "Valuing Risky Projects: Option Pricing Theory and Decision Analysis." *Management Science*, 1995.

[Sullivan1996] K. Sullivan. Software Design: The Options Approach. Presented at the SIGSOFT Software Architectures Workshop. San Francisco, California. 1996.

[Sullivan+1999] K. Sullivan, P. Chalasani, S. Jha, V. Sazawal. "Software Design as an Investment Activity: A Real Options Perspective." In *Real Options and Business Strategy: Applications to Decision Making*. L. Trigeorgis, ed. Risk Books, 1999.

[Sundaram1997] R. Sundaram. "Equivalent Martingale Measures and Risk-Neutral Pricing: An Expository Note." *The Journal of Derivatives*, 1997.

[Taudes1998] A. Taudes. "Software Growth Options." *Journal of Management Information Systems*, 1998.

[Taudes+2000] A. Taudes, M. Feurstein, A. Mild. "Options Analysis of Software Platform Decisions: A Case Study." *MIS Quarterly*, Volume 24, 2000.

[Trigeorgis1994] L. Trigeorgis. *Real Options: Managerial Flexibility and Strategy in Resource Allocation.* MIT Press, 1994.

[Trigeorgis1999] L. Trigeorgis. *Real Options and Business Strategy: Applications to Decision Making*. Risk Books, 1999.

About the Authors

Hakan Erdogmus is a research officer with the Institute for Information Technology, National Research Council of Canada in Ottawa. His current research focuses on software economics. He is the coauthor of several articles on the subject. He holds a Ph.D. in telecommunications from Université du Québec, Institut national de la recherche scientifique, and an M.S. in computer science from McGill University, Montreal. In 1997, Hakan was the program co-chair of CASCON, the annual conference of IBM Center for Advanced Studies. He is the coeditor of an upcoming book entitled *Advances in Software Engineering: Comprehension, Evaluation, and Evolution*, reporting on the research activities of the Consortium for Software Engineering Research.

John Favaro is an independent consultant based in Pisa, Italy. He was general chair of the Sixth International Conference on Software Reuse in June 2000. He is

European co-chair of the IEEE Technical Subcommittee on Reuse and is a founding member of the steering committee of the Society for the Advancement of Software Education. Recently, he has published articles on principles of strategy and valuation for IT investment, including "Value-Based Software Reuse Investment" in the special issue on reuse of *Annals of Software Engineering 1998*; that article became the basis for the material on options in *Extreme Programming Explained*. He has degrees in computer science from Berkeley and Yale.

Hakan and John cowrote the tutorial on the economics of XP given at the XP Universe 2001 workshop "XP for Capitalists," hosted by Kent Beck in Raleigh, North Carolina.