CORE CURRICULUM

HARVARD BUSINESS PUBLISHING

Finance

Mihir A. Desai, Series Editor

READING + INTERACTIVE ILLUSTRATIONS

NPV and Capital Budgeting

TIMOTHY A. LUEHRMAN

5176 | Published: June 10, 2016

Table of Contents

1 Introduction 3

2 Essential Reading 4

2.1 Net Present Value and Investment Decisions 4

2.1.1 Calculating NPV for a Single Risky Investment 5

2.1.2 Case Example of an NPV Calculation: Pharmaco, Inc. 7

2.2 Alternatives to NPV 11

2.2.1 Internal Rate of Return for a Single Risky Investment 12

2.2.2 Payback Period for a Single Risky Investment 16

2.3 NPV, Project Risk, and Expected Return 20

2.3.1 Comparing Expected Return to Required Expected Return 21

2.3.2 Single Versus Multiple Hurdle Rates 24

3 Supplemental Reading 28

3.1 A Set of Hypothetical Investment Projects 29

3.2 Capital Budgeting with Constraints 31

3.2.1 A Simple Fixed Budget Constraint 31

3.2.2 Problems with IRR and NPV Rankings 34

3.2.3 The Profitability Index 35

3.2.4 Limitations of the Profitability Index 37

3.3 Further Complications and Some Limitations of NPV 40

3.3.1 Interdependent Projects 41

3.3.2 Contingent Projects 41

3.3.3 Intertemporal Trade-offs 43

3.3.4 Budget and Managerial Constraints 43

4 Key Terms 45

5 Notation 47

6 Practice Questions 47

7 Index 48

This reading contains links to online interactive illustrations, denoted by the icon above. To access these exercises, you will need a broadband Internet connection. Verify that your browser meets the minimum technical requirements by visiting http://hbsp.harvard.edu/tech-specs.

Former Harvard Business School professor Timothy A. Luehrman developed this Core Reading with the assistance of writer Barbara Wall Lobosco, HBS MBA 1995.

1 INTRODUCTION

In the normal course of business, companies regularly face important decisions about how to deploy their capital to maximize value. Such decisions may involve continuing investment to support and expand existing businesses, changes in business strategy and/or significant new initiatives, or the liquidation of assets developed from prior investments. Many factors considered from multiple perspectives are combined in a corporate process of ***capital budgeting***. The process is typically staged and includes elements such as project definition, a strategic analysis, a risk assessment and scenario analysis, financial analyses, execution planning, and post-investment performance assessments. This reading adopts a fairly narrow corporate finance perspective on the capital allocation problem. Specifically, we examine the valuation tools and financial metrics companies use as they evaluate opportunities within the larger process.

The reading first considers the relatively simple problem of evaluating a single investment proposal (or multiple proposals, each in isolation). It introduces techniques and tools for evaluating a project and deciding whether to invest in it or not. These tools include net present value (NPV), internal rate of return (IRR), and other financial metrics, along with related decision criteria. We will see how NPV is related both to the goal of value maximization and, more broadly, to external capital markets, and why it has become such an important, widely used guide to investment decision making. But we also look at competing metrics, their pros and cons, and how each compares to NPV in execution and likely performance.

The reading then turns to the more complicated problem of comparing and ranking multiple projects. This is essential in many real-world settings, such as when companies face financial or managerial constraints that prevent them from investing in all apparently positive-NPV projects. We will examine different ranking schemes and assess their performance and limitations. Finally, we also consider some real-world complications that can limit the reliability of some basic financial metrics, including NPV and various types of rankings, and consider ways to adapt metrics and methodologies to corporate realities, such as budget constraints, interdependencies, and contingencies.

This reading presumes readers have a good grasp of the mathematics of discounting and compounding as presented in *Core Reading: Time Value of Money* (HBP No. 8299). The reading also makes use of some basic relationships between project risk and project discount rates but does not derive them formally. Readers wishing to understand how to prepare quantitative estimates of discount rates may wish to study *Core Reading: Risk and Return* (HBP Nos. 5220 and 8603)[a] and *Core Reading: The Cost of Capital* (HBP No. 8293) before starting this reading.

2 ESSENTIAL READING

2.1 Net Present Value and Investment Decisions

Net present value **(*NPV*)** is, by design, a measure of the value created by a business investment. It is widely used by managers as part of a process for allocating capital among investment opportunities. The important idea behind NPV can be expressed very simply and intuitively: NPV is the difference between how much an investment is worth today and how much it costs. If a manager identifies an asset that is worth $1,000 but which can be acquired for $900, the difference of $100 is the NPV of investing in that asset. It also is a measure of how much the owner's wealth increases when the asset is acquired for $900. In this narrow financial sense, NPV reflects *value creation* (or value destruction)—the change in wealth associated with a specific investment decision.

Now, the mere fact that $900 was invested in a business asset that is actually worth $1,000 doesn't mean the work is done. Many people may have to work hard over an extended period to implement and execute an appropriate business plan. If they don't succeed—even if they are simply unlucky—the asset could turn out to be worth less than $1,000, maybe even less than $900. So what really creates the value, the original investment of $900, or the subsequent effort to achieve operating performance? Clearly, it is both—neither alone is sufficient—and the difference may seem to be only a semantic one involving the meaning of "value creation." But it is important to understand that the NPV calculation incorporates expectations about future operations. That is, our estimated asset value of $1,000 reflects the effect on today's value of anticipated future performance, including the risk associated with it.

Because it measures the value created by a prospective investment, NPV is widely used as a decision criterion in corporate capital budgeting. Known as the ***NPV rule***, this criterion says simply, *Invest in all projects for which NPV > 0*. The intuition

[a] *Core Reading: Risk and Return 1: Stock Returns and Diversification* (HBP No. 5220) and *Core Reading: Risk and Return 2: Portfolio Theory* (HBP No. 8603).

behind this rule is equally simple. In essence, it says to buy assets that are worth more than they cost. Under certain assumptions, it can be shown that the NPV rule is value maximizing; that is, an owner or a manager who wishes to maximize the value of his or her business should follow the NPV rule to decide which investments to undertake. A corollary to the rule is that firms should not invest in projects for which NPV < 0. Projects for which NPV = 0 are value-neutral cases that neither create nor destroy value.

Although the intuition behind NPV is simple, the practical mechanics of computing it and implementing the NPV rule are often complex for several reasons: real-world business opportunities are complicated and risky; big organizations have to evaluate large numbers of different types of opportunities; time and resources for doing so may be limited; and funds available to invest may be limited. We begin this section of the reading with the mechanics of calculating NPV for a single project, without yet considering the harder problem of comparing multiple projects. Then we will compare NPV to other common investment metrics. Only then do we begin to confront the problem of implementing the NPV rule in a large corporation.

2.1.1 Calculating NPV for a Single Risky Investment

NPV equals the ***present value (PV)*** of future cash flows minus the required upfront investment:

$$\text{NPV} = \text{PV} - \text{investment}$$

Generally, the first part of calculating NPV, estimating present value, is the harder one. The future cash flows of a business are, roughly, the net after-tax operating cash flows the business is expected to generate, often over a very long time, along with the net new investment required to keep the business running. For now, we refer to them simply as "cash flows."

Recall that to estimate the value today of riskless cash flows received in the future, we discount them to present value using the basic discounting relationship developed in *Core Reading: Time Value of Money* (HBP No. 8299):

$$\text{PV} = \sum_{t=1}^{T} \frac{Cf_t}{(1+r)^t}$$

where PV denotes the present value of the stream of T future cash flows; Cf_t denotes the cash flow received at future date t; and r denotes the discount rate and reflects the time value of money, that is, r equals the risk-free rate of interest.

This is the basic discounting formula used to compute the present value component of NPV, but we need to make some important adjustments. First, future business cash flows are *risky*—they are uncertain as to timing and amount. So we must regard the cash flow at a given future date as a random variable, characterized by a probability distribution that reflects the range and likelihood of possible outcomes. Therefore, when cash flows are risky, we use the ***expected cash flow*** at time t in the numerator of the expression above. The expected cash flow is the mean of the probability distribution, denoted as $E(Cf_t)$.

Second, we need to adjust the discount rate in the denominator for risk. Why? If we did not, two projects with identical expected cash flows would have identical present values. But this makes no sense if one project is riskier than the other—investors would always choose the safer investment given identical values. Riskier expected cash flows must be discounted at a higher rate, which (appropriately) reduces their value, all else being equal.

How do we specify the discount rate? It must equal investors' ***required expected return*** on the proposed investment. By this we mean the minimum expected return investors demand for bearing the risk associated with the investment. Investors require higher expected returns on riskier investments, all else being equal, and the discount rate must reflect this. So we use a ***risk-adjusted discount rate***, denoted as k, which contains a ***risk premium***. That is, k equals the risk-free interest rate plus a risk premium: $k = r$ + risk premium. The risk premium is larger for riskier investments and smaller for less risky investments. Of course, saying that investors "require" a certain minimum expected return does not guarantee they will receive it—the investment is risky, after all, and the realized return on it may turn out to be higher or lower than the required expected return.

It makes sense that riskier projects require higher discount rates, but we don't yet know how much higher. For now, it is sufficient to observe that for a value-maximizing investor, the discount rate k must equal the ***opportunity cost of funds***. By this we mean that k equals the expected return that an investor could obtain by investing in an alternative investment with the same time horizon and the same riskiness as the investment under consideration. The intuition for this requirement is clear: No one would agree to invest in a given project if he or she could obtain a higher expected return on a similar project without taking any extra risk. It simply cannot be value maximizing to do so.

The expression below now shows both adjustments needed in the basic present value calculation to accommodate risk.

$$\text{PV} = \sum_{t=1}^{T} \frac{\text{E}(Cf_t)}{(1+k)^t}$$

PV denotes the present value of the stream of T future cash flows, $\text{E}(Cf_t)$ denotes the expected cash flow at future date t, and k denotes our risk-adjusted discount rate equal to the opportunity cost of funds.

These adjustments have only a minor effect on the appearance of the PV expression—it still "looks" the same. However, they require much more care to implement than their simple appearance suggests. In particular, estimating the appropriate risk premium to include in k requires an application of portfolio theory, covered in *Core Reading: Risk and Return* (HBP Nos. 5220 and 8603).

2.1.2 Case Example of an NPV Calculation: Pharmaco, Inc.

Pharmaco, Inc., is a company that manufactures and distributes pharmaceuticals. It is considering buying the rights to make and sell a new drug created by another company, the drug development laboratory R&D Ltd., which develops new drugs but generally lets others manufacture and distribute them. R&D Ltd. will sell the rights to the product to Pharmaco for $1 billion. The net after-tax cash flow generated by the drug is expected to be $150 million in Year 1 and to grow at 20% per year for the following four years, at which point (in Year 5) the remaining value (or ***terminal value***) of the drug is expected to be $500 million, as shown in **Exhibit 1**. Suppose the risk-free rate of return is 4% and Pharmaco has determined that 8% is the appropriate opportunity cost of funds for this project. Pharmaco needs to decide whether to go ahead and purchase the rights to the new drug.

EXHIBIT 1

Projected Cash Flows for Pharmaco's Proposed Investment ($ in Millions)

Time	***t* = 0**	**1**	**2**	**3**	**4**	**5**
Net inflows		$150.0	$180.0	$216.0	$259.2	$311.0
Terminal value						$500.0
$\text{E}(Cf_t)$		$150.0	$180.0	$216.0	$259.2	$811.0
PV @ k = 8.0%		$138.9	$154.3	$171.5	$190.5	$552.0
Sum of PVs	$1,207.2					
Investment	($1,000.0)					
NPV	$207.2					

To analyze Pharmaco's investment, Exhibit 1 calculates the present value of the expected cash flows by discounting each at k = 8% for the indicated number of periods and summing the result (note that both the last operating cash flow and the terminal value occur in Year 5):

$$PV = \frac{E(Cf_1)}{1+k} + \frac{E(Cf_2)}{(1+k)^2} + \frac{E(Cf_3)}{(1+k)^3} + \frac{E(Cf_4)}{(1+k)^4} + \frac{E(Cf_5)}{(1+k)^5} + \frac{\text{Terminal Value}}{(1+k)^5}$$

$$PV = \frac{\$150}{1.08} + \frac{\$180}{(1.08)^2} + \frac{\$216}{(1.08)^3} + \frac{\$259}{(1.08)^4} + \left(\frac{\$311}{(1.08)^5} + \frac{\$500}{(1.08)^5}\right)$$

$$PV = \$1{,}207 \text{ million}$$

Subtracting the initial investment of $1,000 million from the PV of $1,207 million gives the NPV:

$$NPV = PV - \text{investment} = \$1{,}207 \text{ million} - \$1{,}000 \text{ million} = \$207 \text{ million}$$

Should Pharmaco buy the new drug from R&D Ltd.? Yes, because the opportunity has a positive NPV. Even though Pharmaco has an immediate cash outflow of $1 billion to acquire the new drug, the company's value should rise by more than $200 million when it enters this transaction: The new drug is worth more than Pharmaco is paying for it. Once again, this is not to say that the work is done—Pharmaco's employees have a lot to do to make the operating projections come true. But the calculation anticipates their efforts and includes them in the present value of the product opportunity.

Now suppose Pharmaco already has the necessary $1 billion in cash on hand, and further, that the company's board of directors has decided that if it fails to conclude a deal for the new product, it will make no other acquisition but rather simply invest the $1 billion in US Treasury bonds yielding the risk-free rate of 4%. Does this consideration affect the opportunity's NPV? No. To see why not, we will first consider the (incorrect) logic that might lead one to recompute NPV using a discount rate of 4% instead of 8%. (At 4% the PV of the opportunity would be $1,391 million and its NPV $391 million, almost twice as large as the figure computed above.) Recall that the discount rate k is an opportunity cost of funds reflecting what the company could earn on an alternative investment. In this case, the board has declared that the actual alternative under consideration is Treasury bonds: Pharmaco will either acquire the new product or a $1 billion portfolio of T-bonds. As long as the new product returns more than 4%, it will create more value than T-bonds and is a good deal. In fact, pushing this logic still further, Pharmaco should be willing to pay *even* *more* than $1 billion for the new drug—at any price up to $1,391 million, the NPV will still be positive compared to Treasury bonds. Right?

Wrong. Treasury bonds cannot be used as a benchmark for k because they are risk free and the new drug is not. The opportunity cost of funds is the expected return on an alternative investment of the *same risk*. The new drug is riskier than Treasury bonds. If Pharmaco applies a discount rate of 4%, it will overvalue the new drug. Indeed, if it actually pays $1,391 million for the new product, it will *destroy* value by accepting a risk-free return (4%) on a project that is clearly risky. The fact that the board is prepared to invest in T-bonds if it can't buy the new product is irrelevant to the NPV calculation. Does this mean Pharmaco's board shouldn't compare an investment in the new drug to an investment in Treasury bonds? Of course not. But to make the comparison, one must evaluate the T-bonds at *their* appropriate risk-adjusted rate (4%) and the new product opportunity at *its* appropriate rate (8%). The two rates cannot be the same given the differences in risk.

The Pharmaco example illustrates how to calculate NPV and apply the NPV rule. NPV gives us a disciplined way to evaluate opportunities, one that takes into account both project risk and the time value of money and is consistent with value maximization.

Interactive Illustration 1 is a simple NPV calculator that lets you observe the effect on NPV of higher or lower discount rates and/or initial investments. The future cash inflows are fixed at $1,000 per year for five years. Set the discount rate at 8% and the initial investment at $1,000. The present value of expected future cash flows is $3,992.71. Since the required outlay today (Year 0) is $1,000, the NPV equals $3,992.71 − $1,000 = $2,992.71. (Press RUN to see the calculation.) As you raise the required initial investment, the present value of the cash flows stays the same ($3,992.71) but NPV falls, actually becoming negative as the required investment rises past $3,992.71. This result illustrates an important point: A project may be very valuable, but it is not worth doing if it costs too much. In the same way, even a very ordinary, unexciting business can be an excellent investment if it can be acquired for a low enough price.

INTERACTIVE ILLUSTRATION 1 Net Present Value

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2GfCY21

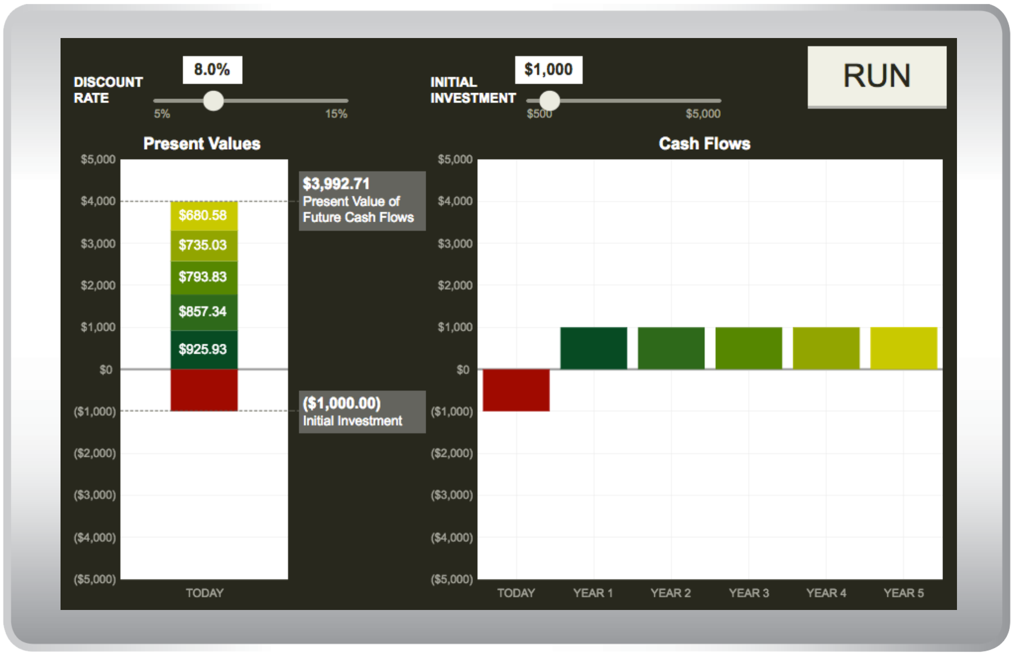

Returning to the illustration, keep the initial investment at $1,000 and observe the effect on present value and NPV with a change in the discount rate. As the discount rate falls from 8% to 5%, the present value of the future cash flows rises from $3,992.71 to $4,329.48—the future cash flows have higher present values at lower discount rates. NPV rises accordingly to $3,329.48. Now move the discount rate up to 15%. What happens? The present value and NPV fall significantly, but NPV is still positive. But when both the discount rate and the initial investment are high, say 15% and $5,000, respectively, the NPV is convincingly negative.

Interactive Illustration 2 provides a more general, more flexible NPV calculator in which you can manually enter cash inflows and outflows beginning today and for each of 10 future years. Then you select a discount rate and NPV will be calculated. The NPV Profile in the bottom graph shows the NPV for the cash flows you selected over a range of discount rates.

INTERACTIVE ILLUSTRATION 2 NPV Profile

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2pHPbWZ

Try using Interactive Illustration 2 to reproduce the NPV from the Pharmaco example presented earlier. The only outflow of $1,000 occurs in Year 0; inflows of $150, $180, $216, $259, and $311 follow for Years 1 through 5. The terminal value of $500 is included as an additional inflow in Year 5 (so the total inflow for Year 5 is $311 + $500 = $811). Cash flows for Years 6 through 10 are zero. The discount rate is 8%. As we saw earlier, the PV is $1,207 and NPV is $207. Using the slider and the input fields in Interactive Illustration 2, you can see what happens to NPV when any of the key variables change: higher cash inflows and/or a lower discount rate raise the NPV; a higher initial outlay and/or a higher discount rate lower the NPV.

2.2 Alternatives to NPV

We have seen how NPV is used as a decision criterion. But NPV is not the sole investment metric companies consider when allocating capital. What about alternatives to NPV? In this section we'll look at two of the most common ones, internal rate of return and payback period, to see how they work and how they compare to NPV.

2.2.1 Internal Rate of Return for a Single Risky Investment

A project's ***internal rate of return*** (***IRR***) is the discount rate at which its NPV equals zero. Alternatively, one may think of it as a "break-even" discount rate: the rate at which the present value of expected future cash flows is exactly equal to the required investment. To find the IRR, we simply compute NPV for a given project's cash flows and required investment using an arbitrary discount rate. If the resulting NPV is positive (negative), we raise (lower) the discount rate and recompute NPV repeatedly until we find the discount rate at which NPV equals zero.

IRR is an alternative to NPV because, in some cases, we could compute the IRR for a project and base a "go or no-go" investment decision on it rather than on the NPV. What is the associated decision rule? When the IRR of a project exceeds the appropriate risk-adjusted discount rate k (i.e., if IRR > k), the project should be accepted. The intuition is that when IRR exceeds k, the expected return on the project is greater than the opportunity cost of funds. In other words, IRR > k when NPV > 0, so the two rules should lead to the same investment decision. To see how this works, we will revisit the Pharmaco example and find the IRR of the proposed investment.

Exhibit 2 reproduces the Pharmaco cash flows and NPV calculation shown in Exhibit 1. Then, to find the IRR, Exhibit 2 recomputes NPV using a higher discount rate—higher because at k = 8% the NPV is strongly positive, so the IRR must be higher than 8%. By trial and error, or by using a financial calculator, we can discover that the Pharmaco project has an IRR of 13.7%. Using 13.7% as a discount rate, Exhibit 2 shows a PV of $1,000 million for the expected cash flows. Subtracting the initial investment of $1,000 million confirms that NPV equals zero.

EXHIBIT 2
Calculating the IRR of Pharmaco's Proposed Investment ($ in Millions)

Time	$t = 0$	1	2	3	4	5
Net inflows		$150.0	$180.0	$216.0	$259.2	$311.0
Terminal value						$500.0
$E(Cf_t)$		$150.0	$180.0	$216.0	$259.2	$811.0
PV @ k = 8.0%		$138.9	$154.3	$171.5	$190.5	$552.0
Sum of PVs	$1,207.2					
Investment	($1,000.0)					
NPV	$207.2					
PV @ 13.7%		$131.9	$139.2	$147.0	$155.1	$426.8
Sum of PVs	$1,000.0					
Investment	($1,000.0)					
NPV	$0					

Once again, to decide whether to invest or not, Pharmaco should compare the IRR of 13.7% to the opportunity cost of funds of 8.0% (some companies refer to the opportunity cost of funds as the ***hurdle rate***). Pharmaco should accept the project because IRR = 13.7% > k = 8.0%. For this simple project, and others that share some key attributes, the IRR rule leads to the same decision as the NPV rule. Unfortunately, this is not always the case; some project attributes can cause IRR and the IRR rule to be unreliable and perhaps even conflict with the NPV rule. So when (and why) might IRR lead us astray?

Problems Associated with IRR

Problem 1: There may be more than one IRR (or none)

The IRR is best suited for projects in which one or more cash outflows are followed by a series of cash inflows *with no further net outflows*. Typically, when the sign of the annual cash flows (positive or negative) changes more than once during the life of a project, there will be multiple IRRs. Further, the more sign changes there are, the more IRRs there typically are.[b] It is also possible that the IRR as we have defined it

[b] In algebra, Descartes's rule of signs is a rule for determining the maximum number of positive real number solutions (roots) of a polynomial equation in one variable based on the number of times that

will not exist. This will happen, for example, if a project has no outflows, just cash inflows. For such a project, IRR does not exist, even though NPV is clearly greater than zero.

To show how there can be multiple IRRs, we'll use **Interactive Illustration 3**, which is an IRR calculator. You can use it to find the IRR of a project by typing in the project's cash flows. The illustration shows the cash flows from the Pharmaco investment and its IRR of 13.7%. The IRR appears as the *x*-intercept in the graph of NPV as a function of the discount rate.

INTERACTIVE ILLUSTRATION 3 IRR

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2uowkFo

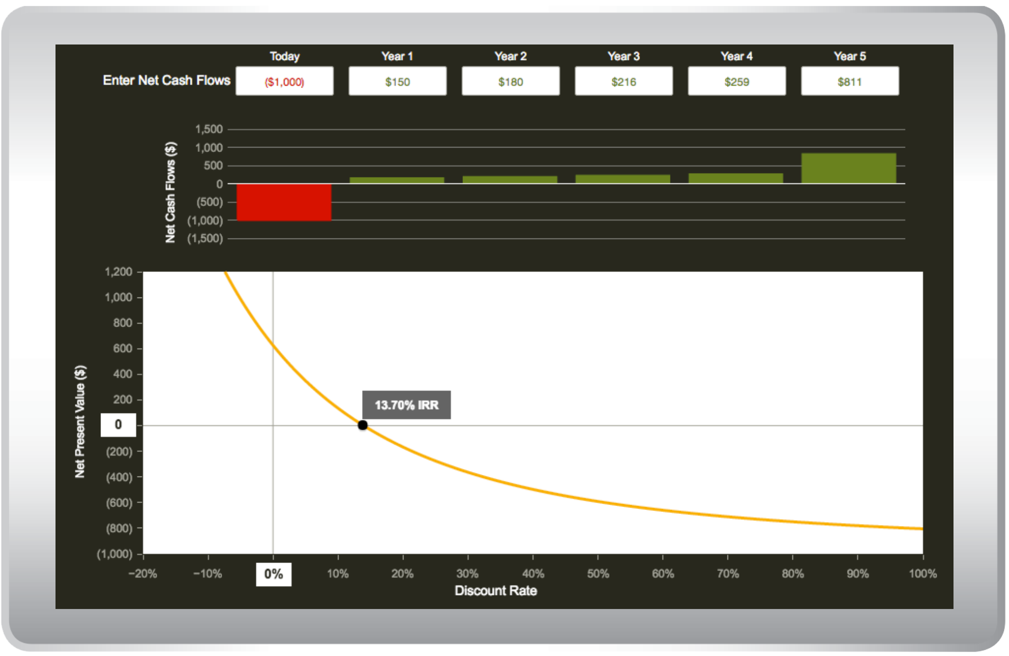

Now suppose the Pharmaco cash flows were a bit different. Specifically, suppose Pharmaco could acquire the rights to the new drug at a much lower price, $375 million instead of $1 billion (today). But in exchange for this price break, Pharmaco must pay R&D Ltd., the seller, a lump sum in Year 5 such that the total net cash flow for that year is negative: –$400 million. Type these two changes into the illustration and see what happens to the graph below the cash flows. Now there are two IRRs, one negative and one positive. Using either as a discount rate will yield an NPV of $0, but one is higher than k = 8% and one is lower. This makes it difficult, perhaps impossible, to safely use the IRR rule. In this case, a solution suggests itself: Why not

the signs of its real number coefficients change when the terms are arranged in the canonical order. The number of positive roots either equals the number of sign changes between consecutive non-zero coefficients or is less than the number of sign changes by an even number.

ignore the negative IRR as irrelevant and simply compare the positive IRR of 10.27% to the required expected return of 8.0%? This expedient happens to work in this case, but just barely, and it will not always be available. Note that the difference between the positive IRR of 10.27% and the opportunity cost of 8.0% might lead us to believe that NPV is comfortably positive. In fact, it is not: The NPV of these cash flows at 8.0% is $8 million, actually very close to zero considering the initial outlay of $375 million (to see this, type in the cash flows and compute NPV using Interactive Illustration 2).

Use Interactive Illustration 3 to experiment with different patterns of cash flows. You should be able to discover patterns for which there is no IRR or there are multiple IRRs. You also will find situations in which the shape of the NPV curve, and consequently the set of associated IRR(s), is extremely sensitive to one or a few inputs—that is, changing a single cash flow can cause a dramatic change in the shape of the NPV graph and its IRR intercepts.

Problem 2: The IRR decision rule changes when outflows follow inflows

Another problem occurs when cash inflow(s) in early years are followed by later cash outflows, even if there is only one change of signs. Consider **Exhibit 3**, which shows an initial $1,000 inflow followed by outflows of $150 for each of the next five years. Basically, the investor receives the benefit of this project immediately and pays for it over the next five years—a reversal of the basic pattern in the Pharmaco project. Now suppose the appropriate opportunity cost of funds is 8.0%. Typing these cash flows into Interactive Illustration 3 results in an IRR = −8.9%, also shown Exhibit 3. This IRR is clearly less than 8%; in fact, it is negative. Should the project be accepted?

EXHIBIT 3 The IRR Rule Is Reversed When a Project Inflow Is Followed by Outflows ($ in Millions)

	Cf_0	$E(Cf_1)$	$E(Cf_2)$	$E(Cf_3)$	$E(Cf_4)$	$E(Cf_5)$
Inflow	$1,000					
Outflow		($150)	($150)	($150)	($150)	($150)
Net cash flow	$1,000	($150)	($150)	($150)	($150)	($150)
IRR	−8.9%					

In this case, as the discount rate increases, the NPV increases as well. Because the investor receives cash up front and pays for it afterward, he or she really is better off if the discount rate is higher: NPV increases as k rises. This project's NPV is positive at all discount rates greater than −8.9%. To see this, return to Interactive Illustration 3 and examine the graph of NPV under the cash flows you entered. NPV rises rather than falls when the discount rate increases, and the horizontal intercept is negative. In short, when positive cash flow(s) precede a stream of outflows, the IRR rule must be *reversed*: a project should be accepted when IRR > *k*.

Proper Use of the IRR

We will return to uses and misuses of IRR later in this reading when we consider the problem of multiple, competing projects. At that point there will be additional problems to enumerate. In the meantime, some summary guidance on careful use of the IRR may be helpful:

- The IRR is best used when project cash flows exhibit a "conventional" pattern: one or more outflows are followed by a series of cash inflows and no subsequent outflows. This is the circumstance in which the IRR rule is equivalent to the NPV rule for a single project.
- In such cases, IRR calculations also may be an efficient way to gauge the sensitivity of NPV—and hence, the "go or no-go" investment decision—to the discount rate.
- The IRR should be used carefully or not at all when the cash flow pattern is unconventional, particularly when the sign of the cash flows changes more than once.

2.2.2 Payback Period for a Single Risky Investment

Another widely computed investment metric is the length of time it takes to recover the original investment in a project, known as the ***payback period***. When used as a decision criterion, the typical payback rule states that a project should be accepted if its payback period is less than or equal to some targeted length of time.

Payback appeals to managers on several levels. First, it is fairly straightforward and simple to calculate—an investment is either recouped within a stipulated period of time or it is not. Further, it seems like an intuitive way to acknowledge the time value of money—investors prefer to be repaid sooner rather than later. A shorter payback period may also give the impression, perhaps incorrectly, that a project with quick payback is less risky than one with a longer payback. However, payback is flawed as a decision criterion for several reasons, and this is true even when there is only a single project under consideration.

Problems Associated with Payback Period

Problem 1: Choosing a target payback period

There is no economic basis on which to choose the targeted payback length, or "hurdle." How long a payback period is too long? We cannot say because payback relies on *timing* rather than *value creation* as the basis on which the decision is made; this makes it unlikely that any payback-based rule would be value maximizing. Just because a project has a shorter payback does not guarantee that it will create more wealth; it is easy to imagine longer-term investments that create much more value than shorter-term ones.

Problem 2: Payback is myopic

Regardless of the target payback period selected, the payback calculation itself ignores all cash flows that occur *after* the point when the initial outlay has been recovered, which makes it shortsighted. Imagine two projects, each requiring a $1,000 outlay. Both have expected cash inflows of $500 in each of the first two years, so both have a two-year payback. But one project terminates at that point while the other continues to generate cash in perpetuity. Clearly the latter is much more valuable, but payback period will not recognize any of this extra value. In fact, if the latter project's inflows were delayed by even a single year, its payback period would be longer than the first project's despite being much more valuable than the first.

Problem 3: Payback ignores risk and time value

A further obvious objection is that the payback calculation itself ignores both the time value of money and risk. Sometimes adjustments are made to the calculation in an effort to address this shortcoming, but these still do not make payback a reliable investment criterion.

To illustrate these problems we will examine the payback calculation, with and without some common refinements, once again using the Pharmaco example introduced earlier. To compute payback period we accumulate the net cash flows year by year to see when the total exceeds Pharmaco's $1 billion investment. So in the first year, for example, the expected cash flow is $150 million, which offsets part of the $1 billion investment; the cumulative net cash flow after Year 1 is −$850 million. **Exhibit 4** presents the full set of cumulative cash flows and shows that Pharmaco will not recoup its investment until Year 5. If management required a payback period of, say, three years for investments of this type, this deal would be rejected despite its positive NPV.

EXHIBIT 4 Payback Period Calculations for Pharmaco's Proposed Investment ($ in Millions)

	Cf_0	$E(Cf_1)$	$E(Cf_2)$	$E(Cf_3)$	$E(Cf_4)$	$E(Cf_5)$	Terminal Value
Payback Period:							
Outlay	($1,000)						
Inflows	$0	$150	$180	$216	$259	$311	$500
Net cash flow	($1,000)	$150	$180	$216	$259	$311	
Cumulative cash flow	($1,000)	($850)	($670)	($454)	($195)	**$116**	

Interactive Illustration 4 presents these cash flows and the associated payback period, and allows you to modify the cash flows to see the corresponding change in payback period. Note that the terminal value of $500 is not included among the annual cash flows, even though it clearly contributes to the project's NPV. Why not include it in the payback calculation? Because we don't know when the associated cash flows occur. The terminal value of $500 represents the value as of Year 5 of all the post–Year 5 net cash flows. But to compute payback period, we need to know the annual cash flows for Years 6, 7, 8, and so forth. For concreteness, suppose the cash flow in Year 5 was only $11 instead of $311 (insert this value in the appropriate box in Interactive Illustration 4). Now the cumulative net cash flow at the end of Year 5 is –$184, so the project will not have earned back the initial investment of $1,000. In other words, the payback period is greater than five years, but we don't know how long the payback will take unless we know annual cash flows for a few more years at least.

INTERACTIVE ILLUSTRATION 4 Payback Period

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2DX0UFu

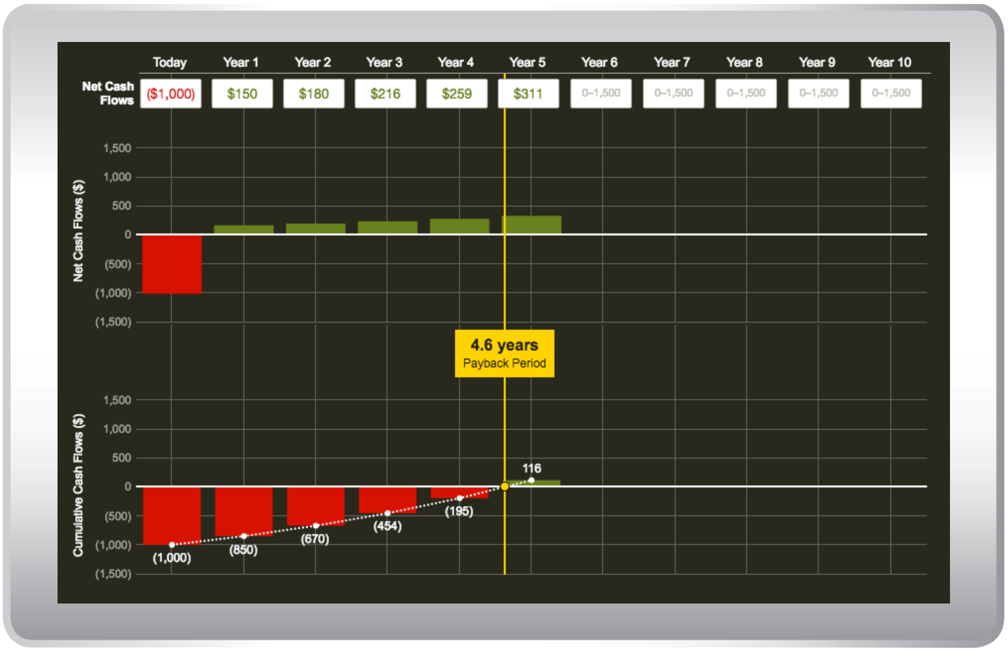

Practitioners sometimes refine the payback approach by discounting the cash inflows to present value before accumulating them to compute the payback period. This approach is called *discounted payback period* and it is shown in the bottom half of **Exhibit 5**. This change does introduce time value and risk into the calculation in a limited way, but it doesn't improve the project's chance of being accepted. On the contrary, the discounted payback period is *even longer* than undiscounted payback, so this valuable project *still* would be rejected whether the target payback period was three years or, even more generously, five years. In short, payback period is a flawed decision criterion even if it is modified to include discounting, and even if target periods are extended as a result—there still is no sound basis on which to choose a target.

EXHIBIT 5 Comparison of Payback Period and Discounted Payback Period for Pharmaco's Proposed Investment ($ in Millions)

	Cf_0	$E(Cf_1)$	$E(Cf_2)$	$E(Cf_3)$	$E(Cf_4)$	$E(Cf_5)$
Payback Period:						
Outlay	($1,000)					
Inflows	$0	$150	$180	$216	$259	$311
Net cash flow	($1,000)	$150	$180	$216	$259	$311
Cumulative cash flow	($1,000)	($850)	($670)	($454)	($195)	**$116**
Discounted Payback Period:						
Discounted cash flow @ k = 8.0%	($1,000)	$139	$154	$171	$191	$212
Cumulative discounted cash flow	($1,000)	($861)	($707)	($535)	($345)	($133)

In summary, neither payback period nor discounted payback period provides much information about value creation. As such, neither is a satisfying or reliable metric for making investment decisions. Nevertheless, some managers use it and, as we will see further below, in some real-world settings payback period may add incrementally helpful information to a complex investment picture.

2.3 NPV, Project Risk, and Expected Return

In general, investment choices may be analyzed in terms of either values (expressed in units of currency) or returns (expressed in percentages). An investor considering whether to buy a share of stock, for example, may ask, "What is the highest price I am willing to pay for this share?" Or, equivalently, "What is the lowest expected return I am willing to accept on this investment?" These two questions lead to the same investment decision. This should sound familiar; we saw earlier that in some circumstances the NPV rule and the IRR rule are equivalent. But then we noted potential problems with IRR (e.g., sometimes there may be no IRR or multiple IRRs). We avoid these problems if we restate the IRR rule as an expected return rule: Invest if the expected return on a project exceeds its required expected return.

2.3.1 Comparing Expected Return to Required Expected Return

Computing investment returns at the end of the investment period is a backward-looking exercise: It answers the question, "How did the investment perform?" But at the time the investment is made, its future performance is uncertain. Investors must base their decisions on forward-looking metrics: *expected future values* or, again equivalently, the *expected return*, denoted as $E(r)$. For a given project, the expected future value is the probability-weighted average of all possible future values; the expected return is likewise the probability-weighted average of all possible returns.[c]

Now look again at our expression for NPV:

$$\text{NPV} = \text{PV} - \text{investment} = \sum_{t=1}^{T} \frac{E(Cf_t)}{(1+k)^t} - \text{investment}$$

Where does the project's expected return appear in this expression? It doesn't, at least not as such. Instead, we have the discount rate k, which is the *required* expected return—the lowest expected return we would be willing to accept on the project. By setting the discount rate equal to the required return, we are computing a reservation price for the asset—the most we would be willing to pay for it. This is how we obtain the functional equivalence noted above between "the most we would be willing to pay" and "the lowest expected return we would accept." When the most we would be willing to pay equals the initial investment, NPV is zero. An NPV of zero does not mean the investor earns a zero return; it means the project's expected return equals the investor's required expected return: $E(r) = k$. When the project's expected return exceeds the required expected return, NPV > 0. So another way of stating the NPV rule is: "*Invest in all projects for which* $E(r) > k$." This is obviously similar to the IRR rule, but expressed in terms of expected returns instead of IRR (and so avoids some of the pitfalls of IRR).

Expressing the NPV rule in terms of $E(r)$ and k is helpful because it allows us to illustrate graphically a general relationship among NPV, project risk, and expected

[c] This expected return may be equal to the IRR, but is unique and well defined (as long as there is an initial investment) even in situations when there are multiple or no IRRs. To see this, imagine that the project in question is traded, like a share of stock, and so has a market price. (A project with a positive NPV is one that can be built or bought for less than the market price.) The expected return over a single period can be computed given expected cash flows during the period and the expected market price at the end of the period. Expected return may be computed in this way even when the pattern of project cash flows causes problems with IRR. For more on computing expected return in this way, see *Core Reading: Risk and Return* (HBP Nos. 5220 and 8603).

return. In **Exhibit 6**, expected return, E(r), is on the vertical axis and risk is on the horizontal axis.[d]

EXHIBIT 6

Investment Projects Plotted According to Expected Return and Risk

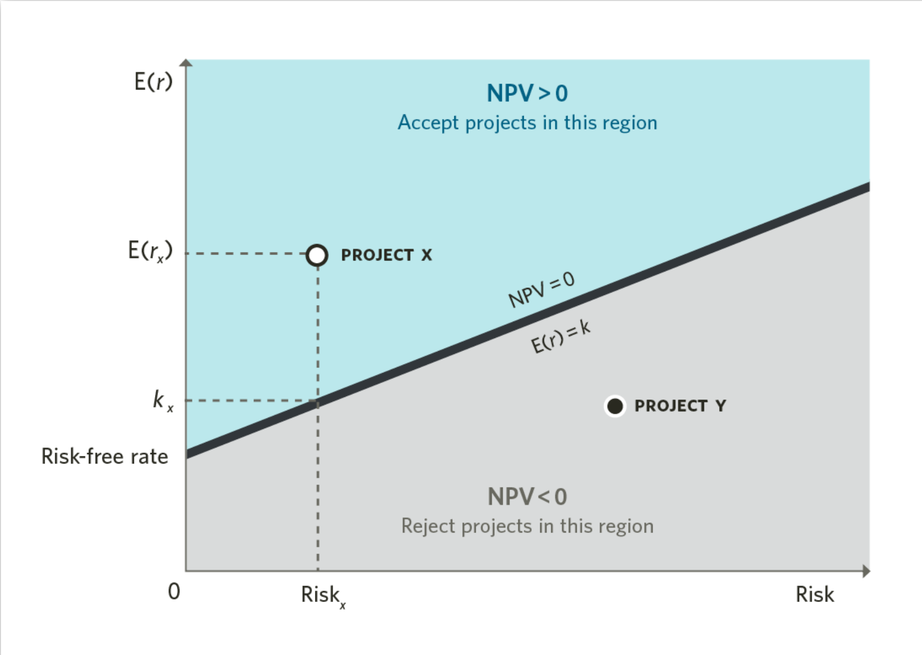

Now any investment opportunity may be plotted in this space according to its risk and expected return. The upward-sloping line describes a boundary between positive and negative NPV investments. On this line, k = E(r) and NPV = 0. Above the line E(r) > k so NPV > 0; below the line, E(r) < k and NPV < 0. Specifically, consider Project X, characterized by a certain degree of risk (call it $Risk_x$), for which the appropriate discount rate is k_x. By comparison, E(r_x), the expected return on Project X, is greater than k_x. In other words, $NPV_x > 0$ and Project X plots above the line. (Note that our Pharmaco example, with its positive NPV, would also plot above the line.) By similar reasoning, Project Y plots below the line (E(r_y) < k_y) and its NPV is negative. In short, all projects plotting above the line have NPV > 0; all those below have NPV < 0. By plotting any large group of projects in this space, we can tell at a glance which ones have positive NPVs.

[d] We could have put k on the horizontal axis instead of risk because we know that k equals the risk-free rate of return plus a risk premium, and that the risk premium increases with project risk. Had we done so, the graph would look the same qualitatively. By placing risk on the x-axis, we highlight the more fundamental relationship between risk and expected return.

To use this graphical representation of the NPV rule, we simply need to characterize projects in terms of their risks and expected returns. We also need to identify the zero-NPV line. You may be wondering how this is done. Actually, quantifying the equilibrium trade-off between risk and expected return (estimating the slope of the line in Exhibit 6) is the subject of *Core Reading: Risk and Return* (HBP Nos. 5220 and 8603) which covers portfolio theory. For now, we need only two key characteristics of the line and we have touched on both already. The first is that it is upward sloping (we have not actually shown that the line is straight because we haven't precisely defined "risk" on the horizontal axis, but for our present purposes, the shape doesn't matter so we drew a straight line). Put another way, for zero-NPV projects, higher risk is always associated with higher expected return.

Second, the zero-NPV line in Exhibit 6 implicitly tells us how to estimate k, the discount rate in any NPV calculation. Recall that k equals the expected return on an alternative asset with identical risk. So to identify the zero-NPV line, we simply need to identify a group of alternative investments that spans the entire range of risk levels from low to high and that are fairly priced. By "fairly priced" we simply mean value equals cost, or $E(r)$ equals k, or NPV equals zero. Where can we find such a group of investments? In the capital market. Indeed, we may even use this idea to *define* the capital market: the set of actively traded assets priced so that NPV equals zero. Put another way, it is the set of actively traded, fairly priced assets. Why would we suppose NPV equals zero in the capital markets? Because buying and selling securities in the market is relatively cheap—transaction costs are low—and investors are constantly looking for positive-NPV trading opportunities; that is, for *mis*priced assets. Their efforts to exploit mispricing involve trades that continually push prices to zero-NPV prices (i.e., selling overpriced assets and buying those that are underpriced).

The important implication of the previous paragraph for corporate capital budgeting is that *the capital market matters*. Not because the treasurer uses it to raise new funds—some companies are self-financing and seldom if ever raise new external capital. Rather, it is because the capital market is where we find an alternative investment with the same risk. The capital markets represent, in large part, investors' opportunity set for competing investments. This is so even if a corporation is private rather than public, even if it is self-financing, and even if managers privately suspect the market is somehow "wrong." In short, the capital market matters because it embodies genuine opportunities to trade: to buy and sell all kinds of securities representing claims on all kinds of business assets. As such, it determines k, the opportunity cost of funds. Indeed, now we can be a bit more precise in our definition of the opportunity cost of funds: It is the expected return on an alternative asset with the same risk *and zero NPV* (i.e., a fairly priced alternative).

2.3.2 Single Versus Multiple Hurdle Rates

Many of the thorniest capital budgeting problems arise when companies have to evaluate multiple projects. This is especially the case when managers must choose *between* projects—that is, accept some but not all—either because some are mutually exclusive or because one or more corporate constraints prevent them from accepting all projects. We will return to those problems below, but first we must take an important detour and consider a simpler problem: how to compute and interpret NPV when a single company is evaluating multiple independent projects but does not face any binding constraints—that is, it is free to accept or reject any project based solely on that project's NPV.

In this setting in the real world, the question often arises: should all projects' NPVs be computed using a single companywide hurdle (discount) rate? Or should each project receive its own discount rate, possibly different from the others? Actually, we have already answered this question. Early in our discussion, we argued that the appropriate risk-adjusted discount rate, k, had to be different for two projects if one was riskier than the other. Otherwise, projects with the same expected cash flows but different levels of risk would have the same value. We also noted that value maximization requires that the risk-adjusted discount rate k equal the opportunity cost of funds—the expected return available on an alternative investment with the same risk and a zero NPV. Both of these observations lead to the conclusion that projects with different risk characteristics *must* have different discount rates.

Although theory is clear on this point, many companies nevertheless find it difficult to implement, and they use a single hurdle rate for almost all projects throughout the firm. Why? In some large organizations the practice arises from a desire to ensure fairness across different lines of business—operating managers want a "level playing field" and some argue that this requires that all investment opportunities face the same discount rate. They find it vexing and unfair when a favorite project is rejected while a competing project with a lower expected return is accepted. Other managers hold as an article of faith that the right discount rate is the firm's ***cost of capital***, and they reason that within a single firm there can be only one cost of capital and hence only one firm-wide discount rate. As we have already seen, however, this prescription leads to mis-valuation of projects with varying degrees of risk.

Some readers may still find this assertion puzzling. After all, if the firm's treasurer can raise funds at a cost of 10%, how could it ever make sense to apply a higher discount rate, say 12%? Put another way, isn't it true that any project with an expected return greater than 10% will create value for such a firm? No. Consider a new project that we will stipulate is riskier than the firm's existing business and let us suppose that the true opportunity cost for the project is 12%. If the project's NPV, computed at 12%, is positive, then it will indeed create value and this is true even if the firm overestimates the NPV by using 10%. However, if the NPV is positive at 10% but negative at 12%, then accepting the project will destroy value—the firm will be worth less for having accepted it than if it had rejected it. A final question: What if the treasurer is very good (or very lucky) at his or her job and manages to raise funds at bargain rates, say 3%? Should that cause us to reconsider discount rates? Again, no. It is true that raising funds cheaply is valuable and makes the firm more valuable. But the skillful treasurer created that value, *not the project*. We should give the treasurer a bonus, but still use only opportunity costs for discount rates. As a more dramatic example, suppose your rich grandmother gave you funds *for free*—at 0% interest—so you could invest in a business. Should you use a discount rate of 0% when you decide which business to invest in? Clearly not. If you did, you would find yourself overpaying for *everything*, which cannot be value maximizing.

How important is this problem of tying discount rates to project risk? To consider this, return to Exhibit 6, on which we plotted two projects according to their risks and expected returns. The upward-sloping line in the graph described zero-NPV projects and separated positive NPVs (above the line) from negative NPVs (below the line). What happens when a firm uses a single hurdle rate for all projects? In **Exhibit 7**, a horizontal line represents the single hurdle rate, k^*, used by such a firm. The shaded triangles are regions in which the single-hurdle-rate firm will make investment mistakes. The triangle on the left (shaded darker blue) shows errors of omission—positive NPV projects (those that fall above the zero-NPV line) that were nevertheless rejected by the firm. These are relatively low-risk projects that deserve a lower discount rate than the company's one-size-fits-all k^*; by using too high a discount rate, the firm incorrectly concluded that the NPV was negative. The triangle on the right side of the graph (shaded darker gray) identifies errors of commission—projects with negative NPVs that the firm accepted anyway, because its single hurdle rate is too low for these relatively risky projects. Using too low a hurdle rate led to the NPV being overstated and the project mistakenly accepted.

EXHIBIT 7 Using a Single Hurdle Rate May Lead to Decision Errors

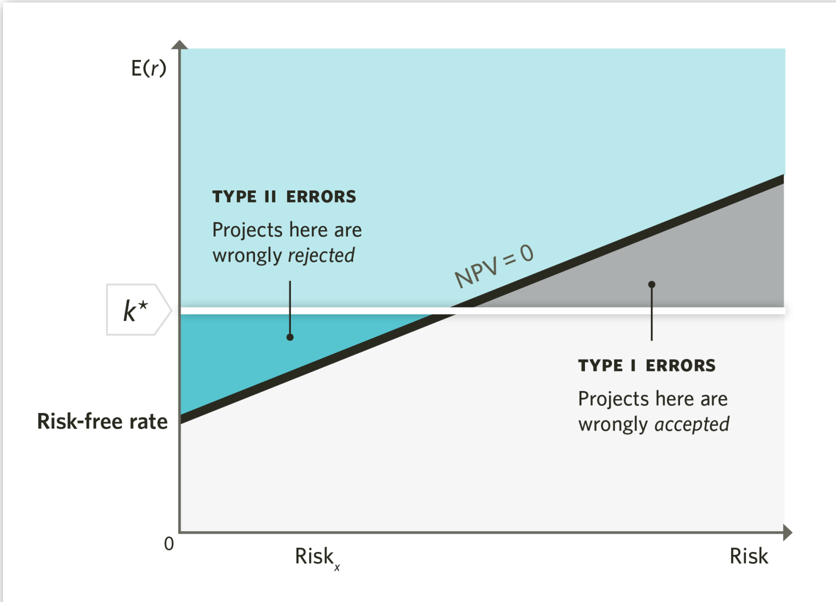

Interactive Illustration 5 allows us to experiment with this problem. It presents sample sets of 30 projects for consideration, representing random (but logical) combinations of risks and expected returns. Each point in the graph represents a project; some fall above the zero-NPV line and some below. Leave the slider for Project Risk Dispersion at its default setting, and focus instead on the slider for Hurdle Rate, which allows you to select a single firm-wide hurdle rate to be used in evaluating all 30 projects. As you raise and lower the selected single hurdle rate, note that none of the possible hurdle rates gets all of the investment decisions right. As we saw earlier, triangles on the left (shaded in darker blue) and right (shaded in darker gray) *almost always* contain some of the 30 projects—those that are mistakenly accepted (white) or rejected (black).

Interactive Illustration 5 quantifies the consequences of these mistakes. For each hurdle rate selected, it shows the accumulated value destroyed by mistakenly accepting negative-NPV projects and forgone by mistakenly rejecting positive-NPV projects. After experimenting with different single hurdle rates, click "New Set." This will result in a new set of 30 projects with different risk and return characteristics so you can repeat the experiment. It is unlikely that you will ever see an opportunity set for which a single hurdle rate does not result in some investment mistakes.

INTERACTIVE ILLUSTRATION 5

Capital Budgeting with a Single Hurdle Rate

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2IVa1Kx

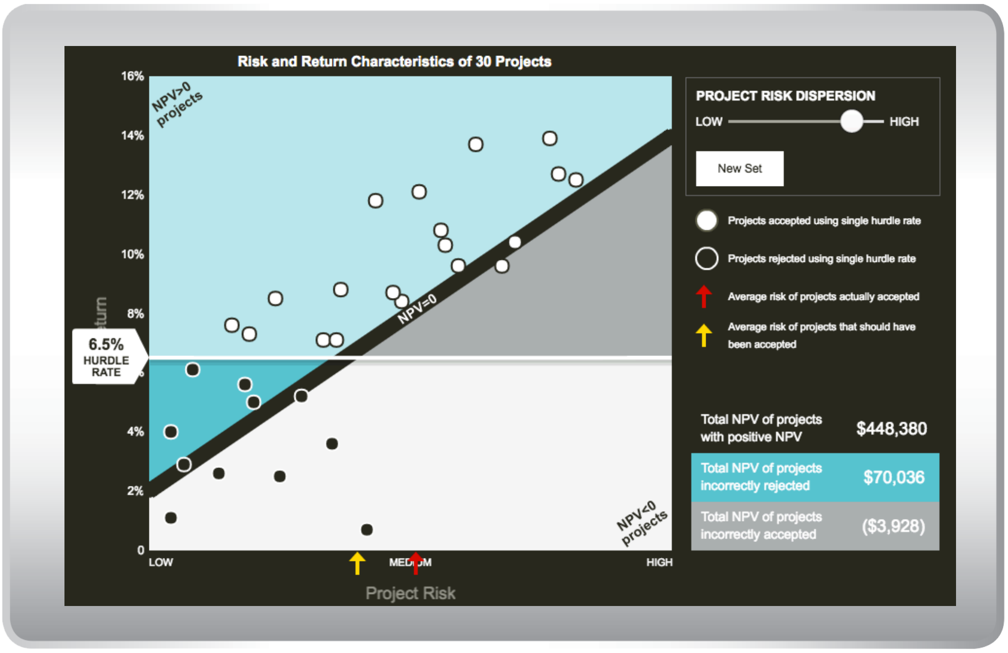

Does this mean that real-world firms are always wrong to use a single firm-wide hurdle rate? No. But doing so is only reasonable for firms facing a fairly homogeneous opportunity set—one containing only projects with virtually identical risk characteristics. To simulate this circumstance, set the "Project Risk Dispersion" slider to low and then click "New Set." Now you will see a set of 30 projects with different expected returns, but substantially the *same degree of risk*—they all pile up roughly vertically over the same point on the *x*-axis. Now it is indeed possible to find a single hurdle rate (perhaps more than one) that will make no investment mistakes (or very few). Use the slider to find it. You can do this very quickly by visual trial-and-error. Or, go the vertical "stack" of projects and find the point on the zero-NPV line that falls within the same "stack"—it has the same risk. Now simply set the single hurdle rate equal to the *y*-coordinate of this point. What is the lesson? Only firms whose projects all have the same risk can hope to maximize value by computing NPV using a single firm-wide hurdle rate. Firms facing even moderate dispersion in the risk characteristics of a large opportunity set will make some investment mistakes.

What are the long-term consequences of such investment mistakes? Once again, imagine a firm that confronts a variety of different types of projects but uses the same hurdle rate for all, as in Interactive Illustration 5 (slide the "Project Risk Dispersion" slider to high, and click "New Set"). Now imagine it follows this practice this year after year for an extended period. What should happen? We can make two general predictions. First, the accumulated investment errors represent a significant and growing departure from value maximization—underperformance. The firm is systematically destroying value with bad investments and failing to create it by forgoing good ones, and it is doing so year after year. You see this in the box in the lower right corner of Interactive Illustration 5. (The illustration does not show the NPV for each individual project, because the graph's axes are Project Risk and Return, but these projects do indeed have NPVs; they are used to compute the values in the lower right boxes.) Second, the firm as a whole should become *riskier* because of its systematic bias in favor of risky projects—the right triangle—and against safer projects—the left triangle. You see this by comparing the red and yellow arrows at the bottom of Interactive Illustration 5: As a group, the projects actually accepted are riskier than the group that should have been accepted. The combination of lower value and higher risk should be alarming to managers and shareholders alike.

3 SUPPLEMENTAL READING

Thus far we have examined different ways of evaluating a single project to decide whether or not to fund it today. This is an important problem, and the tools developed to address it are essential underpinnings for a disciplined, value-based approach to capital budgeting. But corporations face an additional, larger problem, which is how to evaluate a *set* of investment opportunities. Circumstances commonly arise in which a company must choose between projects, investing in some but not others, even though all may have a positive NPV when examined in isolation. Sometimes this is because investment funds are limited, so projects must compete for funds. Or managerial talent or "bandwidth" may be limited, so investment choices are tailored to management's tastes and talents. Or some projects may relate to others in fundamental ways—if two projects are by their nature mutually exclusive, managers need a means of choosing between them. All of these circumstances require that projects be examined together and evaluated relative to one another. The ideas already covered are still important—value maximization, project risk, required expected return, NPV, and so forth—but we need to extend them to the problem of relative valuation in the context of constraints and interdependencies.

3.1 A Set of Hypothetical Investment Projects

To examine the problems that arise when different projects must be evaluated relative to one another, we need some projects to compare. **Exhibit 8** presents a small group of hypothetical projects—generic investment opportunities—with both common elements and some simple differences. To focus on how these different projects are treated by common investment metrics, we will ignore strategic and other non-financial considerations; each project is completely specified by the required outlay, subsequent expected cash flows, and a discount rate corresponding to the project's riskiness. Exhibit 8 summarizes these characteristics for the six Projects, A through F.

EXHIBIT 8 Projected Cash Flows and NPVs for a Set of Six Hypothetical Investment Projects ($ in 000s)

Project	Discount Rate (*k*)	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	NPV	IRR
A	4.0%	($1,000)	$200	$250	$280	$300	$320	$191.8	10.1%
B	4.0%	($300)	$60	$75	$84	$90	$96	$57.5	10.1%
C	7.0%	($1,000)	$0	$0	$0	$0	$2,000	$426.0	14.9%
D	7.0%	($400)	$100	$150	$175	$180	$200	$247.2	25.5%
E	10.0%	($1,000)	$500	$700	($700)	$550	$600	$255.3	20.2%
F	10.0%	($500)	$620	$0	$0	$0	$0	$63.6	24.0%

What do Projects A through F have in common? They all receive expected cash flows for a period of five years after the initial investment. All require a cash outflow in Year 0—this is effectively the cost of the project. Three of the six require the same initial outlay of $1 million; the others—B, D, and F—are somewhat less expensive, with outflows of $300,000, $400,000, and $500,000, respectively. Despite these differences, all six projects have positive NPVs, ranging from a high of $426,000 (Project C) to a low of $57,500 (Project B). Further, IRR > *k* for all six; IRRs range from a high of 25.5% (D) to a low of 10.1% (A and B), but each IRR is greater than the corresponding opportunity cost of funds (discount rate *k*). For each of the six, the NPV rule and the IRR rule give the same decision: Accept the project.

What are the salient differences? The projects differ in their riskiness and hence, in their risk-adjusted discount rates. A and B are least risky, with a discount rate of 4%; C and D are somewhat riskier, at 7%; and E and F are the riskiest, with a discount rate of 10%. Although all the projects entail an initial outlay, the patterns of cash flows in Years 1 through 5 differ. The best way to see this is to examine bar graphs of the cash flows, shown in **Exhibit 9**.

EXHIBIT 9 Cash Flow Graphs for Hypothetical Projects A–F

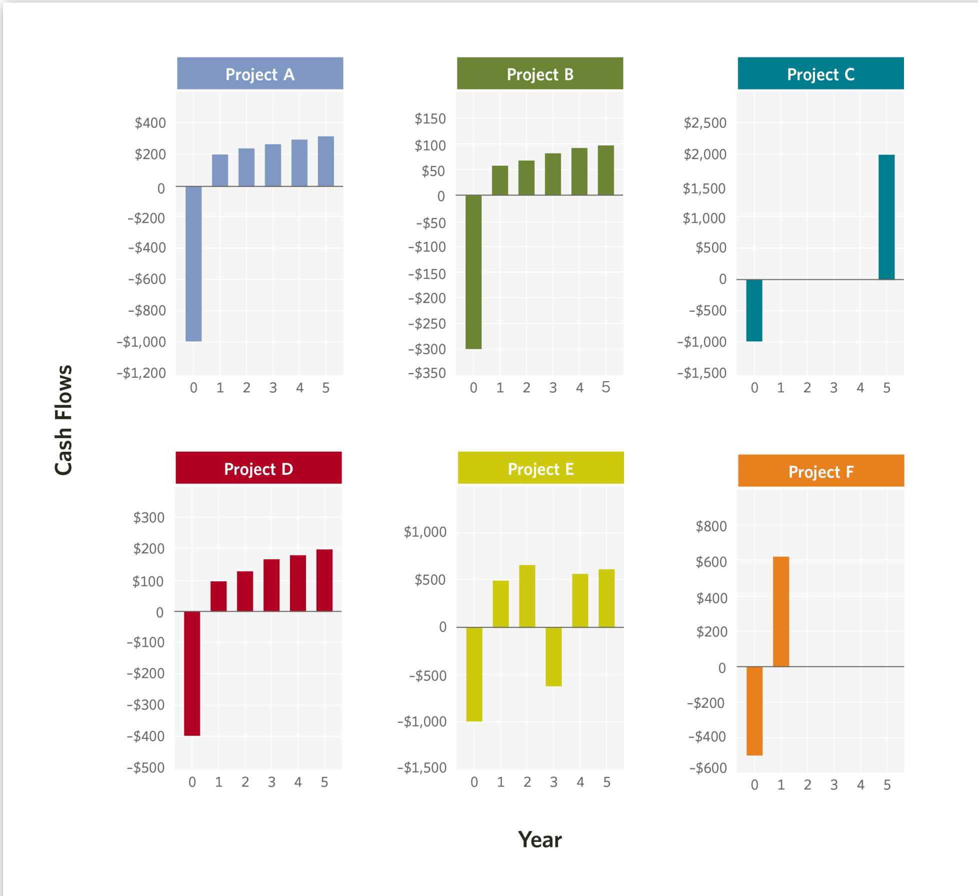

Note that Projects A, B, and D have similar patterns of cash flows: an outflow followed by inflows that grow over time. Projects C and F are similar: a single outflow followed by a single inflow, but the length of the period between inflow and outflow is different for each. Finally, Project E has an *interim* outflow: a negative expected cash flow in Year 3.

3.2 Capital Budgeting with Constraints

So far we have argued that NPV is a good way to measure the wealth a project may be expected to create and that a firm can maximize value by adopting the NPV rule; that is, by investing in all projects for which NPV > 0. In the real world, many companies do not do this. Specifically, they deliberately forgo some apparently positive-NPV projects. Among a number of possible reasons for this, we focus on one here: the existence of one or more ***binding constraints***. Perhaps the most common real-world constraint is a *budget constraint*—a limit on how much a firm can invest in a single period. When investment funds are limited, some valuable projects may go unfunded. The question then is, which one(s)?

Before exploring the problem of constrained capital budgeting, it is worth considering briefly why budget constraints exist. In a perfect world, one with perfect capital markets, they don't. A firm with a valuable project should be able to go to the capital market and raise money to fund it rather than forgoing it; valuable investment opportunities should attract financing. So budget constraints must be the result of one or more imperfections—ways in which real-world markets fall short of the perfect ideal. One example of a possible imperfection is *asymmetric information*: the situation in which managers inside the firm know more about the firm and its opportunities than do outside investors and everyone knows this is the case. In such a situation, the firm may be unwilling to reveal its private information (e.g., its trade secrets) for fear of helping competitors. Or perhaps it simply cannot convince laypeople in the market that its private information is accurate. In either case, the firm may not be able to raise funds at a reasonable cost for a project it knows to be valuable, so it skips it—that is, it underinvests. A quite different possibility is that the firm's managers and employees work harder and smarter and so perform better when resources are scarce. In this view, imposing a limit on funds for investment—even an artificial one—may force the firm to skip a few projects, but it makes all the rest perform better, enough better that the firm is more valuable with the constraint than without it. There are other possibilities as well, all of them involving some way(s) in which the markets or people aren't perfect.

3.2.1 A Simple Fixed Budget Constraint

For now we'll consider one type of constraint—a simple fixed budget—and not worry about why it exists. Suppose a firm is presented with the six projects, A through F. All have a positive NPV and, in an ideal world, all should be accepted and funded. However, the firm has only $2 million to invest. What is the value-maximizing investment decision?

An obvious possible procedure is to rank the projects by NPV—our favorite investment metric so far—and simply accept the project with the highest NPV first, then the next highest, and so forth until we run out of money. In this case we would run out of funds quickly. **Exhibit 10** includes project rankings for the set of six projects. Ranked by NPV, the project ordering is C, E, D, A, F, B. Projects C and E have the highest NPVs ($426,000 and $255,300, respectively) and they exhaust the budget of $2 million all by themselves. The total NPV from investing $2 million in C and E is $681,300. Would a different decision process produce a superior choice?

EXHIBIT 10 The Set of Hypothetical Investment Projects Ranked by NPV and IRR ($ in 000s)

Project	Discount Rate (*k*)	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	NPV	NPV Ranking	IRR	IRR Ranking
A	4.0%	($1,000)	$200	$250	$280	$300	$320	$191.8	4	10.1%	5
B	4.0%	($300)	$60	$75	$84	$90	$96	$57.5	6	10.1%	5
C	7.0%	($1,000)	$0	$0	$0	$0	$2,000	$426.0	1	14.9%	4
D	7.0%	($400)	$100	$150	$175	$180	$200	$247.2	3	25.5%	1
E	10.0%	($1,000)	$500	$700	($700)	$550	$600	$255.3	2	20.2%	3
F	10.0%	($500)	$620	$0	$0	$0	$0	$63.6	5	24.0%	2

As shown in Exhibit 10, ranking the projects by IRR and investing until the budget is exhausted yields a different set of decisions. Ranked by IRR, the project ordering is D, F, E, C, A/B. Investing in the projects with the three highest IRRs (D, F, and E) exhausts the budget (albeit with $100,000 left unspent). The total NPV from this set of investments is $566,100, not as good as the set chosen by NPV rankings.

As it happens, though, the optimal set of projects given a budget of $2 million is neither of the two just examined. Actually, it is D, C, and F. Together they cost $1.9 million and deliver a total NPV of $736,800, about 8% more value and with $100,000 less spent than the set determined by NPV ranking. If neither NPV nor IRR rankings reveal the optimal set, then how can we find it?

In this simple example, we can devise a fairly painless procedure to ensure a value-maximizing choice. To begin, we identify all combinations of projects that fit within our budget. **Exhibit 11** presents the set of feasible choices. Once we have identified the feasible combinations, it is a simple matter to see which results in the highest total NPV—the sum of the NPVs of the selected projects. Exhibit 11 reveals this is C, D, and F, as noted earlier, even though it does not consume the entire budget.

EXHIBIT 11
Feasible Project Combinations with a Budget of $2 Million ($ in 000s)

Project Combination	Required Investment	Total NPV
A + C	$2,000	$617.8
A + E	$2,000	$447.1
C + E	$2,000	$681.3
A + B + D	$1,700	$496.5
A + D + F	$1,900	$502.6
A + B + F	$1,800	$312.9
C + B + D	$1,700	$730.7
C + D + F	$1,900	$736.8
C + B + F	$1,800	$547.1
E + B + D	$1,700	$560.0
E + D + F	$1,900	$566.1
E + B + F	$1,800	$376.4
B + D + F	$1,200	$368.3

To find this result, we needed to identify and examine 13 possible combinations of the six projects (actually, there are more than 13 possibilities, but we did not bother examining certain choices, such as A alone, or C alone, etc., which are feasible but obviously inferior). Given the simplicity of this example, with its small number of projects and possible combinations, we can be confident that we have identified the optimal set of investments. But the process quickly becomes unwieldy with a large number of projects (large companies may have hundreds).[e] Clearly, it would be simpler to rank-order the six projects by some metric and then fund them in this order until the budget was exhausted. We tried this with both NPV and IRR and neither worked. Before going further, we take a brief look at why not.

[e] To appreciate how daunting the number of possibilities can be, note that the number of possible combinations of a set composed of N elements equals $2^N - 1$. So even a set of 10 projects generates more than 1,000 possible combinations of them; 20 projects leads to more than 100,000 possible combinations! Imposing a budget constraint reduces the number of *affordable* combinations, but also imposes the problem of identifying that subset.

3.2.2 Problems with IRR and NPV Rankings

Begin with IRR, for which we have already identified two problems (see Section 2.2.1). Now we can see an additional one.

Problem 3: IRR focuses on return, not value

When the problem is to choose between two projects, or between two versions of the same project, simply comparing IRRs can be misleading because it measures breakeven returns, but not value creation. For example, a firm might be offered a choice between two mutually exclusive investment opportunities, one with an IRR of 100% and one with an IRR of 10%. But it is possible that the project with the lower IRR actually creates more value. Or the projects may have the same IRRs, but their NPVs may vary widely. Our hypothetical Projects A and B illustrate exactly this phenomenon. Their cash flows, NPVs, and IRRs are reproduced in **Exhibit 12**.

EXHIBIT 12
Projects with Equal IRRs May Have Very Different NPVs ($ in 000s)

Project	Discount Rate (*k*)	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	NPV	IRR
A	4.0%	($1,000)	$200	$250	$280	$300	$320	$191.8	10.1%
B	4.0%	($300)	$60	$75	$84	$90	$96	$57.5	10.1%

By construction, every Project B cash flow is 30% of the corresponding Project A cash flow; Projects A and B are alike except that A is more than three times bigger than B. This is reflected exactly in their NPVs ($191,800 versus $57,500) but not in their IRRs—both have an IRR of 10.1%. Since the risk-adjusted discount rate is 4%, both projects will be accepted, whether we use the NPV rule or the IRR rule (10.1% > 4.0%). But what if A and B are mutually exclusive? That is, suppose we can undertake one or the other, but not both? Then the IRR rule is no help because the IRRs are identical. We need NPV to tell us about the scale of the projects, which is critically important if we must choose between them.

IRR can be misleading even for projects of the same size. **Exhibit 13** reproduces cash flows, NPVs, and IRRs for Projects C and E. These projects have the same scale, at least as measured by their required outlays. But subsequent cash flows' signs and magnitudes cause Project E to have the higher IRR, even though it has a lower NPV. (Note that Project E is shown with a single IRR of 20.2%, even though it has *two* changes of sign in its cash flows; this implies it may have additional IRRs—and it does—but others are very negative and clearly absurd from an investor's perspective; they represent no conflict with the decision to invest.)

EXHIBIT 13
IRR Can Be Misleading Even for Projects of the Same Size ($ in 000s)

Project	Discount Rate (*k*)	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	NPV	IRR
C	7.0%	($1,000)	$0	$0	$0	$0	$2,000	$426.0	14.9%
E	10.0%	($1,000)	$500	$700	($700)	$550	$600	$255.3	20.2%

We have seen NPV outperform IRR in other circumstances, so it is perhaps a bit surprising (and disappointing) that ranking Projects A through F by NPV did not reveal the optimal set of investments. There are two main reasons it failed. The first is that, while simple NPV rankings say something about *absolute* value creation across projects, they do not reflect the *relative efficiency* of value creation across projects, which may be critical with a binding budget constraint. By "relative efficiency of value creation" we mean how much NPV is delivered per dollar of investment. The second problem is ***lumpiness***, the situation in which the size of some projects is large relative to the size of other projects in the available set and/or to the size of the overall budget. Projects tend not to be divisible or perfectly scalable, so we cannot simply resize a project to make it fit the budget. When this unchangeable project size is significant relative to the budget, this "lumpiness" can interfere with otherwise sensible rankings of projects. We defer lumpiness for now and introduce a metric that will help with the first problem: relative efficiency.

3.2.3 The Profitability Index

To measure relative efficiency of value creation we can devise a metric—called the ***profitability index*** **(*PI*)**—that often (though not always) results in a useful ranking; that is, by choosing projects according to their profitability index ranking, we will maximize value subject to our budget constraint. Not surprisingly, the profitability index is simply a rearrangement of the components of NPV:

$$\text{Profitability Index} = \frac{\text{NPV}}{\text{Investment}}$$

The profitability index measures the *NPV per dollar of investment*, which is why it often works. The scarce resource is investment funds. The profitability index tells how much NPV each project generates per scarce investment dollar. In other words, it is a measure of how efficiently each project creates one dollar of NPV. Projects that have low NPVs may still be part of an optimal set of decisions if they don't use up much of the budget. The PI shows us how much "bang for the buck" we get from each project. (Note that the PI is usually calculated as NPV per dollar of *initial* investment, which makes sense for projects in which all significant investment occurs up front. Projects with delayed investments [such as our Project E] distort the PI and may limit its usefulness.)

Interactive Illustration 6 is a profitability index calculator. Using it, you can input cash flows and discount rates for two projects, P and Q, and compare their PIs. (Note that the initial investment must be input as a positive number between $300 and $1,000 in order for the PI to be calculated correctly.) The vertical axis shows each project's NPV, and the horizontal axis shows each project's required initial investment. Accordingly, the PI is given by the *slope* of a line passing through the origin and the point described by the *x*- and *y*-coordinates (Initial Investment and NPV, respectively) of each project. Higher PIs are indicated by steeper slopes. An additional visual indication of relative PIs is the size of the "bubble" associated with each project—the size of each project's bubble is proportional to its PI. What happens as you change the discount rate for a project? The PI changes: The project rises or falls in the space and the bubble around it grows or shrinks. To see the impact of negative interim cash flows on the simple PI calculation, try reversing the sign on Project Q's Year 3 cash flow from negative to positive.

INTERACTIVE ILLUSTRATION 6 Profitability Index

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2pHPcdv

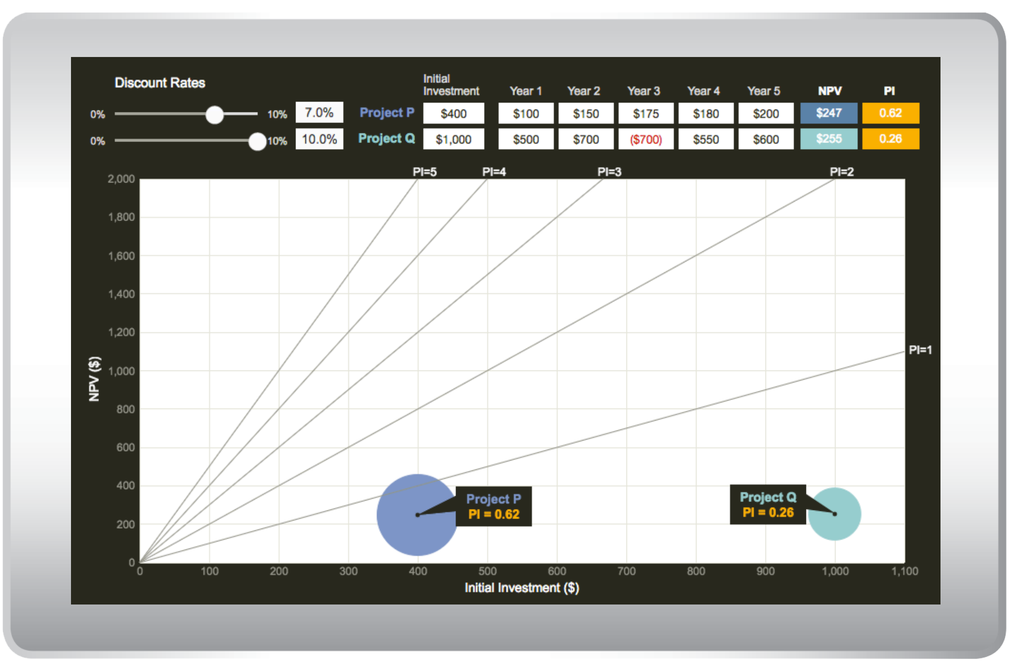

3.2.4 Limitations of the Profitability Index

Unfortunately, as we hinted earlier, the profitability index doesn't always lead so simply or reliably to the value-maximizing solution. In particular, it may not work perfectly when projects are lumpy, as described previously. In that case, it may actually be necessary to examine all feasible combinations and rank them by total NPV as we did in Exhibit 11. Fortunately, a high degree of lumpiness may also mean this exercise is not unreasonably difficult—lumpiness sometimes reduces the number of feasible combinations that must be examined.

Let's apply the profitability index to Projects A–F. Assume the budget constraint is still $2 million. **Exhibit 14** presents project characteristics and rankings by both NPV and PI. You can check the PI calculation for any project by inputting its cash flows and discount rate in Interactive Illustration 6. The first thing to note is that the NPV and PI rankings are indeed different from one another. Project D has the highest PI at 0.62 but is ranked third by NPV. Conversely, the highest NPV belongs to Project C, but its PI ranks second at 0.43.

EXHIBIT 14 Projects A–F Ranked by NPV and Profitability Index ($ in 000s)

Project	Discount Rate (*k*)	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	NPV	NPV Ranking	PI	PI Ranking
A	4.0%	($1,000)	$200	$250	$280	$300	$320	$191.8	4	0.19	4
B	4.0%	($300)	$60	$75	$84	$90	$96	$57.5	6	0.19	4
C	7.0%	($1,000)	$0	$0	$0	$0	$2,000	$426.0	1	0.43	2
D	7.0%	($400)	$100	$150	$175	$180	$200	$247.2	3	0.62	1
E	10.0%	($1,000)	$500	$700	($700)	$550	$600	$255.3	2	0.26	3
F	10.0%	($500)	$620	$0	$0	$0	$0	$63.6	5	0.13	6

Recall that if we select projects by NPV ranking until the $2 million budget is exhausted, we would invest only in C and E, and we discovered that this was not value maximizing. Rather, we found it was optimal to invest in C, D, and F, which require expenditures of only $1.9 million and produce an NPV of $736,800 (= $426,000 + $247,200 + $63,600).

Would the PI rankings have given this result? Unfortunately no, due to lumpiness, but they do point in a helpful direction. The PI rankings tell us that Project D is most efficient, with a PI of 0.62. Next most efficient is Project C, with a PI of 0.43. Projects C and D together use up $1.4 million of the budget. The next highest PI belongs to Project E (0.26), but E requires an investment of $1 million and we can't afford it with only $600,000 left to invest. Only by recognizing the lumpiness in the opportunity set are we able to boost the NPV a bit further. With our remaining budget of $600,000, we can invest in either Project B or F (but not both). F has the higher NPV, so we add it to C and D for a total NPV of $736,800. Project F ends up in the value-maximizing portfolio, even though it is ranked last in PI and next to last in NPV. It is clear that this result is caused by the projects' lumpiness. It may be less clear that the amount of the budget constraint matters, too; had the constraint been $1.6 million or $2.6 million, for example, the PI rankings alone would have given the optimal set of investments.

Interactive Illustration 7 shows the interaction of the lumpiness inherent in Projects A–F with budget constraints of different sizes. It lets us examine which set of projects is value maximizing as the budget constraint varies over the entire relevant range from $0 to $4.2 million. To begin, note that the table at the top of the illustration contains data on the investment, NPV, and PI for all six projects. There are tick boxes for investment, NPV, and PI that allow you to rank-order the projects by the chosen metric. At first none of the boxes are checked, so the order is simply alphabetical: A through F. The bar chart below the table shows the optimal set of projects for each possible budget. Move the budget slider to the right and see what happens. As the budget changes, a different set of projects is highlighted in the chart as the optimal set for that budget. For example, set a budget of $2 million and the illustration shows that the combination of C, D, and F is optimal, just as described earlier. Now check the box for "PI" next to the table. The projects in the table are reordered from highest to lowest PI. At the same time, the projects that would be selected if PI were to be used as a decision criterion are highlighted. Notice that for a $2 million budget and PI as a selection criterion, the highlighted projects are D, C, and B. Why these three rather than D, C, and E (note that E's PI is higher than B's)? It is because E is unaffordable given the budget; so is A, which also has a higher PI than B does. The third selected project is B because it is the affordable project with the highest PI. But even this is not optimal, as the bar chart confirms: The chart shows C, D, and F as discussed earlier; we saw that F was a better final addition to the set than B, because F is affordable and has a higher NPV than B.

INTERACTIVE ILLUSTRATION 7

Capital Budgeting with a Budget Constraint

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2unLaf9

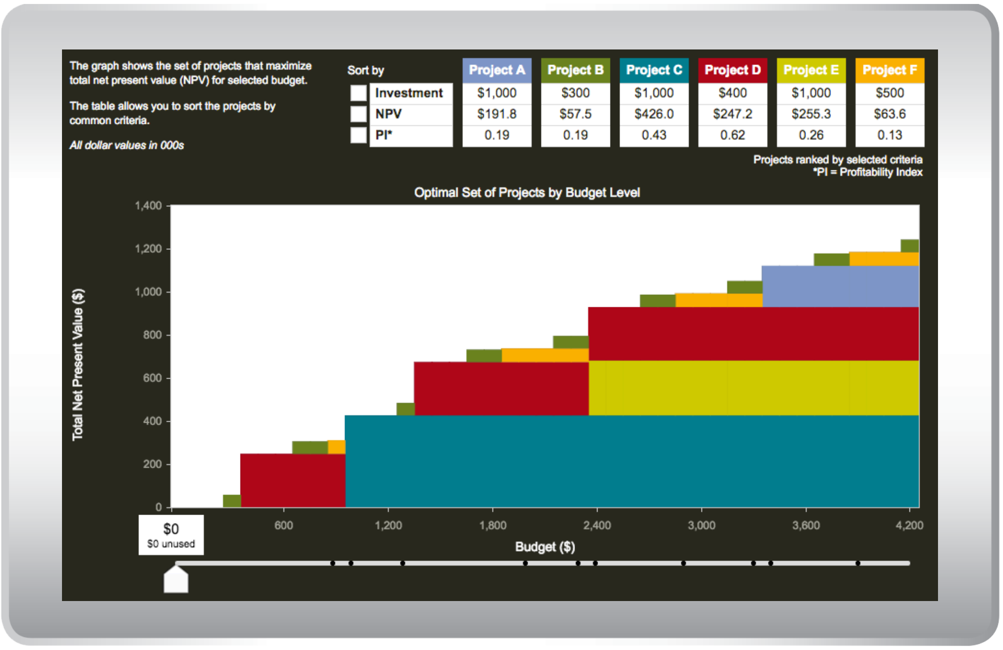

Use the slider and the tick boxes to examine different combinations of budgets and selection criteria. Notice the bar chart never changes—it always displays the optimal set of projects for a given budget. But the ordering of projects in the table changes when the tick box changes, and the selected (highlighted) projects for a given, narrowly applied criterion change continually with the budget.

Finally, there are some interesting transition points. For example, notice that with a budget of $2.3 million, the optimal set of projects to undertake is B, C, D, and F. Now what happens if we increase the budget by $100,000? The optimal set becomes C, D, and E. In other words, E replaces both B and F; we undertake three projects instead of four even though the budget has increased, and one of them, E, has not shown up previously for any lower budget amount. In fact, projects B and F do not both show up again in the optimal set until the budget reaches $3.2 million. A similar swap occurs when the budget goes from $3.3 to $3.4 million: Projects B and F are replaced in the optimal set by Project A. They do not both show up again until the budget reaches $4.2 million, large enough to afford all six projects—in other words, above this level the budget constraint is no longer binding.

Interactive Illustration 7 shows how lumpiness combines with a budget constraint to make project selection difficult, or at least unobvious. An even tougher challenge arises when firms can make ***intertemporal trade-offs***, such as when one or more projects may be saved or deferred from one period to a subsequent period, even if a budget constraint operates in both. Or when cash inflows from previous investments can augment the budget in subsequent years, thereby relaxing a binding constraint. Or when there is more than one constraint that might bind in a given period or for a given combination of projects.

To illustrate such a situation, consider Projects A–F again, but now suppose that opportunities can be saved from this year to the next; that is, any project we skip this year will be available on the same terms next year. Suppose further that this year's budget constraint is still $2 million. We don't know next year's budget, but next year's cash flows from this year's investments will be available for reinvestment. Now, which projects should we invest in this year? The best choice is D, E, and F (previously, it was C, D, and F). Together they consume almost all of this year's budget—$1.9 million—and produce total NPV of $566,100. This is less than the total NPV of $736,800 in the previous example, but note that D, E, and F produce cash flows of $1.22 million in Year 1 (refer to Exhibit 14). This is enough to afford a subsequent investment in Project C, *regardless of next year's new budget*. By investing in D, E, and F first, we *also* get to invest in C the following year, for a total NPV of $566,100 + $426,000 / 1.07 = $964,231. (We discount Project C's NPV for one year at its risk-adjusted rate of 7% because of the delay in undertaking it.)

This example shows another limitation of the profitability index. It is a helpful tool with which to select projects when the budget is constrained in a single period, but it may not work when managers have the scope to make choices across multiple periods.

3.3 Further Complications and Some Limitations of NPV

NPV and the NPV rule work well for deciding whether to invest in a given project at a single decision point. This is an important, useful result and the NPV rule outperforms other simple rules because it takes into account time value and risk while imposing value maximization. Nevertheless, capital budgeting in the real world is characterized by complications that may make NPV difficult to compute and/or make the simplest version of the NPV rule unreliable. Unfortunately, the list of complications is long: Managers in large organizations face many proposals, not just a few; projects may be interdependent and characterized by contingencies; companies face potentially binding financial and non-financial constraints; capital budgeting occurs over many years and intertemporal trade-offs may be not only possible but also commonplace and essential; and managers face significant uncertainty about both future opportunity sets and future constraints. It is, in short, a very complex environment. In the sections that follow, we will examine some of these real-world problems briefly and qualitatively to begin to develop an idea of the limitations of NPV and the NPV rule.

3.3.1 Interdependent Projects

Some opportunities have physical, financial, or managerial characteristics that make them interdependent. Some interdependencies are simple; for example, suppose two projects are mutually exclusive—the firm may undertake either but not both. In this situation, neither project should be evaluated in isolation. In other words, both must be evaluated to make a proper decision, but the NPV calculations may be otherwise straightforward. In contrast, consider two projects that are complementary or synergistic: If we undertake Project J, it enhances the expected cash flows from Project K, and perhaps vice versa. Each project is worth more if undertaken jointly with the other. Once again, value maximization demands that neither project be considered solely in isolation. The same is true for "anti-synergies," which might be present if, for example, one project cannibalized sales from another project or even from the existing business. Usually such interdependencies make NPV harder to calculate rather than making it irrelevant or misleading. But NPV may become *much* harder to calculate; we can imagine large numbers of interdependencies, not only across investment proposals but also between proposals and existing businesses and even across divisions in large corporations. Such interdependencies may be poorly understood and seldom evaluated rigorously. The resulting errors in estimated NPVs may make the NPV rule far from infallible.

3.3.2 Contingent Projects

Some investment opportunities beget more opportunities. A company that invests in a new technology may find it leads to other new technologies. A new product introduction may lead to later expansions into multiple geographic markets. A product that becomes a "hit" with consumers (a movie, for example) may lead to the creation of subsequent, related products (a sequel to the movie or product tie-ins). All of these are examples of what financial economists call ***real options***.

A real option shares some fundamental characteristics with financial options. A call option on a share of stock is a simple financial option: The owner of the call has *the right but not the obligation* to buy a share of stock at some future date at some specified price. The share of stock is the call option's *underlying asset*. For real options, the underlying asset is a *real* asset—a factory, a patent, a movie sequel, or a business, for example—rather than a financial asset, such as a share of stock. But the fundamental optionality is still important: the idea that the owner of the option has the right but not the obligation to do something with the underlying asset (e.g., buy or sell it) at some future date.

Full description and analysis of real options are well beyond the scope of this reading. But we want to note the challenges to NPV and capital budgeting generally that real options present. The first is simply that NPV may be a lot harder to calculate because of the important contingencies. In effect, it may be very hard to estimate expected future cash flows, especially if the real option happens to be embedded within another otherwise conventional project. But, second, even if we could correctly compute expected cash flows, the present value expression we have used in our NPV calculations would not perform the discounting correctly; that is, it would not correctly incorporate the effect of the contingency on present value. To get the present value right, we need an option-pricing model instead.

The final problem associated with optionality concerns the NPV rule rather than the NPV calculation. If a real option has not yet reached its expiration—we are not yet at the point where we must invest or lose the opportunity to do so—then the simple NPV rule doesn't guarantee value maximization. The NPV rule says invest if NPV > 0. For a call option, this is the same as saying, "Invest when the value of the underlying asset is greater than the exercise price"—in other words, when the asset is worth more than it costs. But an unexpired option may actually be worth more if we do *not* exercise it—that is, if we continue to hold it as an option or sell it as such to someone else. If our project is a real option, or contains real options, the NPV rule must be supplemented. To maximize value, we must compare the value of exercising *now* (per the conventional NPV rule) to the value of waiting and holding the option unexercised—saved, in effect, for a *later* go or no-go decision. Option pricing is well beyond the scope of this reading, but will be covered in *Core Reading: Financial Options* (HBP No. 5197). The main point is that NPV alone is too simple to ensure value-maximizing decisions involving options.

3.3.3 Intertemporal Trade-offs

When companies can choose the *timing* of an investment in addition to the simple go or no-go decision, the problem of maximizing value becomes more complicated. This is especially true in the presence of budget or managerial constraints. But it is also implicit in the presence of real options, as we just saw. Earlier in the reading we saw an example of a project (Project C in Exhibit 14) whose feasibility in the future depended on choices made now. It was available one year into the future, but we were only able to afford it if we also accepted Projects D, E, and F right away, which generated the cash needed to fund Project C the following year. Knowing this information at $t = 0$ enabled us to optimize both the combination and the timing of projects from the available opportunity set.

Other examples of intertemporal problems or trade-offs include: deciding whether or not to accept a project with multiyear spending requirements (in the face of uncertainty about future years' budget constraints); deferring a project for later consideration (this may differ from a real option in that it doesn't necessarily involve uncertainty about the future value of the project); and choosing a project with a long time horizon or "tail" in the face of unknown future budget constraints (and the related potential impact on future investment opportunities). All of these complications make NPV (and the profitability index in cases involving budget constraints) much more difficult to use with confidence.

3.3.4 Budget and Managerial Constraints

Budget and managerial constraints are a reality for many, perhaps most, companies. Regardless of size, companies have limited funds for investment, or behave as if this is the case, and only so much managerial time and talent to devote to new initiatives, including the problem of deciding which ones to undertake. We saw earlier that the profitability index, which is derived from NPV, is helpful when there is only one constraint, one budget cycle, and projects are not terribly lumpy. We mention constraints again here because, when overlaid with the complications described just above, they make capital budgeting very challenging—much more difficult than simply separating positive NPVs from negative ones, or rank-ordering by PI.

To close the reading we mention two further possible features of constraints that make their effects on capital budgeting decisions potentially significant and, unfortunately, difficult to analyze rigorously. First, future constraints may be *stochastic*, a phenomenon to which we have already alluded. That is, even if this year's budget constraint is known, subsequent years' may not be. When future constraints are uncertain, intertemporal trade-offs are harder to analyze and more difficult to manage, but they may be extremely important. Second, many companies face constraints that are endogenous, by which we mean that the constraint depends on other variables that are themselves part of or related to the firm's capital budgeting system. This sounds complicated, but there are simple versions of it that are fairly common. For example, systems in which a division's capital budget depends on its recent operating performance are not unusual. Division managers are told, in effect, "The better you perform, the more you get to spend." Some of the effects of such a constraint may be obvious—projects with quicker paybacks may be attractive not because they themselves are more valuable but because they help relax a future budget constraint sooner or to a greater extent. Other effects may be less obvious or less benign; for example, such a constraint might lead to repeated deferrals of valuable long-horizon projects in favor of those that enable greater near-term spending. Such behavior may be mitigated (or reinforced) by other corporate systems such as incentive compensation policies. In short, although some degree of endogeneity is nearly universal among large corporations, it is all but impossible to generalize about its effects on value maximization.

4 KEY TERMS

binding constraint A constraint, such as a budget, human resource, or management-determined constraint, that effectively prevents investment in otherwise desirable (i.e., value-enhancing) projects.

capital budgeting The process by which companies evaluate opportunities and allocate capital; the process by which companies choose which projects to fund.

contingent project A project for which feasibility or value depends on the status or characteristics of some other project.

cost of capital In the context of capital budgeting, the cost of capital usually refers to the discount rate (including both the time value of money and an appropriate risk premium) used to discount expected future cash flows to present value for inclusion in the NPV. For a value maximizer, this cost of capital equals the *opportunity cost of funds.* (Note that in other contexts, this term may mean something different.)

expected cash flows The mean of the probability distribution of uncertain incremental cash flows at a future date.

hurdle rate The rate chosen by a company for comparison to a project's IRR and which the IRR must exceed for the project to be accepted; alternatively, the discount rate selected by a company for use in an NPV calculation.

internal rate of return (IRR) The discount rate at which the NPV of expected future cash flows equals zero.

intertemporal trade-offs Trade-offs reflecting the possibility of changing the timing of an investment or a decision about an investment; for example, the ability to forgo an investment now in favor of making it later. This is in contrast to trade-offs between projects at the same point in time.

lumpy opportunities Opportunities of a discrete size, relatively large compared to other projects and/or to the firm's budget constraint.

net present value (NPV) NPV equals the difference between the present value of expected future cash flows and the immediate required investment (or the present value of required investments if they are spread out over time).

NPV rule An investment criterion: Invest in all projects for which NPV > 0.

opportunity cost of funds In the context of capital budgeting, the opportunity cost of funds equals the expected return that can be earned on an alternative project with the same risk and time horizon, and for which NPV = 0.

payback period The time it takes for a project to earn back its required investment. *Discounted payback period* discounts the cash flows to present value before accumulating them to compute the payback period.

present value The value of expected future cash flows discounted to the present at an appropriate risk-adjusted rate of return.

profitability index (PI) NPV per dollar invested (NPV/investment). A similar metric may be constructed for other constraints, such as people, as NPV per required person, for example.

real option The right but not the obligation to buy or sell a real asset, such as a business, a technology, a patent, etc. Real options are analogous to financial options, such as puts and calls, but involve real rather than financial assets.

required expected return The lowest expected return at which investors are willing to undertake a risky investment. *See also* opportunity cost of funds.

risk premium The portion of a risk-adjusted discount rate that compensates investors for bearing risk. The risk premium equals the difference between the risk-adjusted discount rate (or expected return) and the risk-free rate.

risk-adjusted discount rate A discount rate composed of the time value of money (the risk-free rate) and a premium for the systematic riskiness of the cash flows being discounted.

terminal value The discounted value at a terminal horizon (a selected point in the future) of all expected cash flows generated by a project beyond the terminal horizon.

5 NOTATION

Cf_t	Cash flow at time t
$E(Cf_t)$	Expected cash flow at time t
$E(r)$	Expected return
IRR	Internal rate of return
k	Risk-adjusted discount rate
NPV	Net present value
PI	Profitability index
PV	Present value
r	Risk-free rate
T	Terminal point

6 PRACTICE QUESTIONS

Scan this QR code, click the image, or use this link to access the interactive illustration: bit.ly/hbsp2IY5CGH

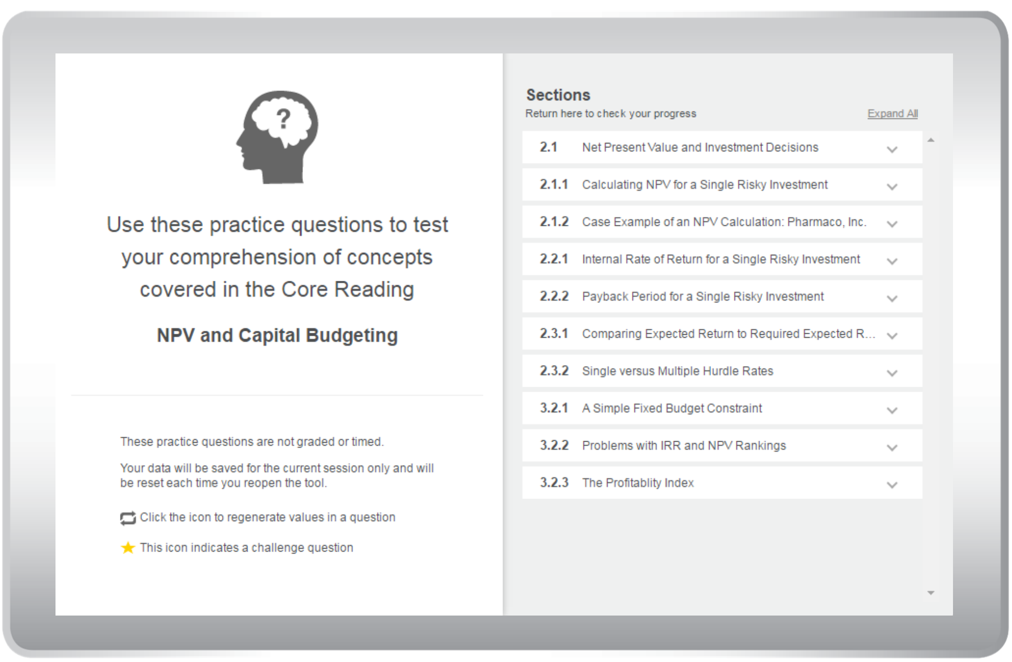

7 INDEX

anti-synergies, 41

binding constraint, 31, 35, 39, 40, 43, 45
budget constraint, 31, 35, 37, 38, 39, 40, 43

call option, 42
capital budgeting, 3, 45
contingent project, 45
cost of capital, 24, 45

discounted payback period, 19, 46
discounting, 3, 5, 8, 19, 42

expected cash flows, 6, 8, 12, 17, 24, 29, 30, 41, 45
expected future values, 21
expected return, 21, 22, 24, 26

hurdle rate, definition, 13, 45
hurdle rate, examples of single rate, 25, 26, 27
hurdle rate, single versus multiple, 24

internal rate of return (IRR), 3, 12, 45
intertemporal trade-offs, 40, 43, 45

lumpy opportunities (lumpiness), 35, 38, 45

managerial constraints, 43

net present value (NPV), 3, 4, 45
NPV alternatives, 11
NPV calculation, for a range of discount rates over 10 future years, 10, 11
NPV calculation, for a single risky investment, 5
NPV calculation, for effect of higher or lower discount rates and/or initial investments, 9, 10
NPV calculation, in Pharmaco case example, 7
NPV limitations, 40
NPV rule, 4, 5, 9, 13, 16, 20, 21, 29, 31, 34, 40, 42, 45

opportunity cost of funds, 6, 7, 9, 12, 13, 15, 23, 24, 46
option pricing, 42

payback period, 16, 17, 46
portfolio theory, 7, 23
present value, 5, 46
profitability index, calculator, 36
profitability index, definition, 35, 46
profitability index, limitations, 37

real option, 41, 46
required expected return, 6, 20, 21, 46
risk-adjusted discount rate, 6, 12, 24, 29, 34, 46
risk premium, 6, 7, 46

simple fixed budget constraint, 31

terminal value, 7, 8, 11, 18, 46

value creation, 4, 17, 20, 34, 35