

Capital Budgeting Applications and Pitfalls

Tips and Tricks

Applying the concepts of NPV and IRR in the real world can be very difficult. This chapter explains many of the nuances and pitfalls in their application. It will help you avoid many common mistakes that many companies commit almost every day — mistakes that cost them value.

13.1 So Many Returns: The Cost of Capital, the Expected Rate of Return, the Internal Rate of Return, and the Hurdle Rate

Before we begin, let us just recap the four commonly used rates of return in finance: the *internal rate of return*, the *cost of capital*, the *expected rate of return*, and the *hurdle rate*. A common conceptual problem for students (and even some faculty) is that they end up being the same number in a perfect capital market. (This assumes an *arms-length* transaction. Of course, you can give your cousins a sweet-heart deal at a different price if you want to.) The four numbers are *not* the same in the real world. We need to tease out the difference.

1. Cost of capital: This is the rate of return managers have to offer to entice investors to give them funds to undertake project.

When managers have to raise funds externally, then this cost of capital is determined by the rate of return investors could earn in similar projects elsewhere. Therefore, it is driven by the demand and supply for capital in the economy — the expected rate of return that investors can obtain. A good way to remember this is to mentally attach the word “opportunity” — the opportunity cost of capital. Also, always remind yourself that the cost of capital is not a promised rate of return but an expected rate of return.

Because US debt markets are very competitive, with even small firms being able to approach dozens of banks and request quotes, we often think of competitive loans as having zero net present value. Firms’ cost of capital comes from paying the “required expected rate of return.”

Large publicly-traded firms also enjoy a competitive equity-purchasing market, where they can obtain high and fair prices for their shares. This may not be true for smaller firms that live in imperfect capital markets. They may have to accept higher costs of capital, perhaps from friends and family or from venture capitalists.

In the real world, some large firms are sitting on a hoard of cash, with diffuse investors being unable to force entrenched managers to disgorge the money back to them. In this case, such managers may consider their costs of capital to be lower than the competitive rate — not necessarily in the interests of their investors.

2a. Expected rate of return: The expected rate of return is a generic term, but in the corporate context, it usually means your project's expected rate of return. As such, it has nothing to do with capital markets and all with the machines, suppliers, customers, etc., of the project. (Exception: for a financial services firm, such as a bank, the projects are themselves financial and thus related to capital markets.) The “internal rate of return” of a project is often a near-synonym for this project's expected rate of return.

2b. Internal rate of return (IRR): The internal rate of return is a characteristic of project cash flows (hence “internal”) and, as already noted, usually has nothing to do with capital markets. This is IRR's big advantage — you can calculate it before you ever look at the cost of capital (itself often based on external capital markets). The IRR is a descriptive summary measure with an internal focus that is based on the payoff patterns of the project.

It is very important to remember that you must use *expected* cash flows and not *promised* cash flows to calculate the IRR. Admittedly, expected cash flows are not as easy to assess — but (possibly empty) promises are really not what you would ever care about.

A project creates value for the firm if its internal (expected!) rate of return exceeds its cost of capital. This is what makes it a positive NPV project. If you are the manager running a good firm in the best interest of the owners, you should take all projects whose expected rate of return is above their cost of capital. As you work your way down the list of good projects, the very last marginal project should have an expected rate of return just about the same as the cost of capital. Projects before the last marginal one are sometimes called infra-marginal.

The hurdle rate is quite different:

3. Hurdle rate: The hurdle rate is a choice that managers decide on. It is the expected rate of return above which management decides to accept and go forward with the project. Good management should set a hurdle rate that makes them accept all projects that have positive net present values. Typically, managers set hurdle rates above what they believe their cost of capital to be. Bad management could choose any arbitrary, or even outright idiotic, hurdle rate.

A lot of confusion about the hurdle rate arises from the common professorial assumption that they are just explaining concepts — the situation in which managers behave as if they operate in a perfect market and perfect world. In a perfect capital market, perfect managers should set projects' hurdle rates to be equal to their costs of capital. Thus, all three numbers — the cost of capital (the external rate of return), the internal expected rate of return, and the hurdle rate — would be identical.

In the real world, this is not the case — not even in large companies that face near-perfect capital markets. Managers have to deal with flawed employees. They may have enough self-awareness to realize that they have rose-colored views of their projects' chances of success or profits — or that their subordinates are self-interested and want to get as much funding for their favorite undertakings as possible. Agency problems (This is indeed a pervasive problem!) In smaller companies, the capital markets are not so perfect, either, and even

a small miscalculation could potentially lead to the firm's demise. "Better safe than sorry" in such cases means choosing higher hurdle rates.

A natural way to deal with many real-world issues is to be conservative and set a hurdle rate that is above the cost of capital. Managers then pass up projects that seem to have rates of return higher than what they could finance them with, i.e., projects with a mildly positive presumed NPV. John Graham's 2019 survey of good-sized publicly-traded corporations — taken just before Covid — suggested that the typical hurdle rate was about 400-500 basis points above the cost of capital. **Begin Important**

Firms typically choose hurdle rates that are about 4-5% higher than their weighted average cost of capital (WACC). **End Important** The most important impact on their hurdle rates were their cost of capital estimates — every time their weighted average cost of capital goes up or down, the hurdle rate shifted likewise. Thus, the 4-5% wedge remained fairly constant.

Venture capitalists regularly employ project hurdle rates as high as 30%. They understand themselves well enough to know that their average project return will be far worse than what they expect it to be. This forces them to be more conservative in their investment selections.

(Omitted solvenow)

13.2 Promised, Expected, Typical, or Most Likely?

By now, you know that you must always distinguish between promised and expected numbers. In particular, models like the CAPM are about expected rates of return and simply do not tell you anything about credit risk. When you want to apply the present value formula, you must use the *expected* cash flows in the numerator (adjusted for credit risk), not the *promised* cash flows. When it comes to your risk judgment, it goes into the PV numerator first. Never, ever discount promised cash flows with (CAPM) costs of capital!

Promised and Expected Returns

Let's recap this difference. Say the world is really as perfect as the CAPM suggests and you have a B-rated corporate zero-bond that promises \$1,000 next year and has a beta of 0.2. Assuming you believe the risk-free rate is 5% and the equity premium is 3%, you can still not compute the bond price as

(Omitted eq)

Yes, in a perfect CAPM world, the expected rate of return on this bond should be $5\% + 3\% \cdot 0.2 = 5.6\%$. (In an imperfect world, you would have to add the liquidity and tax premiums.) Yet, to determine the price, it is not enough for you to know the *promised* bond cash flow. You need the *expected* cash flow, a number that is always less than \$1,000. The same problem arises, of course, not only in the context of bonds but also in the context of

corporate projects. You cannot simply discount the “good-scenario” cash flows. You must discount the project’s expected cash flows!

The same mistake appears sometimes in another form when managers use the IRR capital budgeting rule. This rule says “accept the project if its IRR is above the hurdle rate.” The common mistake here is that the cash flows from which the IRR must be computed are not the promised cash flows, but the expected cash flows. Of course, you can also compute a number from the promised cash flows, but you should probably call it the “promised IRR” to distinguish it clearly from the “expected IRR” — and you should never compare the promised IRR to a hurdle rate based on the expected rates of return of other projects in the economy when you want to determine whether you should accept the project or not. Always keep in mind that the promised IRR — the one that is most seductive because it is easiest to compute from numbers provided — should not be used for capital budgeting purposes.

(Omitted solvenow)

Expected, Typical, and Most Likely Scenarios

Managers often commit a related (but milder) error in applying NPV. They tend to confuse expected values with “typical” or “most likely.” (Statistically speaking, this means that they confuse the mean with the median or the mode of a distribution.) If you do this, you will fail to consider low-probability events appropriately: a plane crash, a legal suit, an especially severe recession, or a terrific new client.

For example, your business may have the following payoffs:

Event	Probability	Value
Good Business	46%	\$1,200,000
Normal Business	44%	\$1,000,000
Lawyers Sue for Punitive Damages	10%	−\$10,000,000

The most likely payoff is \$1,200,000. The median payoff is \$1,000,000. The expected payoff, however, is only

(Omitted eq)

An NPV analysis requires *this* expected payoff. If you run this business 100 times, you would on average receive \$1.2 million 46 times, \$1 million 44 times, and lose \$10 million 10 times. Fortunately, if the statistical distribution is symmetric — as it is in the case of the normal bell-shaped distribution — then the center of the distribution is all three: mean, median, and mode. Unfortunately, few businesses are immune to low-probability shocks, often negative, so you need to think about whether the distinction between mean, median, and mode is applicable to your business.

(Omitted solvenow)

13.3 The NPV of Linked Projects

How should you think of projects that depend on one another?

For example, think of buying a laboratory to develop a drug. The laboratory can be sold after the R&D stage is done. If and only if the drug development succeeds, then it can be marketed and sold. Say the laboratory costs \$10 million and has a cost of capital of 10% (it can be partly mortgage financed), of which \$8 million can be recovered in a sale after the development is done (unless something truly rare like a lab fire occurs). Assume that after one year, the lab work will be complete. The lab has a probability of going up in flames of 1% (accidents happen), and the drug has a 5% chance of working (as is typical for human drugs).

The drug, if it works, is expected to yield \$2.5 million in a sale to a large drug distributor. Of course, this being (legal) drugs, the scientists will emphasize that the drug can be sold to Pfizer for \$50 million if it works. Assume investments in drugs of this kind (in development or not) offer an average *expected* rate of return of 15%. Note that 15% is the expected return, and therefore already includes the probabilities of their failures, too. (In the context of the incorrect CAPM, the interpretation of 15% being above 10% could be the probably equally incorrect belief that drugs earn lower rates of returns when the market rate of return is low, e.g., because Medicare reduces reimbursements in dire times.) We will also assume that we will just sell the drug to Pfizer after one year of negotiations if it works.

Don't read the answer first. Instead, try to figure this one out by itself.

The answer is to consider the projects separately. In millions of dollars, the laboratory has a value of

(Omitted eq)

The drug has a value of

(Omitted eq)

You have to assume that you can only take the two projects together, leaving you with a sum-total negative NPV here. Don't take this project!

This type of payoff pattern in linked projects — where one project can only be taken if another earlier one has been successful — is common in many risky project contexts. For example, if you want to sell communication bandwidth, you first have to strap your satellite to a rocket and hope it won't blow up. If you want to sell a product, you first have to set up a research and production system. And so on.

13.4 Incorrectly Blended Costs of Capital

One of your first lessons about NPVs was that you can add them if projects are independent. Yet, believe it or not, although most managers know that it is impossible to add value by merely combining independent projects, in practice they often make exactly this mistake. This error arises most commonly in contexts in which costs of capital need to be blended across multiple projects, and especially when projects are financed with different levels of debt and equity. As always, the concept is straightforward, but the devil is in the details. It is easy to overlook the forest for the trees. Let's make sure you do not commit this mistake.

Does Risk Reduction Create Value?

Recall the insight from Section that companies cannot create value by reducing risk via diversification into multiple businesses. However, some mergers can add value due to synergies, which will be discussed in the next section. But these synergies are not a result of the plain diversification effect. Many researchers believe that the most common but unspoken rationale for mergers are not synergies but the fact that managers like to take over other firms. They prefer the reduced idiosyncratic firm uncertainty and higher salaries associated with larger firms to the higher risk and lower salaries in sharply focused, smaller firms. To justify a merger, managers will want to argue for a lower cost of capital for the target any way they can — including incorrectly using the acquirer's cost of capital. (This is an example of an agency conflict, which will be explained later in this chapter.) There is also good evidence that in the real world, diversified firms often do not operate as efficiently as stand-alone firms (e.g., due to limited attention span of management or more bureaucratization). Many mergers actually *destroy* firm value.

(Omitted anecdote)

► Does Corporate Risk Management Create Value?

Although risk management is discussed in more detail in the companion chapter on options, let me give you a brief preview. Firms can reduce their own overall risk by hedging. A hedge is an arrangement that reduces the firm's volatility. For example, a refinery could purchase crude oil today in order not to suffer if the future oil price were to increase. (This is further discussed in the companion chapter on risk management and hedging.)

Remarkably, a firm with a high cost of capital and risk could even transform itself into a firm with a low cost of capital! (Some hedge funds do this.) The firm can hedge away market risk by selling the stock market itself. [S&P 500 futures contracts](#) make shorting the stock market exceptionally easy. Shorting stocks Whenever the stock market goes up, the futures contract goes up in value. The futures contract sold by the hedging firm goes down in value. Put differently, the firm's hedge contract has a negative market exposure. The hedged firm is now a bundle, consisting of the unhedged firm plus this contract. Therefore, the market exposure of the hedged firm would be lower than the market exposure of the unhedged one. If it wished, the firm could even make its own market exposure zero or negative. Usually, being hedged against market risk would also reduce the firm's total risk (stock price volatility). Some firms may hedge against other risks. For example, Southwest Airlines has often purchased jet fuel far in advance (through futures contracts), though it is not altogether clear whether Southwest's intent was to hedge or to speculate.

But would this hedging contract create firm value in a perfect market? No. The firm has not given its investors a new positive-NPV project. If investors had wanted less exposure to the overall stock market, *they could have shorted the stock market themselves*. Alternatively, investors can simply undo a firm's hedging — they can buy the financial markets contracts that the firm has sold. This undoes any corporate hedge from the investors' perspectives. So, in itself, in a perfect market, trading fairly priced hedging contracts neither adds nor subtracts value. It is only if the market is imperfect that a hedge may allow a firm to operate more efficiently. For example, the extra cash from a hedge contract could help the firm to avoid running into a liquidity crunch in situations in which more funding would be difficult to raise. Or the firm may have inside information concerning what the future will hold and thus whether the hedged good is underpriced. In this case, risk management could add value.

Begin Important

In a perfect market, the following holds:

- If two firms are independent, then combining them into a conglomerate usually reduces the overall firm risk, but does not create value for investors. Investors can easily diversify risk themselves.
- Adding independent projects to the firm cannot create value if these projects are not positive-NPV in themselves.

In an imperfect market, the value effects of hedging are complex. Hedges could indeed add (or subtract) value.

End Important

(Omitted solvenow)

How to Misuse Costs of Capital

This brings us to a common simple NPV mistake: forgetting that the NPVs of independent projects are additive. Sounds obvious, but here is how it gets lost in the details: In a perfect market, NPVs are only additive if you use each individual project's own costs of capital. You cannot use the firm's overall cost of capital for its individual projects.

► When Acquiring Another Company

Your old firm, cleverly named *old*, is worth \$100 million and has a cost of capital of 5% (maybe because its business is mostly holding debt). At a fair price, it expects to pay off \$105 million next year. A potential acquisition target (or just a new project), cleverly named *new*, costs \$10 million this year, expects to pay out \$11 million next year, and has a cost of capital of 15%. The simplest method to compute the value of acquiring project *new* relies on the fact that the NPVs of independent projects are additive. You can value the new project using its own expected cash flows and its own cost of capital. Who owns *new* should matter little: The project is worth what it is worth. Therefore, in million dollars, the true NPV of project *new* is

(Omitted eq)

Therefore, if *old* adopts *new*, the original owners of *old* become \$430,000 poorer than they would have been otherwise. (If you want to practice the CAPM, think of a beta of 0.5 for the old project, a beta of 3.0 for the new project, a risk-free rate of 3%, and an equity premium of 4%.)

Unfortunately, in some firms, it is standard policy to evaluate *all* projects by the firm's overall cost of capital. Would such an *old* firm take the *new* project now? Evaluated incorrectly at a cost of capital of 5%, the *new* project looks a lot better, at $-\$10 + \$11/(1 + 5\%) \approx \$0.48$ million.

If the *old* firm did take project *new*, how would its value change? The true present value of the combined firm, in million dollars, would be

(Omitted eq)

This is \$430 thousand less than the original value of \$100 million plus the \$10 million acquisition cost of the new project. Taking *new* makes the *old* owners \$430,000 poorer.

Of course, not all acquisitions are driven by such mistakes. Don't make the mistake of reflexively thinking everything is a perfect market. Thus, it is not always true in the real world that mergers *never* add value on the cost-of-capital side. If capital markets are not as efficient for small target firms as they are for large acquiring firms, it would be possible for a large acquirer to create some value also on the cost of capital side. For example, if a target previously had no access to a perfect capital market, then the cost of capital to the target can change when it is acquired. The correct cost of capital for valuing the acquisition (the target), however, is still *neither* the cost of capital of the acquirer *nor* the blended post-acquisition cost of capital of the firm. Instead, the correct cost of capital is the lower rate that is appropriate for the target's projects, given the improved access to capital markets. For example, if an entrepreneur¹ inventor of holographic displays previously had faced a cost of capital of, say, 60%, primarily due to access only to personal credit card and credit-shark financing, and if this inventor's business is bought by Intel with its cost of capital of 6.5%, the proper cost of capital is neither Intel's cost of capital nor a blended average between 60% and 6.5%. Instead, once part of Intel, the holographic project division should be evaluated at a cost of capital that is appropriate for projects of the risk class "holographic display projects." This can add value relative to the 303% earlier cost of capital. (Of course, there are also many examples of large corporations that have destroyed all innovativeness and thereby all value in small companies that they had taken over.)

► When Acquiring Another Project

It is not only firms to be acquired, but also smaller or sub projects themselves that can have components with different costs of capital. For example, when firms keep cash on hand in short-term U.S. Treasuries, such investments have a lower expected rate of return. These bonds should not need to earn the same expected rate of return as investments in the firm's risky long-term projects. (The presence of this cash in the firm lowers the average cost of capital for the firm by the just-appropriate amount.)

Here is another application, which shows how you can decompose projects into categories with different costs of capital: Assume that you are considering buying a rocket to launch a telecom satellite next year. It would take you 1 year to build the rocket, at which point you would have to pay \$80 million. Then you launch it. If the rocket fails (50% chance), then your investment will be lost. If the rocket succeeds, the satellite will produce a revenue stream with cost of capital of, say, 13%, beginning immediately. (Telecom revenues may have a high covariance with the market.) The telecom's expected cash flows will be \$20 million *forever*.

The correct approach is to think of the rocket as one project and of the telecom revenues as another. The rocket project has only idiosyncratic risk. Presumably, its risk can be diversified away by many investors, its beta is close to zero, and it may have a discount factor that is close to the risk-free rate of return — say, both rates being 3%. The rocket value (in millions of dollars today) is

(Omitted eq)

You can think of this as the cost of storing the \$80 million in Treasuries until you are ready to proceed to your second project. The telecom revenues, however, are a risky perpetuity. With telecom-like costs of capital of 13% and cash flows that appear only if the rocket succeeds (a 50-50 probability), its value is

(Omitted eq)

Consequently, the combined project has an NPV of about $-\$1$ million. If you had mistakenly discounted the rocket's \$80 million cost by the same 13%, you would have mistakenly valued it at $-\$80/1.13 + \$76.9 \approx +\$6.1$ million.

(Omitted solvenow)

Differential Costs of Capital — Theory and (Agent) Practice

It is clearly correct that projects must be discounted by their project-specific costs of capital. Yet Graham and Harvey found in their 2001 CFO survey that just about half of surveyed CFOs *always* — and often *incorrectly* — used the firm's overall cost of capital rather than the project-specific cost of capital! And even fewer CFOs correctly discounted cash flows of different duration or riskiness within projects. (They sometimes do and sometimes do not take into account that cash flows farther in the future typically require higher expected rates of return — they should!) The easy conclusion is that CFOs are ignorant — and many CFOs may indeed incorrectly use a uniform cost of capital simply because they are ignorant.² sect:CFOdecide2001 CFO survey

Yes, they should use duration and leverage adjustments depending on the type of project, but what really saves the day (well, the company) is that most of them probably don't use their cost of capital estimates much. Many use a much higher hurdle rate than their cost of capital estimate. This limits the damage that this error can inflict

However, there is yet another reason. Even some intelligent CFOs use one-and-the-same discount rate quite deliberately on many different types of projects. Why? You already know that it can be difficult to estimate the appropriate cost of capital correctly. In theory, markets are perfect and we know the cost of capital. In practice, this may or may not be a good approximation. Do you really know the correct expected rate of return for projects of this specific type? (Do you really even know the correct expected cash flows? Remember — this is not physics where we understand all the driving processes from the mechanics of the spinning wheel.) In addition, you have not even yet considered such issues as the influence of liquidity and tax premiums¹ sect:liqtaxpremium Imperfect markets premiums on your cost of capital. Quite simply, you must be aware of the painful reality that our present value methods are usually just not as robust as we would like them to be.

Together, your uncertainties distort not only your overall corporate cost-of-capital estimates, but also your relative cost of capital estimates across different projects. Consequently, the problem with assigning different costs of capital to different projects may now become one of disagreement. Division managers can argue endlessly about why their projects should be assigned a lower cost of capital. Is this how you want your division managers to spend their time? And do you want your managers to play revenue games? Managers could even shift revenues from weeks in which the stock market performed well into weeks in which the stock market performed poorly in order to conjure up a seemingly lower market beta. The cost-of-capital estimate itself then becomes a pawn in the game of agency conflict and response — all managers would like to convince themselves and others that a low cost of

capital for their own divisions is best. What the overall corporation would like to have in order to suppress such “gaming of the system” would be immutable good estimates of the cost of capital *for each division and potential project* that no one can argue about. In the reality of corporate politics, however, it may be easier to commit to one-and-the-same immutable cost of capital *for all projects* than it would be to have different costs of capital for each division and project. This is not to argue that this one cost of capital is necessarily a good system, but just that there are cases in which having *one* systemwide cost of capital may be a lesser evil.

In sum, a good rule of thumb in real life is not to worry too much about differential costs of capital across projects of similar horizon and financing class, unless your projects are vastly different. (A good rule of thumb in job interviews is to understand what you must do in a perfect world, though — you will be asked. Make sure to answer that each project needs its own cost of capital.)

b Interview Questions:

[//www.amazon.nl/Heard-Street-Quantitative-Questions-Interviews/dp/1991155417/](http://www.amazon.nl/Heard-Street-Quantitative-Questions-Interviews/dp/1991155417/)

Errors: Do Projects Really Need Their Own Costs of Capital?

But let's not get carried away. Does every project *really* need its own cost of capital? Don't miss the forest for the trees. Yes, in theory, each component must be discounted at its own discount rate if you want to get the value (and incentives) right. However, in practice, if you want to value each paper clip by its own cost of capital, you will never come up with a reasonable firm value — by focusing on individual trees, you will lose track of the forest. You need to keep your perspective as to what reasonable and unreasonable errors are. The question is one of magnitude: If you are acquiring a totally different company or project, with a vastly different cost of capital, and this project will be a significant fraction of the firm, then the choice of cost of capital matters and you should differentiate. However, if you are valuing a project that is uncertain and long-term, and the project is relatively small, and its cost of capital is reasonably similar to your overall cost of capital, you can probably live with the estimation error. It all depends.

Begin Important

- Theoretically, all projects must be discounted by their own costs of capital, and not by the firm's overall cost of capital.
- In practice, most companies discount not by their cost of capital, but by a hurdle rate that is above their cost of capital. In an imperfect market, there are good reasons for doing so. It also moderates the impact of incorrect use of cost of capital estimates.
- Practically, the effort involved, the uncertainty in your estimates, the distraction from getting your expected cash flows in the PV numerator right, and the “gaming” by division managers may prevent you from discounting every project — every paper clip — by its own cost of capital.
- Depending on the situation, you may be better off assigning the same cost of capital to all cash flows of similar maturity, perhaps with only a modest holistic risk adjustment, based on such aspects as how it is financed, how long-term it is, and how perfect a capital market you are facing.

It is up to you to determine when it is important to work with different costs of capital and when it is better to use just one cost of capital.

End Important

13.5 The Economics of Project Interactions

If projects are independent, you have the luxury to consider them in isolation. You can compute separately the costs and benefits necessary to make a decision whether to accept or reject each project. However, in the real world, projects are not always independent.

Let's assume that you are the only person who can service a market and that you assess your potential profits in different states to be \$120,000 in NY, \$60,000 in CA and \$40,000 in PA *if* you enter only one of them. However, it may cost an extra \$70,000 to develop states on different coasts simultaneously, but the cost of developing two nearby markets may be sharable among neighboring states. For example, say that the potential profit is not \$160,000 but \$200,000 if you develop NY and PA. So, how do you select the best set of projects? (You could think about negative consequences, too. For example, if your best reseller in CT threatens to withdraw business if you develop either NY or PA [and even more if you develop both], you would have to figure this revenue loss into developing these two states.)

Begin Important

The ultimate project selection rule: Consider all possible project combinations and select the combination of projects that gives the highest overall NPV. **End Important**

Optimal project selection is easier said than done. It is easy for the basic example with these three states (take NY and PA, skip CA), but this is rarely the case. For two projects at a time, there are usually only 2^2 options to consider: take neither, take one, take the other, or take both. But the complexity quickly explodes when there are more projects. For three projects, there are $2^3 = 8$ options. For four projects, there are 16 options. For 10 projects, there are about a thousand options. For 20 projects, there are over a million options. For 50 projects, there are quadrillions. And even the simplest corporate projects can easily involve hundreds of decisions that have to be made. Mathematically, it is an impossible task to find the perfect combination.

To help you determine which projects to take, you need to find some rules that help you make a decision. Such rules of thumb are called heuristic[heuristics]—that is, rules that simplify your decisions even if they are not always correct. One common heuristic algorithm is to consider project combinations, one at a time. Start with the project combination that would give you the highest NPV if you were only allowed to take two projects (one pair from a set of many different projects). For example, start with the state that has the highest profit. There are only 50 of them. Now consider adding each state. There are only 49 possible choices. Then take this pair as fixed, that is, treat it as a single project. Now see which of the remaining 48 states adds the most value to your existing pair. Continue until adding the best remaining project no longer increases value. Computer scientists call this the greedy algorithm. It is a good heuristic, because it drastically cuts down the possible project combinations to consider and usually

gives a pretty good set of projects. There are many possible enhancements to this algorithm, such as forward and backward iterations, in which one considers replacing one project at a time with every other option. Full-fledged algorithms and combinatorial enhancements that guarantee optimal choice are really the domain of computer science and operations research, not of finance. Yet many of these algorithms have been shown to require more time than the age of the universe, unless you make simplifications that distort the business problem so much that the results are likely no longer trustworthy. Fortunately, finance is in the domain of economics, and economics can help simplify the project selection problem.

Project Pairs

Considering projects in pairs is not only common practice, but also clarifies many economic issues. With two projects, you can decompose the total net present value into three terms:

(Omitted eq)

For example, the original two state project (NY+PA) project choice yielded

(Omitted eq)

The final term reflects the interaction of the two projects. It suggests that you can classify project combinations into one of three different categories:

1. Projects with zero interactions
2. Projects with positive interactions
3. Projects with negative interactions

An interaction is also sometimes called an externality in economics, because one project has an external influence on another project — sometimes imposing external costs, and sometimes providing external benefits. Let's consider these three cases separately.

► Zero Project Interactions

Most projects in this world are independent — they have no mutual interactions. For example, for Walmart, opening a mall in Japan probably has no effect on opening a warehouse in Canada. Independent project payoffs permit separate evaluation of each project. This makes decision making much easier:

- Taking any positive-NPV project increases firm value.
- Taking a zero-NPV project leaves firm value unchanged.
- Taking any negative-NPV project decreases firm value.

If projects are independent, then the project interaction term is zero, and project NPVs are additive. Project independence makes decisions a lot easier: For 20 projects, only 20 independent decisions (accept or reject) have to be made, not a million.

Begin Important

You can simply add the project NPVs of independent projects. **End Important**

► Positive Project Interactions

Positive interaction[Positive interactions] mean that the sum of the parts is worth more than the parts individually. If one project has a positive influence on the NPV of another project, you cannot value it without taking into account this positive influence. For example, think of a new product as one project and of an advertising campaign as another. The advertising campaign project is of lesser use without the product, and the product is of lesser use without the advertising campaign. You must consider creating a product and an advertising campaign together. Such positive externalities are even more plentiful in smaller decisions. For example, a computer keyboard is less useful without a computer, and a computer is less useful without a keyboard. Many projects or products make sense only if bought together. In this case, producers may bundle them for their consumers.

In the corporate context, investment in *infrastructure* is another classic example of positive project interactions. For example, building a road, hiring a security firm, or laying a fast Internet connection could enhance the values of many divisions simultaneously. The firm should factor in the increase in value to *all* divisions when deciding on how much infrastructure to add.

Don't take positive externalities too lightly: On a philosophical basis, positive project interactions are the reason why firms exist in the first place. If there were no cost savings to having all resources combined in the firm, all of us could work as individuals and dispense with firms altogether.

Begin Important

When deciding whether to take an additional project, you must credit all positive interactions to this additional project. The overall NPV is higher than the individual project NPVs alone. **End Important**

Internal conflict and cost allocation procedures (discussed further as “agency conflicts” in Section ??) often hinder corporations from taking advantage of many positive externalities. For example, in real life, your division managers might argue that they should not be charged for the Internet connection, because they did not request it and therefore do not really need it (even if it were to increase their divisions' values). After all, division managers would prefer getting the Internet for free from the company instead of paying for it out of their own divisional budgets.

Nowadays, managers who want to acquire other companies usually claim the presence of large positive externalities. Synergies are the managerial term for positive externalities between an acquirer and a potential acquisition target. It has become an important managerial buzzword. For example, in the [announcement of the sale of Vauxhall and Opel](#), the parties explained

PSA Group and General Motors (GM) have come to an agreement over the sale of the latter's European business, specifically the Opel and Vauxhall brands, with the total price coming in at €2.2 billion (\$2.33 billion),

PSA is targeting €1.7 billion of annual synergies from the combination within ten years, with targets of Opel returning to profitability for the first time this century by 2020. PSA expects much of this savings to come from synergies in manufacturing, purchasing, the supply chain and R&D.

Of course, whether enough synergies are ever realized to outweigh the acquisition costs is yet another question. (Most acquirers have performed quite poorly after their acquisitions, never realizing the fantastic synergies which they originally touted.)

► Negative Project Interactions

Negative interaction[Negative interactions] mean that the sum of the parts is worth less than the parts individually. In this case, projects have negative influences on one another and thereby decrease one another's value. Economists sometimes call such negative externalities diseconomies of scale and/or scope. Here are a few examples.

Pollution and congestion: Think of an airline company with two divisions, but only one maintenance facility.

One division handles cargo; the other handles passengers. If the cargo division wants to expand, it will use more of the maintenance capacity. This will leave the passenger division with longer service waiting times. In the extreme, the extra delays may cost the passenger division more than the extra profits that the expanded cargo operation adds.

Cannibalization: If buying and installing a new Apple MacOS-based computer system can produce \$100,000 in NPV (e.g., from specialized software not available on Windows) compared to the older Windows machine that produced only \$70,000 in NPV, how should you credit the Apple machine? The answer is that the Apple would eliminate the positive cash flows produced by the existing Windows machine, so the cash flow of the project “replace Windows with Apple” is only \$30,000: the \$100,000 minus the \$70,000 that the now-unused Windows machine would have produced. Be careful what you consider cannibalization, though. For example, in the 1970s, IBM did not produce personal computers, fearful of cannibalizing its mainframe computer business. IBM's mistake was that it did not realize that other computer manufacturers were able to step in and eat much of IBM's mainframe business for themselves. Put differently, IBM had not realized that the present value of its mainframe business's future cash flows had already changed with the advent of new technology in the competitive market that it was in.

Complexity: As more and more projects are adopted, management will find it increasingly difficult to make good decisions, and do so in reasonable time frames. As you just learned, projects can often impact other projects, and no manager knows every project and cares about them in the right mix.

In trying to deal efficiently with more scale and complexity, larger organizations typically adopt more detailed processes and bureaucracy. The cost is that “*process*” itself consumes resources and can reduce cash flows for all divisions. A good example of bureaucratic destruction of projects can be found on Moishe Lettvin's blog. (To find the url, remember that Google is your friend). A programmer who worked for Microsoft for 7 years, Lettvin describes how it took between 24 and 43 people, separated by six layers of management, over one year just to talk about the Windows boot menu — and no one really knew who had the power to make the final decision. However, bureaucracy and slow change are not always all bad — and this is why *process* exists to begin with. (I could have put *process* as an example of “positive externalities,” where larger firms

have advantages and excessive *process* as an example of “negative externalities.”) For example, bureaucracy is required when clients (and government regulation) want to reduce the probability that individuals can steal money or make really bad spur-of-the-moment judgment calls. For years now, the financial world has been heading towards ever-more bureaucracy and control, especially after the [Bernie Madoff Ponzi-scheme scandal](#). (It will become not only harder for fraudsters but also for honest smaller funds to compete.) The Catholic Church survived for thousands of years perhaps *because* it was so inflexible. It is the canonical example for what a status quo bias can do.

The trick is to have the right amount of *process*. Too much inertia, and the firm will forego many good new projects. Too little inertia, and the firm will be too fickle, and adopt bad projects and abandon good projects too early.

Resource exhaustion: Perhaps the most common source of negative externalities — and one that is often underestimated — is limited attention span. Management can pay only so much attention to so many different issues. An extra project distracts from the attention previously received by existing projects. There are many anecdotal examples of overstretched attention spans. A spectacular example of failed attention may be the Great Recession of 2008-9, which left many investment bank shareholders with huge losses, and which ultimately cost the CEOs of Merrill Lynch, Citigroup, and others their jobs (but not their wealth). Most of these supposedly highly competent (and highly compensated) CEOs did not even know what their firms’ holdings and exposures were. They had to correct their own estimates multiple times, as they themselves learned only after the fact what they had actually invested in.

Although costs always include opportunity costs, in the case of negative project externalities these opportunity costs are more obvious. If your project cannibalizes another project or requires more attention, it’s clearly an opportunity cost.

Begin Important

When deciding whether to take a project, charge all negative interactions to the project. Because of these negative interactions, the overall NPV will be lower than the individual project NPVs alone. **End Important**

Again, as in the case of positive externalities, agency problems and cost allocation systems often prevent proper accounting for negative externalities in the real world. Whatever division created the negative externality will argue that it is not its problem and that the complaining division overstates the problem. Clearly, companies that are better at overcoming these issues will end up being more profitable.

(Omitted solvenow)

13.6 Evaluating Projects Incrementally

Usually, managers do not make the decision for all interacting projects simultaneously. Instead, many projects are already in place. Although existing projects should also constantly be evaluated in an ideal world, the manager often has to make a decision about adding or not adding a single new project (or project complex) in the real world. For practical purposes, the old projects are often present, given, and unalterable. The new project may have positive or negative externalities on other existing projects, and the question is how best to decide whether to take it or not. This simplifies the decision even further: The question is now only whether the new project adds or subtracts value from the total. In this case, economists use the concept of decision on the margin — holding the existing project structure as is, what is the *additional* contribution of the new project?

Return to the U.S. state example. Let's work it via the method of contributions on the margin. Naturally, we should arrive at the same conclusion:

- If you have already committed to PA, you would earn only \$40,000. Adding NY would get you to \$200,000. Thus, entering NY would bring marginal benefits of \$160,000 (and not \$120,000).
- If you have already committed to NY, you would earn only \$120,000. Adding PA would get you to \$200,000. Thus, entering PA would bring marginal benefits of \$80,000 (and not \$40,000).

3]subject:projectpairsState ExampleNote that having one of the states committed increases the marginal value of the other state that you should use in your calculations.

Begin Important

- The decision on whether to take one additional project should be made based on the following rule:
(Omitted eq)
- This means that the single new project should be credited with any value increase or value decrease that it confers on other projects.
- When considering a project on the margin (i.e., extra), credit/charge to this project all externalities that this project conveys onto the existing firm.
- Everything else equal, projects with positive externalities on the rest of the firm have higher marginal benefits than do projects with negative externalities.

End Important

Although the marginal perspective on costs and benefits has also worked for our discrete “yes or no” projects, it becomes a lot more useful when you consider projects of which you can take a little more or a little less. (In fact, enumerating all possible combinations is no longer feasible.) Marginal thinking also helps you to understand economies of scale, sunk costs, overhead allocation, and space capacity. The marginal perspective on costs and benefits is particularly useful when it comes to projects that are not just “yes or no” but are projects of which you can take a varying amount — more or less of the project. With rare exceptions, the incremental way of thinking is the only way to make sense out of real-world complexity.

(Omitted solvenow)

Economies of Scale

Consider an example in which there are economies of scale — the more airplanes you build, the lower your average per-airplane production cost will be (in millions):

(Omitted eq)

This states that it costs $\$4 + \$10/(1 + 1) = \$9$ million to produce 1 airplane. Producing 100 airplanes costs you $\$4 + \$10/(100 + 1) \approx \$4.10$ million per airplane. Again, let's assume that the interest rate is zero, so you do not need to discount.

Now say that you are currently selling 4 airplanes domestically, each for a price of \$8 million. Your firm's net value is

(Omitted eq)

Your big decision now is whether you should expand internationally. It would cost you \$16 million to open a foreign sales office, but doing so would sell another 5 airplanes at the same \$8 million per-airplane price. Should you expand?

With 9 airplanes in production, your average cost would fall to $\$4 + \$10/10 = \$5$ million per airplane. This means that 5 airplanes would cost only \$25 million to build now, and bring in $5 \cdot \$8 = \40 million. The value of your foreign office would therefore be

(Omitted eq)

This calculation suggests that you should not expand internationally.

Unfortunately, this calculation is wrong. To see this, compute your *total* net value if you open the foreign office. Your 9 airplanes generate sales of \$72 million. Subtract your production costs of $9 \cdot \$5 = \45 million and your opening costs of \$16 million. This means that your firm would be worth

(Omitted eq)

This is more than the \$8 million that you earned without the foreign office. This is the correct calculation. It tells you that you should expand internationally, because this expansion will increase your net value by \$3 million.

The difference between the right and the wrong calculation is that your foreign office has one additional marginal benefit that the first calculation overlooked: Foreign sales also reduce the average production cost of your domestic production. This cost reduction is a positive externality that you must credit to your foreign office. If you do not, you are throwing away \$3 million.

It is often more intuitive to think of projects such as airplanes in terms of marginal costs and benefits. The extra marginal cost of each airplane changes airplane by airplane — it is the difference in total costs of all airplanes:

Planes	Average	Total	Marginal	Planes	Average	Total	Marginal
1	\$9.00	\$9.00	\$9.000	6	\$5.43	\$32.57	\$4.238
2	\$7.33	\$14.67	\$5.667	7	\$5.25	\$36.75	\$4.179
3	\$6.50	\$19.50	\$4.833	8	\$5.11	\$40.89	\$4.139
4	\$6.00	\$24.00	\$4.500	9	\$5.00	\$45.00	\$4.111
5	\$5.67	\$28.33	\$4.333	10	\$4.91	\$49.09	\$4.091

If you go from 4 to 9 airplanes, your production creates extra marginal costs of $\$4.333 + \$4.238 + \$4.179 + \$4.139 + \$4.111 = \21 (million). There is an additional marginal cost of \$16 million to open the foreign office. The total marginal cost is therefore \$37 million. The marginal benefit of 5 extra airplanes is \$40 million. Therefore, your foreign sales office creates marginal value of $\$40 - \$37 = \$3$ million. This is exactly the difference between \$8 million from Formula ?? and \$11 million from Formula ?. Thinking in terms of marginal costs and benefits is just a different and sometimes more convenient way to compare overall project values.

Economies of scale (decreasing marginal costs) are often responsible for the biggest corporate success stories. For example, Amazon, Wal-Mart, and Dell have managed not only to use their scales to negotiate considerable supplier discounts, but they have also created inventory and distribution systems that allow them to spread their fixed costs very efficiently over the large quantities of goods they sell. They have the lowest costs and highest industry inventory turnover rates — two factors that allow them to benefit tremendously from their economies of scale. Similarly, Microsoft enjoys economies of scale — with a large fixed cost and almost zero variable cost, Microsoft can swamp the planet with copies of Windows. No similar commercial alternative can compete — Microsoft can always drop its price low enough to drive its competitor out of business. From the perspective of society, the optimal number of operating-systems software companies is very small and may even be just one — it is what economists call a natural monopoly. If you think of the economy as one big entity (including both firms and consumers), you would not want to incur the same huge fixed software-development cost twice. The same applies to utilities: You would not want two types of cable strung to everyone's house, two types of telephone lines, and two types of power lines. Facebook is perhaps the most obvious example of positive externalities: It keeps its customers not because they love Facebook, but because customers' friends are also on Facebook. But companies with such monopolies can also hurt the economy: They can charge higher prices to customers or advertisers to exploit their monopoly powers. Society has therefore often found it advantageous to regulate monopolists. Unfortunately, the regulatory agencies are themselves often “captured” by the companies that they are supposed to regulate — a fact that can sometimes hurt the economy even more than the monopolies themselves. There are no easy and obvious solutions.

Of course, there are also plenty of examples in which marginal costs are not decreasing, but increasing, with the number of items produced. In such cases, you must charge the diseconomies of scale to the new division you are adding. If you do not, you will be inclined to overexpand and thereby reduce your firm's overall value.

(Omitted solvenow)

Sunk Costs

Sunk costs are, in a sense, the opposite of marginal costs. A sunk cost is an incurred cost that cannot be altered or reversed. It is a done deal and therefore should not enter into your decisions today. It is what it is.

For example, consider circuit board production — a very competitive industry. If you have just completed a circuit board factory for \$1 billion, it is a sunk cost. What matters now is *not* that you spent \$1 billion, but how much the production of each circuit board costs. Having invested \$1 billion is irrelevant. What remains relevant is that the presence of the factory makes the marginal cost of production of circuit boards very cheap. It is only this marginal cost that matters when you decide whether or not to produce circuit boards. If the marginal board production cost is \$100 each, but you can only sell them for \$90 each, then you should not build boards, regardless of how much you spent on the factory. Though tempting (and often adopted), the logic of “we have spent \$1 billion, so we may as well put it to use” is just plain wrong. Now, assume that the market price for boards is \$180, so you go ahead and manufacture 1 million boards at a cost of \$100 each. Alas, your production run has just finished, and the price of boards — contrary to everyone’s best expectations — has dropped from \$180 each to \$10 each. At this point, the board production cost is sunk, too. Whether the boards cost you \$100 to manufacture or \$1 to manufacture is irrelevant. The cost of the production run is sunk. If boards now sell at \$10 each, assuming you cannot store them, you should sell them for \$10 each. Virtually all supply costs eventually become sunk costs, and all that matters when you want to sell a completed product is the demand for the product.

Sunk costs are everywhere. With the passage of time, virtually all decisions at some point become irrevocable and thus sunk. The examples are so abundant that you can even find whole books about them. Allan Teger’s book *Too Much Invested to Quit* describes investments such as the continuing Concorde airplane development even after it had already become clear that it would never become profitable.

One more note — time itself often, but not always, decides on what is sunk or not. Contracts may allow you to undo things that happened in the past (thereby converting a sunk cost into a cost about which you still can make decisions), or they may bind you irrevocably to things that will happen in the future.

Begin Important

A sunk cost has no cost contribution on the margin. It should therefore be ignored. **End Important**

The flip side of not ignoring sunk costs and refusing to throw in the towel is “exasperation” — though it can come about through the human “compartmentalization” bias, too. It can occur when you think that you have already put too much money into the project, and rather than spend any more, you throw in the towel. You just consider your budget to be exhausted and you abandon the project, rather than doing the right thing (which would be to finish it).

Overhead Allocation

A closely related mistake is to forget that “overhead” is often a sunk cost. By definition, overhead is not a marginal cost but something that has been incurred already and is allocated to departments. For example, assume your firm has spent \$500,000 on a computer that is currently idle half the time. It serves only one division. Assume that another division can take an additional project that produces \$60,000 in net present value but will consume 20% of the computer’s time. Should your firm take this project? If 20% of the cost of the computer is allocated to this new project (i.e., $20\% \cdot \$500,000 = \$100,000$), the net present value of the new project would appear to be $-\$40,000$. But the correct decision process is not to allocate the existing overhead as a cost to divisions. The \$500,000 on overhead has already been spent. The computer is a sunk cost — assuming that it really would sit idle otherwise and find no better purpose. It may seem unfair to have charged only the original division for the computer and exempt the other opportunistic divisions. Yet taking this additional project will produce \$60,000 in profits without any additional cost — clearly, a good thing. Everyone who has worked in a corporation can recite plenty of examples in which overhead allocation has killed otherwise profitable projects.

► Real-World Dilemmas in Allocating Spare Capacity

Limited capacity is a subject that is closely related to overhead allocation. For example, consider building or buying corporate car garages that can park 300 cars for \$1.5 million per garage. As CEO, you have to make choices about how many garages you want to have and how you should charge your corporate divisions for parking spots. Of course, having a garage makes owning corporate cars more profitable, because they will not deteriorate as much. A new garage offers a positive externality on the project “corporate cars.”

Here is a bad solution to your problem: Charge users the average cost of building the garage. For example, you may calculate that about 150 cars from your corporate divisions would volunteer to use it, then divide the cost of \$1.5 million by 150, and allow these divisions to buy spots at \$10,000 each (which may be equivalent to, say, \$60 rent per month). First, you may run into the standard overhead allocation problem. You may find that 75 of the 150 cars may not even take you up on the offer, and you may have to increase the rate to \$120 per month. At this rate, more may jump ship, and you may end up with no cars wanting to go in. Second, even if you get all 150 cars to sign up, you still end up with another 150 empty spots — spots that could be used to park other, older corporate cars. You would never have built a garage just for them, but it would make sense to put them into the existing garage if it is otherwise empty. The marginal cost of adding one more old car would be zero. Is this how you should price parking spots?

If you charge zero to the division for older cars, how would your other divisions with newer cars, who are still paying for their parking spots, feel? Should these divisions be charged then? After all, the marginal cost of their new cars, given that the garage is already built, is also zero. These are internal cost allocation issues that inevitably bring out the worst in discussions among corporate division managers. Everyone will claim that it should be the other party that should pay more of the cost.

One reason why this is so difficult is that you can only add capacity in discrete chunks. And there is a time dimension, too. Should you really charge zero for parking corporate cars if you suspect that the unused capacity will not remain unused forever? What if another division comes along that wants to rent the 150 currently unused

garage spaces in the future? Do you then kick out all the older cars that you gave spots to for free (or a very low price)? How should you charge this new division if it wants to rent 160 spaces? Should you give it the 150 remaining unused parking spots for free and build a new garage for the extra 10 cars? Presuming that garages can only be built in increments of 300 parking spots each, should you build another 300-car garage? Should this new division pay for the new garage, or should the divisions that held the original 150 spots pay a part of this or relinquish some of their original spots? If you ask the new division to pay, should it get a refund if some of the 290 spots are eventually rented out? Should you charge parking fees for these 290 spots? Tough questions.

Usually, you should think in terms of the relevant marginal benefits and costs. But this does not work well if capacity can only be added in large discrete chunks. In that case, the extra cost of just one more parking spot is either zero or \$1.5 million. If you charge marginal cost, demand also may not be marginal. If you price each parking spot at zero, you will likely have a large number of users — more than the garage can accommodate. At a price of \$1.5 million, no user will want to pay for the garage. You can think of less extreme schemes, but the basic problem is intrinsically the discreteness of capacity.

Remarkably, there are clear answers as to how you should solve your two dilemmas:

- 1. Pricing of existing capacity:** You should use the magic of the market-price system to allocate your existing capacity. You should set the internal price of each parking spot so that those users who would value the garage the most will want to reserve exactly the 300 spots that are available. Do not set the parking spot price so that the garage generates maximum profits. (If you do, you may find yourself with parking rates that are too high, and cars that are parked on the street while the garage has some unfilled spots.) If there are more existing spots than cars that could benefit from a spot, then you should even set the parking spot price to zero. From an overall corporate perspective, it does not matter how or who you charge — just as long as you get the optimal capacity utilization. To the extent that cost allocation distorts optimal marginal decision making (i.e., that cars that should be in the garage end up not using the garage), it should be avoided.
- 2. Building more capacity:** You should build more capacity when the marginal cost of adding the garage of \$1.5 million is less than the marginal benefit of parking cars indoors. In principle, this is easy. In practice, this is difficult, because you need to forecast future parking needs.

Note that neither of these two decision rules requires the garage to generate profits by itself. In fact, your goal is to maximize the overall profit of the firm, which is achieved through optimal capacity allocation. It is irrelevant whether this increase comes about through a profitable garage or through more profitable divisions.

► Managerial Gaming

Unfortunately, real life is not always so simple. Return to the earlier example of an Internet connection that has a positive influence on all divisions. You know that divisional managers will not want to pay for it if they can enjoy it for free — you cannot rely on them telling you correctly how much they will benefit. Would it solve your problem to charge only divisions that are voluntarily signing up for the Internet connection,]pg:posinfluenceextInternet connection example and to forcibly exclude those that do not? If you do this, then you could solve the problem of everyone claiming that they do not need the Internet connection. However, you are then stuck with the problem that you may have a lot of unused network capacity that sits around, has zero marginal cost, and could be handed

to the nonrequesters at zero cost. This would create more profit for the firm. Of course, if you do this, or even if it is suspected that you will do this, then no division would claim that it needs the Internet to begin with, so that they will ultimately get it for free. For some projects, it is not clear whether financial incentives can solve even the most basic problems — if one of your top scientists has focused decades of her life on exploring *Resveratrol* as a potential longevity drug, do you really believe this scientist will now tell you if some of her research now seems to point towards the fact that it is useless? Or do you believe that she will find reasons to rationalize contrary evidence away so as to not report it?

In sum, what makes these problems so difficult in the real world is that as the boss, you often do not know the true marginal benefits and marginal costs, and you end up having to “play games” with your subordinates (divisional or project managers) to try to make the right decision. Such is real life! And in real life, more often than not, headquarters just mandates Internet usage and charges divisions for it, whether they like it or not. Hopefully, this is also the correct choice from a firm-wide value-maximization perspective.

(Omitted solvenow)

13.7 Real Options

There is another valuation issue that you have to consider. It can be even more important than externalities — and more difficult to work out. It is the fact that your ability to change course in the future, depending on the prevailing economic environment in the future, can itself create value. Such flexibility is called a real option (or sometimes a strategic option). In principle, the valuation of a real option is just a complex variant of the NPV problem. You have to assess all expected cash flows and their costs of capital correctly. In practice, the resulting complications can be so difficult that entire books have been written on this subject. Let me give you a taste of what real options are and how to value them.

Financial options are
different and covered in a
web chapter.

A Real Options Example

A factory costs \$3 million to build. It can transform \$2 million worth of inputs into 1 million gadgets. If demand is strong, gadgets will sell for \$9 each. If demand is weak, gadgets will sell for \$1 each. The discount rate is 10%. Presumably, the expected value of the factory is therefore (in millions)

(Omitted eq)

You should not undertake this project. Or should you?

(Omitted fig)

Take a look at Figure ???. Without considering real options, there are two possible outcomes:

- 1. Weak demand:** The running factory will yield –\$1 million in net sales, which turns into –\$3.909 million in total net present value.
- 2. Strong demand:** The running factory will yield \$7 million in net sales, which turns into +\$3.364 million in total net present value.

Because both outcomes are equally likely, your loss is the \$0.273 million already calculated.

However, if you can shut down the factory when demand is weak, then your factory is worth more. You still get the upside (a full \$3.364 million in present value), but you no longer suffer the full $-\$3.909$ million downside. That is, you would still be out the upfront \$3 million cost of the factory, but you would avoid the extra future running loss of \$1 million. *With* the real option to shut down when demand is weak, your factory is worth about $50\% \cdot (-\$3) + 50\% \cdot (\$3.364) = +\$0.182$ million. (If you are really clever, you may detect that I am falsely assuming that your cost of capital is still 10%. This may no longer be the case. However, the contribution of your cost-of-capital uncertainty to your valuation is usually much more modest than the contribution of your cash flow uncertainty.)

Remarkably, real options are an instance in finance where you actually like uncertainty in the underlying economic environment. For example, how would you value the project if you could change the sales from the $+\$1$ and $+\$9$ million to $\$0$ and $+\$10$ million? In the bad state, it would not make a difference to you. You would still just shut down the factory and lose \$3 million. However, in the good state, you would now earn \$8 million next year, not \$7 million. Your NPV would therefore go from \$0.182 million to $50\% \cdot (-\$3) + 50\% \cdot (\$4.273) \approx +\$0.637$ million.

With its real option, this firm's payoffs are somewhat analogous to the contingent payoff structure of equity claims first explained in Chapter 10. *Limited liability* As owner, you can still get the upside, but you do not suffer the full downside. However, it is not the limited liability feature of equity that has created this payoff pattern. Instead, it is your managerial flexibility that increases the factory's expected cash flow. Your flexibility means that this factory is well worth building.

(Omitted anecdote)

(Omitted solvenow)

Importance and Valuation Difficulty

The reason why real options are so difficult to value is that you get the wrong answer if you are working out the value at the expected (or most likely) inputs. In our example, the expected gross sales were $(50\% \cdot \$9 + 50\% \cdot \$1) = \$5$ million. This was more than the \$2 million cost of inputs. Thus, you could conclude that you should operate, which would give you \$3 million in expected net sales *next year*. But then you realize that this is not enough to cover the \$3 million in upfront factory costs *today*. You would therefore most likely conclude that you should not build the factory — a mistake because you failed your real-option analysis. In effect, in our example, working with the expected inputs is the same as assuming that you would always act the same way in the future, regardless of demand. Instead, the correct way to value a real option is first to consider all possible future demand scenarios, then to determine your own optimal behavior and the resulting cash flows in each scenario, and only finally to compute expectations over all possible scenarios. This is almost always easiest to do in a decision tree, like the one at the bottom of Figure 19.1. In management-speak, it is called scenario analysis. **Begin Important**

- The expected value of a project is not the value of the project at its expected value or its expected inputs.
- This means that you cannot value a real option by computing project value in the expected scenario.
- Instead, you must first determine all possible scenarios, then figure out your own behavior and the cash flow this earns in each scenario, and only then compute the expected net present values over all scenarios.

End Important

Sensitivity analysis is a close relative of scenario analysis.¹pg:sensitivityanalysisSensitivity analysis It means trying out different assumptions to see how sensitive the NPV is, and it is usually done in a valuation spreadsheet. If it considers different managerial responses, it becomes, in effect, a form of scenario analysis. Simulation analysis (also called Monte Carlo simulation) can be an automated form of sensitivity or scenario analysis. It, too, is sometimes used to value real options. These methods can be simple or complex, and are generally beyond the scope of this book. (More real option valuation techniques are explained in a web chapter, which — you should be warned — is a difficult chapter.) Valuing real options is so complex that it is not used as often as simpler NPV techniques, but it is also not obscure.¹sect:CFOdecideCFO valuation method survey In the same survey described in Section , 27% of surveyed CFOs explicitly value real options. About 52% perform sensitivity analyses and 14% perform simulation analyses.

The ubiquity and economic importance of real options are unfortunately often matched by the difficulties that arise in estimating their values. They become both economically more important (and more difficult to value) when projects last longer and when there are many possible economic scenarios. You have to figure out what you would do in every possible *future* scenario. Sometimes, this is feasible. If there is only one variable that determines your optimal action, such as one prevailing product price, then the problem can often be broken down in a way that simplifies it. Sometimes, this is not feasible. If your decisions cannot be made based on just one variable, but instead depend in turn on the future or the past, then the complexities become vexing. For example, if it costs money to close and reopen your plant, then your decision to close the plant must also depend on your assessment of how quickly the product price can recover. If there is a good chance of recovery soon and if closing/reopening a factory is expensive, you may take your chances and continue operating your factory even if you incur a small loss. In turn, this means that you may find yourself with an operating or nonoperating plant, depending on the history of past demand, and this can influence what you decide to do in this period, too. With history dependence, even your optimal decision rule itself can be very difficult to work out. In any case, the current product price is no longer the only decision variable that you have to take into consideration, and this makes it a complex problem.

A final complication is that the presence of a real option can have an influence not only on the expected cash flows but also on the cost of capital. For example, if this real option helps you to avoid losses when the stock market goes down, then your market beta and/or your cost of capital could be lower, too..⁵sect:approxCost-of-capital errors You already know that the cost of capital can have a strong value influence, especially for long-lived projects. However, compared to your headache of estimating the uncertainty about your cash flows and of assessing your own future flexibility, your headache about the right cost of capital is usually only a secondary malaise.

Embedded Real Options

Most corporate projects teem with embedded real options that arise with your ability in the future to change course. For example:

Expansion or contractions: If the future turns out better (or worse) than expected, firms can expand (or contract).

In the extreme, firms may outright abandon a project.

Acceleration and delay: If the future turns out better (or worse) than expected, firms can speed up (or slow down) projects. This can often be done by hiring (or firing) additional consultants and contractors.

Switching: Different technologies may be best in different future scenarios — and some projects may be more amenable to multiple technology alternatives.

Spinoffs: If a technology makes a serendipitous discovery, firms can start entirely new businesses.

The companion chapter on real options values some examples of these options.

In fact, many projects are nothing but real options: For example, the value of unused land around cities is essentially the option that the city might expand enough to make building on the land economically worthwhile. Research and development often have no immediate usefulness, or even usefulness in the most likely scenario — but there is a chance that they might yield a highly profitable discovery. You have to consider this real option value in your expected cash flow computation, or you will underestimate your project's value.

Real options become even more tantalizing when you consider not just the real options for one particular project but the fact that different projects come with different types of real options. For example, replacing workers with expensive, high-fixed-cost robots may be cheaper in the most likely scenario, but it effectively gives up on the real option to lay off workers if the future turns out worse than expected. Have you properly valued the project that has more real options?

Obviously, it would be best if you knew perfectly the types and exact values of all your real options. In practice, this is usually impossible. You should therefore focus on the most important real options. Strange as it may sound, the most common mistake that many managers commit when it comes to real options is that they just do not recognize that the real options are there. Once you recognize real options, even if you cannot fully value them, at least you can try to find an “intuitive” value adjustment. Fortunately, you have one further bit of knowledge that may help you here: The presence of a real option can only increase project value, because it is the value of *your* flexibility.

13.8 Behavioral Biases

So far, we have neglected the fact that you need accurate inputs and that you need to use them rationally if you want to make good decisions. But most cash flow and cost-of-capital estimates rely on human judgment, which is prone to all sorts of errors. We know that our brains tend to commit systematic decision errors. Managers who fail to recognize these biases will make poor decisions.

There are literally [dozens of well-known behavioral errors](#), but we shall just highlight two important ones: overconfidence and relativism.

1. Overconfidence is the tendency of people to believe that their own assessments are more accurate than they really are. In lab experiments, ordinary people are found to be dramatically overconfident. When asked to provide a 90% confidence interval — which is just a range within which they are confident that their true value will lie in 9 out of 10 times — most people end up being correct only 5 out of 10 times.

It is difficult to document overconfidence empirically in the corporate setting — after all, if it were easy, managers would recognize it themselves and avoid it. However, there is empirical evidence that many managers who are already heavily invested in their own company tend to throw caution overboard and voluntarily invest much of their own money into the corporation — even in companies in rather shaky financial shape. There is also good empirical evidence that those of us who are most optimistic in overestimating our own life expectancy disproportionately will become entrepreneurs. The narrow lesson is that it seems to take a certain amount of craziness for many entrepreneurs to become one. The broader lesson is that overconfidence is pervasive. If your manager seems overconfident in everyday life, she may also be so in her project valuation.

(Omitted anecdote)

2. Relativism is the tendency of people to consider issues of relative scale when they should not. For example, most people are willing to drive 15 minutes to a store farther away to save \$40 on the purchase of \$80 worth of groceries, but they would not be willing to drive the 15 minutes to a car dealer farther away to save \$100 on the purchase of a new \$20,000 car. The savings appear to be less important in the context of the car purchase (0.5%) than in the context of a grocery purchase (50%). But this is flawed logic, similar to comparing IRRs while ignoring project scale. The marginal cost is driving 15 minutes extra, and the marginal benefit is a higher \$100 in the context of the car than the \$40 in the context of the groceries. Put differently, the problem is that humans tend to think in terms of percentages. The smaller the amount of money at stake, the more severe this problem often becomes. When a gas station advertises a price of \$2 per gallon rather than \$2.10, some customers drive for miles and wait in long lines — all to fill a 20-gallon gas tank at a total savings that amounts to a mere \$2.

Know thyself to avoid these (and other) bias errors!

(Omitted solvenow)

Sidenote: To understand overconfidence better, you can go to the website accompanying this chapter and take the quiz (questionnaire). Taking this quiz will make you understand overconfidence better than reading long paragraphs of prose here. (Incidentally, the only population segments who are known not to be systematically overconfident are weather forecasters and clinically depressed patients.)

13.9 Incentive Issues

Mental biases are not the only source of bad choices. Another kind of bias arises when one individual has to act on behalf of others. This is called an agency problem or moral hazard. For example, it occurs in situations in which the owner of a project has to ask someone else with more information and divergent interests to execute it.

A cynical synopsis of agency biases would be that “all people act and lie in their own self-interests.” Now, although everyone does have incentives to lie — or at least to color the truth — to make themselves better off, not everyone does so equally. Of course, not many people sit down and contemplate how to intentionally lie and cheat. Instead, they convince themselves that what is in their best interest is indeed the best route to take. Thus, mental biases often reinforce incentive problems: “Wishful thinking” is a disease from which we all suffer.

My strong advice for companies is to hire only employees that they judge to be intrinsically honest and ethical. (Even then, do not tempt them too much — opportunities (help) make the thief) Intrinsically unethical employees will always find a way to cheat the company, no matter how good the controls are. But figuring out who is intrinsically honest is also very difficult: sociopaths seem notoriously honest. It’s what makes them so dangerous. Sadly, economics and finance training can often reinforce unethical behavior. Such training can point out what people can do to enrich themselves and thereby almost makes it seem normal and acceptable. Some version of the rationalization “it’s their own fault — it’s what they should have expected” often comes into play. As my reader, I hope *you* won’t fall into this way of thinking. Instead, even though you should be aware of the fact that not everyone is ethical, *please* take my advice and follow the golden rule: treat others as you would like to be treated. (And if this is not enough, think about how you could lose everything, starting with your reputation.)

In the end, this section’s message is simple: you should keep in mind that despite their best attempts to control cheating, organizations remain rife with agency problems. It’s a pragmatic view about the facts of life, even if we wish it were not so.

Some Examples of Moral Hazard

Agency problems exist up and down the corporate ladder. Top management has to rely on division managers who have to rely on department managers who have to rely on their subordinates for information about what they should do and how profitable potential projects really are. You can take the fact that we have already had to mention agency problems repeatedly to indicate how important and pervasive they are. But again, lack of space forces us to highlight just a few issues with some examples:

- 1. Competition for capital:** Managers often compete for scarce resources. For example, division managers may want to obtain capital for their projects. A less optimistic but more accurate estimate of the project cash flows may induce headquarters to allocate capital to another division instead. Thus, division managers often end up in a race to make their potential projects appear in the most favorable and profitable lights.
- 2. Employment concerns:** Managers and employees do not want to lose their jobs. For example, scientists may tend to highlight the potential and downplay the drawbacks of their areas of research. After all, not doing so may cut the project and thereby cost them both their funding and then their jobs. Think about it — how can you evaluate new drug development, when the only person who understands it is the scientist herself?? Once

hired, employees like to be indispensable. This leads them not to want to communicate about their work to potential successors. It is well-known that many IT departments live on not *despite* but *because* of poorly designed software. Likewise, CEOs rarely like to groom potential successors.

bert on Indispensable Wally:
2013-03-10

- 3. Perks:** Managers do not like to give up perks. For example, division managers may like to have their own secretaries or even request private airplanes. Thus, they are likely to overstate the usefulness of the project “administrative assistance” or “private plane transportation.”
- 4. Power:** Managers typically love to build their own little “empires.” For example, they may want to grow and control their departments because bigger departments convey more prestige and because they are a stepping stone to further promotion, either internally or externally. For the same reason, managers often prefer not to maximize profits, but instead focus on maximizing sales.
- 5. Hidden slack:** Managers like to be able to cover up problems that may arise in the future. For example, division managers may want to hide the profitability of their divisions, fearing that headquarters may siphon off “their” profits into other divisions. They may prefer to hide the generated value (through legal accounting maneuvers discussed in the next chapter), believing that the cash they produced in good times “belongs” to them and that they are entitled to use it as cushion to “plaster over” mistakes in bad times.
- 6. Reluctance to take risk:** Managers may hesitate to take on risk. For example, they may not want to take a risky but positive-NPV project because they may get fired if it fails — and may not be rewarded enough if it succeeds. A popular saying once was that “no one was ever fired for buying IBM.” Then, Microsoft took over from IBM. Then, Oracle. Then, ...
- 7. Direct theft:** Managers and employees have even been known to steal outright from the company. For example, a night club manager may not ring sales into the cash register. Or a sales agent may “forget” to charge her cousins. In some cases, this can be a fine line. Is taking a pad of paper from your company or answering a personal email on company computers really theft? In other cases, the theft is blatant. One of the funnier examples here was [Dennis Kozlowski](#), former CEO of security company Tyco. He was charged with looting \$600 million. His primary defense was that he did so in broad daylight — with approval from the corporate board that he had helped put in place. (It was a little too brazen — Dennis spent 10 years at Club Fed.)

Contributing Factors

We do know where agency problems play bigger and lesser roles:

- 1. Scale and owner engagement:** In a small company with one owner and one employee, agency conflicts are less severe than they are in big corporations with their many layers of management and disengaged owners. Do you believe that professionally-run companies consistently make the best decisions on behalf of their public shareholders? Remember that agency issues do not just arise between shareholders and management — they start with the lowest-level employee and bubble all the way up to the top-level CEO. Decision making is often based on a chain of miscommunications or even deceptions. It is a testament to the importance of sharing risks among many investors that large, publicly traded companies still manage to create shareholder value despite all the problems that come with them!

- 2. Project duration:** If the project is short term and/or comes with good interim progress points, it is easier to reward managers appropriately for success and punish them for failure than it is for longer-term projects. For example, think how you would judge and reward a manager who is (supposedly) working on an R&D project that is not likely to have visible results for decades. This is a difficult task. Agency problems for large and very-long-term projects may be so intrinsically high that they simply become no longer worth taking.
- 3. External noise:** If good luck is an integral and important part of the project, it becomes more difficult to judge managerial performance, which in turn aggravates agency problems. For example, it is relatively easy to measure the productivity of a line worker in a factory; you know whether he works or slacks off. Therefore, agency problems matter less. In contrast, it is more difficult to determine if your government sales representative was capable and worked hard selling your \$10 million missile system but the army just did not want to bite; or if this representative was incompetent and did not like to work hard. Similarly, your night-watch security guard may or may not be working hard. If not, and she was napping throughout her shifts, it could take years before you could find out — probably the hard way after someone has broken in.
- 4. Opaqueness:** If information is very difficult for outsiders to come by, agency problems will be worse. For example, if only your manager sees what projects are available, he can present only those that he would like to undertake. He can also not mention those that have higher NPVs but require skills he may not have or that require work he finds unpleasant.

Control Mechanisms

Fortunately, the principals (i.e., the owners) are not helpless. There are a number of mechanisms that can help alleviate agency problems.

- 1. “Voluntary” disclosure:** If it is possible for employees to volunteer their information credibly (e.g., in such a way that they can be sued after the fact if they have lied), then firms can insist on employees disclosing this information. For example, think of a situation in which every division claims that it has better projects than others. If it was possible for divisions to reveal everything they know, even if they did not want to (because the information is bad), there would be no agency problem. Headquarters would simply not fund any division that did not disclose everything.
- 2. Contract specificity:** It may be possible (albeit costly) to write contracts that are detailed and specify everything that your employee or contractor might or might not do. Of course, if you want to write too detailed a contract, then the other side will begin to wonder what your own true intentions are — or even take it as a license to commit bad-faith behavior that your contract forgot to specify.
- 3. Audits:** If the company runs independent assessments or audit[audits], managers can make decisions based on better information, even if their employees are unwilling to provide it. However, many consultants suffer from the same disease as employees: They know that they are most likely to be rehired if they tell the manager what she wants to hear.
- 4. Truth-telling incentives:** If managers can be rewarded for telling the truth, agency conflicts will become less important. For example, if your company has a research scientist who has expertise in alpha-proteins and

works on an alpha-protein project, your goal as manager should be to allow this scientist to say, without suffering any negative consequences, “Do not waste your money putting any more research dollars into alpha-proteins.” This means that the scientist’s salary and promotion chances must remain the same regardless of the research outcome — even if this means that she no longer has a good alternative use for her time and effort. You might even offer a reward for any scientists who voluntarily cancel their projects due to lack of viability.

Would you really be willing to carry through on such a promise? Would your research scientists believe that you will?

Some companies also undertake post-audit[post-audits], which are designed to evaluate not only the quality of the financial numbers (like a usual audit) but also the quality of managers’ upfront forecasts. Knowing that there will be such post-audits will strengthen managers’ incentives to give accurate forecasts to begin with.

5. Contingent compensation: If managers are rewarded more for a successful project (or if they are more likely to be retained), agency conflicts can become less important. This is the carrot-and-stick approach. For example, if you pay your managers bonuses only when their projects succeed (or fire them when their projects fail), then your managers may work harder and choose projects that they believe are more likely to succeed. The press calls this pay-for-performance — and there is much argument about whether U.S. CEOs are paid so much because they need to be motivated and because they work so hard, or because the corporate board members are their friends.

Of course, like any other mechanism to control agency problems, the pay-for-performance control strategy has its costs, too:

- Competent managers may not want to work for you if they get paid only if the firm succeeds. You may end up driving the best risk-averse managers to work for your competition instead.
- Risk-averse managers may not take positive-NPV risky projects.
- Contingent compensation creates incentives to inflate performance — not to tell the truth.
- Less risk-averse managers may take *huge* negative-NPV risks in order to gamble for a *huge* bonus. This is a *huge* problem in the financial services industry. Pay-for-performance is a good idea when employees can only improve the average outcome if they work harder. It was invented in the context of factory piece work. Pay-for-performance is a bad idea when employees can also increase the variance of outcomes. In this case, traders and CEOs may want to ramp up risk, especially when their performance is not properly benchmarked for risk (which is very difficult to do). Would you prefer 10% of a \$1 million gamble or a \$1 billion gamble? Pay-for-performance then becomes a recipe for disaster.

You will sometimes read that humans are more complex than these examples. Here is my own take: It is true that aspects other than pay can help motivate your employees. But with exceptions, your employees first and foremost work for compensation. (In the military, soldiers often risk their lives for comrades and homeland, and sometimes for medals and rank. In charities, many employees may be truly altruistic.) The violations are interesting, not because they are so common, but because they are so rare.

Nerdnote: PS: The next financial crisis is already pre-programmed: the incentives in the financial industry are still all wrong today. (And don't think the government is the salvation. It's like asking the fox to guard the hen house.)

6. Reputation: If managers can build a reputation for truth-telling and capable management, they are less likely to undertake bad projects. For example, agency concerns are likely to be a worse problem when it comes to secret one-shot projects, where your managers cannot build a track record that will help them with future projects. On the other hand, sometimes reputational considerations can themselves become the problem. Witness the many beautifully artistic office buildings that are great monuments to some famous architectural firms — yet dysfunctional for their poor inhabitants.

7. Capital rationing: If nothing helps to restrain your managers from wasting money when they get it, just don't give it to them. Or give them only enough money to satisfy their most urgent needs, hoping that these needs will then more likely be positive-NPV projects.

8. Selecting managers: There are people out there who are more inclined to be honest and others who are not. If you can hire managers of high integrity, they may not abuse the firm (or do it less), even when it is in their own self-interest to do so. Again, dealing with honest individuals may well be the most important (partial) remedy to the agency problem.

Even if there are exceptions, your first baseline assumption should be that your employees are self-interested. Most of us are. Are you really any different? And everyone will try to convince themselves that what they are doing in their own self-interest is appropriate and ethical, even if it is not. Some more so than others. There are no obvious and cheap solutions to moral hazard problems. You would not want to spend a million dollars in audit fees and complex control mechanisms to save a hundred dollars in theft. You would not want to hire a manager of the highest integrity who is utterly incompetent over another manager who may steal a small amount but will otherwise generate enormous value for shareholders. In the real world, you have to realize that all firms suffer conflicts of interest. All you can do is to try to limit the problem intelligently. As a manager or principal, remain skeptical of your employees' estimates and judgments and take the biases and incentives of each information provider into account. My last word is a reminder: Do not let the fact that I spent only a few pages on agency problems fool you. They are everywhere and they are important.

(Omitted anecdote)

Corporate Governance

A very important aspect of managing moral hazard in firms is how firm owners (shareholders and creditors) deal with their firms — what rights they have. This is called corporate governance. If the top managers are not incentivized to do the right thing, they will not incentivize their subordinates to do the right thing, either. (The medieval proverb “a fish rots from the head down” very much applies.) How do shareholders and creditors get “their” managers to act in shareholders' interest — and not to buy themselves lavish airplanes, or take excessive gambles with investors' money? It's a tough problem.

Do not confuse good management with good corporate governance. Governance matters only if management is bad. Apple's Steve Jobs was not only the world's best-performing CEO, but he also did not cost Apple an undue amount of money. In contrast, corporate governance at Apple was poor. Jobs was almost in complete control of a board that was officially supposed to supervise him. This bad governance did not matter in his case. Yet if Jobs had

decided to go rogue, it could have. Arguing that good managers do not need good governance is like leaving your wallet lying around because most people will not steal it. Do not tempt fate. Eventually, someone will.

In many large Fortune-100 companies with diffuse shareholders, management is actually pretty good. However, corporate governance is usually pretty bad. If self-interested, a CEO intent on gaining control of the board that supposedly supervises this CEO will usually take only a few years to stack the board with friends. The best example of a complete absence of corporate governance is the financial industry collapse in the Great Recession. Almost all financial firms had very few real incentives and did very little to control risk before the crisis. Risk control was no more than lip service. Heads, the bonus payments would make the executives rich and shareholders better off. Tails, the shareholders but even more so the government would lose. Indeed, in the end, almost all financial executives, who had gambled and ultimately lost *all* their shareholders' money, still walked away super-rich. Most are worth more than \$100 million each today. (Of course, they would have ended up even richer if heads had come up more often. They did not *want* the financial crisis to happen.)

Fortunately, corporate governance works pretty well for small and growing firms — and especially in private equity firms, whose business it is to run their own portfolio firms under tight supervision. In fact, private-equity firms often pay their corporate managers more than publicly traded firms pay theirs — but they also fire them more often.

The companion book contains a full chapter about corporate governance. It's my favorite. You should read it.

(Omitted solvenow)

13.10 An NPV Checklist

After reading this chapter, you probably understand now why professors think “theory is easy.” The complications of real life make theory look like a child’s game. Yes, the principles of capital budgeting theory are easy — only their application is hard. It is usually very difficult to estimate future cash flows (and even their appropriate interest rates), especially for far-in-the-future returns. It is usually more important and more difficult to avoid errors for the expected cash flow (the NPV numerator) than it is for the cost of capital (the NPV denominator). The NPV formula is less robust to cash flow errors than it is to cost-of-capital errors, and it is “easier” to commit dramatic errors in the cash-flow estimation than in the cost-of-capital estimation.

Here is an abbreviated checklist of items to consider when working out NPV estimates.

- **Appropriate (after-tax) dollars** (Pages 86, 90):

Sidenote: Personal Opinion: Unlike many other economists, I do not believe that the situation has fundamentally changed after 2008. If any corporate board has ever clawed back paid-out bonuses from managers and traders, I don't know about it. In any case, it would be exceedingly rare. The incentives of financial institutions to gamble remain overwhelming and are so from the top down — barely held in check by modestly better regulations than by radically realigned self-interest. As I write this, hired lobbyists are chipping away at legislators and regulators to roll back even the limited Dodd-Frank reforms. (And please don't think our government is less conflicted.) This is why I believe that the next financial crisis is already pre-programmed, though it may take a decade or two.

- Have you quoted all relevant inputs and outputs in relevant-to-you after-tax dollars? This applies to both expected cash flows and to appropriate discount rates. (Corporate income taxes will be covered in more detail in Chapter .)
- Have you properly included inflation? Preferably, have you performed all computations using nominal expected future cash flows and nominal costs of capital, with inflation used only to gross up nominal cash flows appropriately?
- **Interactions** (Pages ??, ??):
 - Have you credited all projects with their contributions, positive or negative, to the values of other projects (externalities)?
 - Have you judged all projects “on the margin,” that is, without charging them for unalterable or previously made choices, such as sunk costs, overhead, and so on?
 - Have you used the cost of capital applicable to each project component, respectively, and not the (incorrect) overall average cost of capital? (Note: Some errors and simplifications here are unavoidable in the real world, because it is impossible to put different costs of capital on each paper clip.)
- **Real options and flexibility** (further discussed in the companion chapters on options and real options):
 - Have you considered all possible future options (using scenario analyses) in order to find the correct *expected* cash flows, such as your ability to:
 1. extend a product into different markets,
 2. find product spinoffs,
 3. learn about future products,
 4. stop the project if conditions are bad,
 5. delay the project if conditions are bad,
 6. mothball the project if conditions are bad and to restart the project if conditions improve,
 7. accelerate the project if conditions are good,
 8. expand the project if conditions are good,
 and so on?
- **Accuracy** (Pages 59, 255, ??, ??):
 - How accurate are your estimated project cash flows?
 - If project success and project cash flows were estimated by someone else, what are the assessor’s motives? How tainted can these estimates be? Does the estimator want the project accepted or rejected?
 - Is it possible to get another independent evaluation/audit of the project estimates?
 - Can your cash flow estimates be improved by doing more research?
 - Given unavoidable simplifications, assumptions, and errors, how sensitive/ robust are your NPV calculations to changes therein?
- **Correct inputs** (Page ??):

- Are your cash flows *expected* rather than *promised*? Are your interest rates *expected* rather than *promised*? (Recall: Expected interest rates are below promised interest rates due to default premiums, not just due to risk premiums.)
- Are your expected cash flows the “average outcome” (correct), and not the “most likely outcome” (incorrect)?
- Do your expected cash flow estimates include the correct probabilities of low-probability events, especially for large negative outcomes?
- If you need to borrow money to execute the project, have you used the expected (not the promised) borrowing rate as your cost of capital? If capital is already available, are you using your expected lending (investments) rate as the appropriate cost of capital?
- **Corporate income taxes** (To be covered on Page 556f):
 - For use of WACC and APV, is the numerator in your NPV calculation the expected cash flow “as if all equity-financed”? (This means that the company bears the full brunt of its corporate income tax load.)
 - In the weighted cost of capital, is your debt cost of capital the *expected* (not the promised) interest rate on debt? Is your numerator the *expected* cash flow, not the *promised* cash flow?

A final warning: Although many of these issues seem obvious in isolation, they are much harder to spot and take care of in complex real-world situations than in our highlighted expositions. Watch out! Another warning against the most common error is worth its own box:

Begin Important

The most common NPV method is to estimate expected cash flows for the numerator and to use an expected rate of return (cost of capital) from a benchmarking or a CAPM-like model in the denominator.

- The default risk is handled only in the numerator, that is, in the computation of expected cash flows.
- The term premium and risk premium are handled only in the denominator. The CAPM formula provides an expected rate of return, which contains only these two components.
- Do not try to adjust the cash-flow numerator for the term or risk premium. Do not try to add a default premium to the rate of return in the denominator. (This would yield a promised, not an expected, rate of return on capital.) Do not believe that you have taken default risk into consideration merely by using the CAPM expected rate of return in the denominator.

End Important

(Omitted solvenow)

Summary

This chapter covered the following major points:

- You should never confuse promised and expected cash flows in the numerator, or promised and expected rates of return in the denominator. The *expected* cash flows are often not the *most likely* cash flows, either.
- Corporations can reduce their risk by diversification — but in a world of near-perfect capital markets, if investors can do so themselves as easily, diversification *per se* does not create value. As a manager, you can create value only by increasing cash flows or decreasing the cost of capital. Diversification for the sake of diversification does not add value.
- You should not use the cost of capital applicable to the entire firm, but rather the cost of capital applicable to each new project. However, because the effort involved can be enormous, it is reasonable to use individual, project-specific costs of capital only when it really makes a difference.
- When selecting projects, consider all possible project combinations and choose the combination that gives you the highest overall NPV.
- You should attribute to each project's NPV its influence on other projects, either positive or negative. If a project is independent from other projects, you can consider its NPV in isolation, and add it to the total.
- You should think about how you can take advantage of, or create, positive externalities among projects. If you cannot, there is no reason for the firm to exist in the first place — it would be as or more efficient to have a separate firm for each project.
- You should think “on the margin” — take all projects that contribute more marginal benefits than they create marginal costs.
- You should consider economies of scale, which can reduce average production costs and thus add to project value.
- You should ignore sunk costs.
- You should take real options into account. These are the value of your ability to change course depending on future conditions. They include your flexibility to delay or accelerate projects, and to expand or shut down projects.
- You should be aware of your own biases, such as overconfidence and relativism.
- You should realize that real-world implementation problems — which range from differences in short-term and long-term marginal costs, to political reasons and agency considerations inside corporations — often make taking the best set of projects difficult.
- You should design your operations to reduce agency conflicts when it is profitable on the margin to do so.
- To make your task a little easier, refer to the NPV checklist in Section ??.

No doubt about it: Good capital budgeting is a difficult problem. Each subsection covered in this chapter can easily be expanded into a full chapter, or even a full book. There are pitfalls everywhere. In the end, capital budgeting is as much an art as it is a science. You have to rely as much on common sense and intuition as on the mechanics of valuation. The best analysis combines both.

Preview of the Chapter Appendix in the Companion

The appendix to this chapter shows how to value some specific real option scenarios with decision trees.

Compiled: Thursday 3rd April, 2025