11

Market Imperfections

Information/Opinions, Market Depth, Transaction Costs, and Taxes

So far, we have assumed no differences in opinions (and thus information), no transaction costs, no taxes, and a large market with many competitive sellers and buyers — a "perfect market." We discussed uncertainty, risk, and the CAPM (like most finance formulas in the real world) in this framework. They are not just ancient and obsolete theory, but they are also not the end-all. If the assumptions do not hold, then these very same formulas, used by practitioners and academics alike, might be simply inapplicable.

Why are the perfect markets assumptions so important? You will learn that it is because they give us one unique, appropriate, expected rate of return — whether you want to borrow someone else's money to finance your projects or lend your money to someone else undertaking projects. Breaking the assumptions causes havoc: Without a unique expected rate of return, the project value depends on the (cash position of the) owner. What does "project value" even mean without a unique price?

Of course, as wonderful as perfect markets are, they do not exist. They are conceptual, not real. For large publicly traded firms, some financial markets can come very close to perfection. For small firms, they almost never do. Entrepreneurial finance is really just one example of "financing in imperfect markets."

So, in this chapter, you are leaving our beautiful, frictionless, utopian world. You will have to contemplate how to think about financial questions in the real world. Fortunately, many of your tools (and specifically NPV) still work — remember, for a tool to work in a more complex scenario, it is a minimum sanity condition that it also work in a simpler scenario. The trick of this chapter, then, is to learn how you apply your tools with more caution and to appreciate their limitations.

rtups:
06-01

11.1 Causes and Consequences of Imperfect Markets

So far, we have not distinguished between the cost of capital at which you can borrow money to finance your projects and the rate of return at which you can save money. In "perfect markets," these two rates are the same. Again, this is the purpose of all four perfect markets assumptions. It is only to guarantee one fact on which everything else rests:

> *Perfect markets cause equal borrowing and lending rates.*

When this is not the case, the implications are far-reaching. If these rates are not equal, then you cannot move in and out of an investment as often as you like because you lose on each round-trip transaction. More fundamentally, even the value of a project stops being unique. Instead, a project may be worth any number in a whole range of possible values. Indeed, the whole concept of one project value may become meaningless. Value can depend on who owns the project, what the tastes of the individuals' relatives are, or even what time of day it is. You could not even claim that the value of a project is its PV. Present value may itself be meaningless. But let's take this one step at a time.

(Omitted solvenow)

Judging Market Perfection for Intel Shares and Houses

Start by contemplating the four perfect markets assumptions for a stock like Intel:

1. **No differences in opinion:** Recall that this assumption does not mean that there is no uncertainty, but that investors do not disagree about the uncertainty. Objective, rational traders with access to the same kind of information should come to similar conclusions about Intel's value. They should agree on the distribution of prices that Intel shares will likely sell at tomorrow, which in turn defines share value today. For the most part, it is unlikely that rational traders would disagree much about the value of Intel shares — they should realize that it is not very likely that they can predict the price of Intel much better than the market. Any disagreements would likely be minor. Of course, if some traders have insider information, then they could predict tomorrow's price better, and the perfect market would be no more — but trading on inside information is illegal.

2. **Infinitely many investors and firms:** On a typical trading day in 2020, the median trading volume for Intel (INTC) was 27.5 million in shares, many millions of transactions, and $1.5 billion in dollars. This is a lot of buyers and sellers. Thus, Intel shares appear to trade in quite a competitive market, in which no single buyer or seller influences the price. There are lots of potential buyers willing to purchase the shares for the same price (or maybe just a tiny bit less), and lots of potential sellers willing to sell the shares for the same price (or maybe just a tiny bit more). Similarly, there are many thousands of firms offering shares, although they are admittedly not all perfect substitutes on all dimensions.

3. **No transaction costs:** Trading Intel shares does incur transaction costs, but these are modest. A typical total round-trip transaction cost spread for Intel is about 5 cents on a $50 share price, which is 10 basis points. An institutional trader may even be able to beat this. There are no searching costs for finding out the proper price of Intel shares (it is posted online), and there are very low costs to locating a buyer or seller — a click away on your retail broker's website.

4. **No taxes:** This may be the most problematic perfect market assumption in this context. Fortunately, we need this assumption of no taxes primarily for one purpose: The return to a seller owning Intel shares should not be different from the rate of return to a buyer. Here is what I mean.

   Consider an extreme example in which Intel starts out at $20 per share and happens to end up at $80 per share two years later. Assume that the capital gains tax rate is 20% and the risk-free discount rate is 5%. How much value is saved if you hold shares for two years versus if you sell them to me at the midway price of $50? If you keep the shares, the taxable capital gains would be on $80 – $20 = $60. At a 20% capital gains tax rate, Uncle Sam would collect $12. If you instead trade them to me at $50 after the first year, the capital gains

> consequences would be on \$30 first for you ($20\% \cdot \$30 = \$6$), and then on \$30 at the end for me (\$6 again). This violates the perfect market assumption, because if you hold the shares for two years, the present value of the tax obligation is $\$12/1.05^2 \approx \$10.88$. If you sell them to me, it is $\$6/1.05 + \$6/1.05^2 \approx \$11.16$. Thus, shares are worth more to you (the seller) if you hold onto them than if you trade them to me (the buyer).
>
> But the difference in how we value shares is really only in regard to the interest on the interim taxation. It is only 28 cents on a gain of \$60. Moreover, this example is extreme not only in the 300% rate of return, but also in assuming a worst-case taxation scenario. This chapter later explains that many capital gains can be offset by capital losses and that investor tax-timing discretion can further lower taxes. Furthermore, most shares are now held by institutions. Many of these are pension funds, which are entirely tax-exempt and therefore face no tax implications when trading.

The market for Intel shares may indeed be close enough to being perfect to allow you to use perfect markets as a first working assumption.

Unfortunately, not everything is traded in a perfect market. For example, think about selling your house. What is its value? What if your house is in a very remote part of the country, if potential buyers are sporadic, if alternative houses with the same characteristics are rare, or if the government imposes much higher property taxes on new owners (as, e.g., California does)? Intuitively, the value of your house could now depend on the luck of the draw (how many potential buyers are in the vicinity and see the ad, whether a potential buyer wants to live in exactly this kind of house, and so on); your urgency to sell (depending perhaps on whether you have the luxury to turn down a lowball first offer); or whether you need to sell at all (as current owner, you enjoy much lower property taxes, so your house may be worth a lot more to you than to a potential buyer). The value of such a house can be difficult to determine because the market can be far from perfect — and the house value may not even be one unique number.

The range in which possible values lie depends on the degree to which you believe the market is not perfect. For example, if you know that taxes or transaction costs can represent at most 2-3% of the project value, then you know that even if value is not absolutely unique, it is pretty close to unique — possible values sit in a fairly tight range. On the other hand, if you believe that there are few potential buyers for your house, but that some of these potential buyers would purchase the house at much higher prices than others, then it depends on your financial situation as to whether you should accept or decline another buyer's lowball offer.

Not all financial markets are close to perfect either. Information differences, the unique power of large buyers or large sellers in the market, transaction costs, or special taxes can sometimes play a role. For example, many corporate bonds are traded primarily over-the-counter. Just a small number of financial traders may make a market in them. If you want to buy or sell such a corporate bond, you must call a designated in-house desk trader. These traders are often your only market venue, and they will try to gauge your expertise when negotiating a price with you. You could easily end up paying a lot more for a bond than what you could then sell it back to them for just one minute later.

To repeat — no market, financial or otherwise — is ever "perfectly perfect." However, for some financial instruments, it is very close.

Begin Important

For many financial securities — for example, for large, publicly traded stocks — the assumption that the market is perfect is reasonable. For other financial securities and many nonfinancial goods, this assumption is less accurate.

End Important

(Omitted solvenow)

Perfect Market Assumptions and Violations

Now think more rigorously about what happens when each of the perfect market assumptions is violated:

1. No differences in opinion (information): This assumption means that everyone interprets all uncertainty in the same way in a perfect market. How could this assumption be violated? Here is an example. If your bank believes that there is a 50% chance that you will go bankrupt and default, and you believe that there is only a 10% chance, then your bank will lend you money only if you pay a much higher interest rate than what you will think appropriate. You will then consider your borrowing rate to be too high. Of course, this also breaks the equality of one fair rate at which you can borrow and lend. Your expected rate of return is now lower when you lend than when you borrow.

To avoid such situations, our perfect markets assumptions include one that posits that *everyone has the same information and agrees on what it means.*

2. Infinitely many investors and firms: This assumption really means that the market is very "deep." By itself, the assumption of the presence of many buyers and sellers defines a competitive market — one in which no buyer or seller has any unique market power. If buyers or sellers are heterogeneous, then this assumption must be slightly modified. It must be that you can easily find many of the most eager types of buyers and sellers. For example, say a truck is worth more if it is owned by a truck driver. This assumption then states that there must be a large number of truck drivers.

How could this assumption be violated? If there is only one bank that you can do business with, then this bank will want to exploit its monopoly power. It will charge you a higher interest rate if you want to borrow money than it will pay you if you want to deposit money — and you will have no good alternative.

To avoid this, our perfect markets assumptions include one that posits that *there are infinitely many buyers and sellers.*

3. No transaction costs: Transaction costs here are defined in a very broad sense, and they include indirect costs, such as your time and money to search for the best deal. In a perfect market, you can buy and sell without paying *any* such costs.

How could this assumption be violated? If it costs $1,000 to process the paperwork involved in a loan, you will incur this cost only if you borrow, but not if you save. Similarly, if it costs you 3 days of work to find the appropriate lender, it means that you will effectively have to pay more than just the borrowing rate. You will have to factor in your 3 days as a cost. Any such transaction costs make your effective borrowing interest rate higher than your effective savings interest rate.

To avoid this, our perfect markets assumptions include one that posits that *there are zero transaction costs.*

4. No taxes: More accurately, this means that there is no distorting government interference (such as government regulation), and that there are no tax advantages or disadvantages to buying or selling securities. Specifically, neither trading of the good nor its possession by one particular owner should change the total tax consequences.

How could this assumption be violated? If you have to pay taxes on interest earned, but cannot deduct taxes on interest paid, your *de facto* savings rate will be lower than your borrowing rate. Similarly, if the total taxes paid are higher when shares are traded, shares could be worth more if they were never traded to begin with. Another violation could be a government regulation requiring you to file lengthy legal documents with the SEC every time you have to sneeze — well, every time you have to execute some transaction.

To avoid this, our perfect markets assumptions include one that posits that *there are no taxes.*

These four assumptions are actually "overkill," but if they hold, you are safe. Thinking about them helps you judge how close to perfect a given market actually is. However, the real usefulness of the perfect market is *not* that you should believe that it exists in the real world. Instead, its usefulness is that it gives you some simple first-order methods and tools that help you value goods. If these assumptions do not hold, borrowing and lending rates may or may not be similar enough to allow us to still use perfect market tools or variations thereon. (And, as I already mentioned, almost all common real-world finance formulas rely on them.)

If these assumptions are far from the situation in the real world, nothing will work anymore. In fact, markets may cease to function entirely. For example, if you fear that other parties you would be transacting with are *much* better informed than you, you could only lose — the other party would take full advantage of you, selling to you only if the price is too high. You should never trade in such a market 'rigged' against you, if you can avoid it. Such a market collapse may have happened in the market for corporate bonds *for retail investors.* These bonds are traded over-the-counter, which means that the Wall Street trader on the other side of the phone tries to gauge how much an ordinary retail investor actually knows about the correct value of these bonds. As a result, retail investors are so systematically disadvantaged that it makes no sense for them to buy corporate bonds directly. Instead, they are better off buying bond funds, where someone else who does not suffer a knowledge disadvantage (a bond mutual fund) buys and sells corporate bonds on their behalves. Similarly, if transaction costs are extremely high, there may be no market in which anyone could profitably buy or sell. Fortunately, such total market collapses tend to occur only if the perfect market violations are large. With modest violations, the benefits of transacting tend to outweigh its costs to buyers and sellers, and so markets can still function. This is the kind of situation that this chapter considers.

(Omitted solvenow)

Ambiguous Value in Imperfect Markets

Why is an inequality between borrowing and lending rates so problematic? It is because it breaks the "unique value aspect" of projects. In a perfect market, project value depends *only* on the project, and not on you personally or on your cash position. You can think of this as a clean separation between the concepts of ownership and value. It also leads to the "separation of investments and financing decisions." Project owners can make investment choices based on the quality of the projects themselves, not based on their personal wealth or financing options.0.6]subsect:separating-decisionsInvestment consumption separation Indeed, the NPV formula does not have an input for your identity or current wealth — its only inputs are the project's cash flows and the rate of return on alternative investments.

For example, assume that you can lend (invest cash) and borrow money (receive cash) at the same 4% in a perfect market. What is the net present value of a project that invests $1 million today and returns $1.05 million next period? It is $9,615. It does not depend on whether you have money or not. If you do not have the $1 million today, you borrow $1.009615, invest $1 million, and hand the $1.050 million to the lender next year. But if the financial market is imperfect and the borrowing and lending rates are *not* the same, then the value of the project does depend on you, because it depends on your cash holdings. For example, assume that you can lend money (invest cash) at 3% and borrow money (receive cash) at 7%. What is the net present value of the same project that invests $1 million today and returns $1.05 million next period?

- If you have $1 million and your alternative is to invest your money in the bank, you will get only $1.03 million from the bank. You should take the project rather than invest in the bank so that you can earn $20,000 more.
- If you do not have the $1 million lying around, you will have to borrow $1 million from the bank to receive $1.05 million from the project. But because you will have to pay the bank $1.07 million, you will lose $20,000 net. You should not take the project.

The value of the project and your best decision whether to take the project or not now depends on how much cash you have. Consequently, the separation between your project choice and your financial position breaks down. Having to take your current cash holdings into account when making investment choices makes capital budgeting decisions more difficult. In this example, it is still fairly easy: If you have a lot of money lying around, you should take the project. If you do not, you should not take it. But think about projects that have cash inflows and outflows in the future and how your decisions could interact with your own wealth positions in the future. This can become vexingly difficult. You can also see that the project value is no longer unique in imperfect markets. In our example, it could be anything between +$19.42 thousand ($1.05 million discounted at 3%) and –$18.69 thousand ($1.050 million discounted at 7%). The same ambiguity applies to ownership. Your capital budgeting decision can be different depending on whether you already own the project versus when you are just contemplating buying it. Again, your identity matters to the value of the project.2]pg:rationingcredit rationing

Begin Important

If the market is not perfect, the separation of ownership and value breaks down. Therefore, project value is no longer unique. It can depend on who owns the project. **End Important**

➤ Do You Always Get What You Pay For?

Reflect a little on the insight that projects may not have unique values. You surely have heard the saying that "it's only worth what people are willing to pay for it" and the claim that some item "is worth much more than it is being sold for." Which is correct? Are there any good deals? The answer is that both can be correct and neither can be correct. The first claim is really meaningful only to the extent that markets are *perfect*: If a market is perfect, items are indeed worth exactly what buyers are willing to pay for them. The second claim is meaningful only to the extent that markets are *imperfect*: If a market is imperfect, items have no unique value. Different people can place different values on the item, and some third party may consider an item worth much more than what it was sold for.

Thus, when someone claims that a stock or firm is really worth more than he or she is selling it for, there are only a small number of explanations:

1. There may be pure kindheartedness toward any buyer, or a desire by a seller to lose wealth. Not very likely.
2. The seller may not have access to a perfect market to sell the goods. This may make the seller accept a low amount of money for the good, so depending on how you look at it, the good may be sold for more or less than what the seller thinks it is worth.
3. The market is perfect and the seller may be committing a conceptual mistake. The good is worth neither more nor less than what it is being sold for — it is worth exactly how much it is being sold for.
4. The seller may be lying and is using this claim as a sales tactic.

Sales: 06-16

(Omitted solvenow)

Social Value and Surplus

Perfect markets are not just privately useful but are also socially useful. If a market is perfect, buyers and sellers need not worry that one deal is better than another — that buying is better than selling, or vice-versa. For example, consider gasoline and imagine that you do not yet know when and where on your road trip you will need to pump more gas. Unlike shares of stock, gas is not the same good everywhere: Gas in one location can be more valuable than gas in another location (as anyone who has ever run out of gas can testify). But in populated areas, the market for gasoline is pretty competitive and close to perfect — there are many buyers (drivers) and sellers (gas stations). This makes it likely that the first gas station you see will have a reasonable price. If you drive by the first gas station and it advertises a price of $3 per gallon, it is unlikely that you will find another gas station within a couple of miles offering the same gas for $2 per gallon or $4 per gallon. Chances are that "the price is fair," or this particular gas station would probably have disappeared by now. (The same applies, of course, in many financial markets, such as those for large company stocks, Treasury bonds, or certain types of mortgages.) As long as the market is very competitive — or better yet, perfect — most deals are likely to be fair deals.

There is an important conceptual twist here: If you are paying what an item is worth, it does not necessarily mean that you are paying what you *personally* value the good at. For example, if you are running out of gas and you are bad at pushing a 2-ton vehicle, you might very well be willing to pay $10 per gallon — but fortunately, all you need to pay in a competitive market is the market price. The difference between what you personally value a good for and what you pay for it is called your "surplus." Although everyone is paying what the good is worth in a perfect market, most buyers and sellers can come away being better off — only the very last marginal buyer and seller are indifferent.

(Omitted solvenow)

11.2 Opinions, Disagreements, and Insider Information

What can you do if you think each one of the perfect market assumptions fails? You need to learn both how to judge the degree to which markets are imperfect and how to deal with them as a real-world investor or manager. (Even if there is no unique value, you can still learn how to think about maximizing your own wealth.) The remainder of the chapter thus explores the extent of market imperfections, what can mitigate them, and how you should work when they don't hold.

We begin with the effects of disagreements, the violation of the first perfect market assumption that everyone has the same opinion. Like all perfect-market assumptions, this one works well in some situations and poorly in others.

Expected Return Differences or Promised Return Differences?

The assumption of no disagreement is only relevant in a world of uncertainty — it would be absurd to believe that differences in opinion could exist if there were no uncertainty. So what happens if the lender and borrower have different information or different judgments about the same information? Most prominently, they could disagree about the default risk. For example, if you have no credit history, then a lender who does not know you might be especially afraid of not receiving promised repayments from you — from the perspective of such a lender, you would be extremely high-risk. Your lender might estimate your appropriate default probability to be 30% and thus may demand an appropriate default premium from you of, say, 10% — an interest rate similar to what credit card vendors are charging. On the other hand, *you* may know that you will indeed return the lender's money, because you know that you will work hard and that you will have the money for sure. In your opinion, a fair and appropriate default premium should therefore be 0%.

When your potential lender and you have different opinions, you will face different expected interest rates depending on whether you want to save or borrow. You can use your knowledge from Chapter to work an example to understand the difference between a perfect and an imperfect market scenario.

Perfect Markets: Assume that the bank and you agree that you have a 20% probability of default, in which case you will not repay anything. For simplicity, assume risk neutrality and that the appropriate interest rate is 5%. Solving $80\% \cdot r + 20\% \cdot (-100\%) = 5\%$ for the interest rate that you would have to promise results in $r = 31.25\%$. This gives the bank an expected rate of return of 5%. In contrast, the bank is government-insured, so if you deposit your money with it, it would be default-free.

	Promised	**Expected**
Your Savings Rate	5%	5%
Your Borrowing Rate	31.25%	5%

Although your quoted interest rate is higher by the credit1]subsect:deconstructdefaultpremiumCredit spreads spread, if you want to borrow, your cost of capital is still the same 5% either way.

Imperfect Markets: Now assume that the bank and you disagree about your default probability. The bank believes that it is 30% — it could be that it has experienced such a default rate for borrowers who seemed to look similar from the perspective of your bank. In contrast, you believe that your default probability is 10%. The bank will therefore quote you an interest rate of $70\% \cdot r + 30\% \cdot (-100\%) = 5\% \Longrightarrow r = 50\%$. Alas, you believe that the expected rate of return at the 50% quoted interest rate is $90\% \cdot 50\% + 10\% \cdot (-100\%) = 35\%$.

	Promised	Expected
Your Savings Rate	5%	5%
Your Borrowing Rate	50% from the bank's perspective	5%
Your Borrowing Rate	50% from your perspective	35%

The disagreements (information differences) are now causing differences in *expected* returns. The borrowing and lending *expected* rates of return are no longer the same. If the bank is wrong, your cost of capital now depends on whether you want to borrow or lend. And even if the bank is right, from your (wrong) perspective, you are still facing different borrowing and lending rates.

Begin Important

- The fact that credit spreads reflect a default premium — a difference between the *promised* rate of return and the *expected* rate of return — is not a market imperfection.
- The fact that credit spreads reflect differences in opinion between borrower and lender — a difference about the two assessed *expected* rates of return — is a market imperfection.

End Important

Can we estimate the default premium, which is the difference between promised yields and "the market's" expected rate of return? Well, "sort of." As with stocks and the equity premium, we cannot measure the expected rate of return, but we can measure the ex-post realized average rates of return. Over millennia, if everything were stable, averaging the latter would give you a perfect proxy of the former. Over decades, which is all we have, it can only give us an imperfect proxy.

For example, Vanguard has been selling three different bond funds which differ in including bonds with different default risks: the `VFITX` government bond fund, the `VFICX` investment grade corporate bond fund, and the `VWEHX` junk-bond corporate bond fund. All three buy and hold intermediate-term bonds, with maturities and durations of about 5-6 years on average. From 2005 to 2021, a typical quoted spread over safe government bonds (`VFITX`) was about 130 bps for `VFICX` and 400 bps for `VWEHX`, though with dramatic spikes during the Great Recession. Yet, for a buy-and-hold investor, `VFICX` beat `VFITX` by a more modest 100 bps (4.7% vs. 3.7%) and `VWEHX` beat `VFITX` by 230 bps (6.0% vs 3.7%). After state and local income taxes on distributions (say 10% in high-tax states), the realized performance spreads further shrink from 100 bps to about 50 bps, and from 230 bps to about 170 bps. To earn these most modest 50 bps and 170 bps yields (rather than the promised 130 bps and 400 bps quoted arithmetic-type spreads), Vanguard's investment-grade and speculative-grade bonds had to take on the default risk and various imperfect-market problems (such as liquidity premia). Are these 17 years representative of what investors might have expected? They included a lot of good years but also the Great Recession of 2008-9, in which many non-investment grade bonds defaulted. Thus, this sample may have been reasonably representative of good times and bad times — in which case the default (and risk and liquidity and other) premia could be assessed as 50 bps for investment grade bonds and 170 bps for non-investment grade bonds.

(Omitted solvenow)

Differences in Information: Covenants, Collateral, Credit Ratings

If you are an entrepreneur who wants to start a company, what can you do to reduce your cost of capital? The answer is that it is in your interest to disclose to the lender all the information you can — provided you are the type of entrepreneur who is likely to pay back the loan. You want to reduce the lender's doubt about future repayment. Unfortunately, this can be very difficult. The lender can neither peer into your brain nor give you a good lie detector test. Even after you have done everything possible to reduce the lender's doubts about you (provided your credit history, collateral, and so on), there will still be some residual information differences — they are just a fact of life. To the extent that you can reduce such information differences, your firm will be able to enjoy lower costs of capital. Also, if you as a borrower fail to give your best to convince the lender of your quality, then the lender should assume that you are not an average company but instead the very worst — or else you would have tried to communicate as much as possible.

There are at least three important mechanisms that have evolved to alleviate such information differences. The first mechanism is covenants, which are contractual agreements that specify upfront what a debtor must do to maintain credit. They can include such requirements as the maintenance of insurance or a minimum corporate value. The second mechanism is collateral, which are assets that the creditor can repossess if payments are not made — anything that inflicts pain on the debtor will do. For example, if defaulting debtors were thrown into debtors' prison (as they often were until the nineteenth century), the promise to repay would be more credible and lenders would be more inclined to provide funding at lower rates. Of course, for the unlucky few who just happened to suffer incredibly bad luck ex-post, debtors' prison had some definite drawbacks.

(Omitted anecdote)

The third mechanism to alleviate repayment uncertainty is a credit rating, which is a history of past payments to help assess the probability of future default. This is why you need to give your Social Security number if you want to take out a substantial personal loan — the lender will check up on you. The same is true for large corporations. It may be easier to judge corporate default risk for large companies than personal default risk, but it is still not easy and it costs both time and money. You already learned about these credit ratings in Section 2]subsect:creditratingsCredit ratings.

Unfortunately, although bond rating agencies update their ratings if the condition of the firm changes, the empirical evidence suggests that bond ratings are not very good in helping an investor earn better rates of return.

In fact, the ratings seem to respond more to *past* drops in the value of the underlying bonds than vice-versa. The rating agencies seem to be more reactive than proactive. (The poor quality and systematic manipulation of debt ratings by investment banks also played an enabling role in the Great Recession.)

Let me close with a philosophical observation: U.S. and European financial markets are truly amazing. People who would never lend their neighbors a few thousand dollars (fearing that they would not get it back) have no second thoughts about lending total strangers in anonymous markets their entire lives' savings. It is the combination of the governance of repayments and risk-spreading that has allowed our financial markets to finance our real enterprises and thus to develop our economies so well, even in the presence of great risk and uncertainty for these undertakings. Yes, it will never be perfectly perfect, of course. Yes, there are problems in the U.S. financial markets, too, but their relative magnitudes have been fairly small. By and large, issues of fraud, credit, and trust seem to be under control most of the time. Banks are a vital component of our economic system. In contrast, many hundreds of million Indians and Africans do not have access either to convenient borrowing or saving institutions and markets even as of 2020. Many are forced to keep their lives' savings in gold under their mattresses. This leaves them with fewer opportunities and more exposure to theft and corruption.

nate
o to
orld.

(Omitted solvenow)

Differences in Opinion

Not all differences in valuation are the results of objective assessments with different information. The premise of Behavioral Finance is that executives and investors are human and subject to predictable biases. Perhaps the most important human bias is overconfidence. The evidence strongly suggests that financial markets are not stupid — it is very difficult to earn an unusual rate of return. Thus, most stocks should be fairly valued. Insiders have more confidential information and thus should have a different perception. Logically, half of them should think that the market is too pessimistic and half should think the market is too optimistic. Instead, regardless of recent stock-market condition, most CFOs consider their own companies undervalued. This means that either the stock market undervalues most companies (unlikely!), or most CFOs operate under mistaken perceptions. (Entrepreneurs are even more notoriously overoptimistic than CFOs.)

(Omitted tbl)

11.3 Market Depth and Transaction Costs

Our second perfect market assumption states that markets are very deep, consisting of many buyers and sellers. If there were only one lender, this lender would have market power over you. Of course, such a lender would exploit her power by charging you a higher borrowing rate and offering you a lower deposit interest rate. This extreme form of market power is called a monopoly, but there are many milder forms of such power, too. For example, if you are already shopping in a grocery store, this store has a degree of market power over you. Even if the milk is 3 cents more expensive than in another store, you will still buy the milk where you are. Or say there is only one ATM close to you. In principle, you could get capital from any number of banks, but locally there is really only this one provider. Fortunately, such uniqueness of capital provision is rarely an important issue in the United States for corporations, especially large ones.

So let's move on to the third perfect markets assumption: the role of transaction costs. Transaction costs drive a wedge between borrowing and lending rates. For example, if it is difficult and costly to administer loans, an investor must charge you a higher borrowing rate than deposit rate just to break even. This is the subject of this section, in which you will learn how corporations and individuals should handle transaction costs.

Typical Costs When Trading Real Goods — Real Estate

When you engage in transactions — that is, purchases or sales — you face costs to facilitate them. One way to think about the magnitude of transaction costs is to compute how much is lost if you decided that you have made a mistake the instant after a purchase, which you now want to undo by reselling. Real estate — most people's biggest asset — is a perfect example to illustrate transaction costs. What does selling or buying a house really cost?

Direct costs such as brokerage commissions: Housing transaction costs are so high and so important that they are worth a digression. In the United States, if a house is sold, the seller's broker typically receives 6% of the value of the house as commission (and splits this commission with the buyer's real-estate agent). Thus, if a real-estate agent sells your house for $300,000, her commission is $18,000 (which she usually splits with the buyer's broker). Put differently, without an agent, the buyer and seller could have split the $18,000 between themselves.

Although only the seller pays the broker's cost, it makes sense to think of transaction costs in terms of round-trip costs — how much worse off you are if you buy and then immediately sell. You would be mistaken if you thought that when you buy a house, you have not incurred any transaction costs because the seller had to pay them — you have incurred an implicit transaction cost in the future when you need to resell your investment. Of course, you usually do not sell assets immediately, so you should not forget about the timing of your future selling transaction costs in your NPV calculations.

If you borrow to finance the investment, transaction costs may be higher than you think. The real-estate agent earns 6% of the house value, not 6% of the amount of money you put into the house. On a house purchase of $500,000, the typical loan is 80% of the purchase price, or $400,000, leaving you to put in $100,000 in equity. Selling the house the day after the purchase reduces your wealth of $100,000 by the commission of $30,000 — for an investment rate of return of –30%. This is not a risk component; it is a pure and certain transaction cost.

How good is your purchase if the house price decreases or increases by 10%? If house prices decline by 10% (or if you overpaid by 10%), the house can only be resold for $450,000, which leaves $423,000 after agent commissions. As the house owner, you are left with $23,000 on a $100,000 investment. A 10% decline in real estate values has reduced your net worth by 77%! In comparison, a 10% increase in real estate values increases the value of the house to $550,000, which means that $517,000 is left after real estate commissions. Your rate of return after this equally-sized magnitude is thus only 17%. If a 10% increase and a 10% decrease are equally likely, your instant expected loss is 30%!

(Omitted anecdote)

In addition to direct agent commissions, there are also many other direct transaction costs. These can range from advertising, to insurance company payments, to house inspectors, to the local land registry, to postage — all of which cost the parties money.

Indirect costs such as opportunity costs: Then there is the seller's and buyer's time required to learn as much as possible about the value of the house, and the effort involved to help the agent sell the house. These may be significant costs, even if they involve no cash outlay. If the house cannot be sold immediately but stays empty for a while, the foregone rent is part of the transaction costs. The implicit cost of not having the house put to its best alternative use is called an opportunity cost — the cost of foregoing the next-best choice (here, renting it out). Opportunity costs are just as real as direct cash costs.

Typical Costs When Trading Financial Securities

Transactions in financial markets also incur transaction costs. If an investor wants to buy or sell shares, the broker charges a fee, as does the stock exchange that facilitates the transaction. In addition, investors have to consider their time to communicate with the broker to initiate the purchase or sale of a stock as an opportunity cost.

Direct costs: Still, the transaction costs for selling financial instruments are much lower than they are for most other goods.

First, as of 2022, most brokers do not charge fixed commission fees per transaction to most investors. Computers don't care much if you transact 5 times or 10 times. You have to do most of the work typing in the orders anyway. Instead, brokers make money by earning small kickbacks (called payment for order flow) for routing your trade to a particular market-maker who earns the bid-ask spread (the difference between the bid price and the ask price).[*]

For example, on Dec 31, 2021, the mid-price for Intel (`INTC`) shares was $51.51. But you could neither buy nor sell at $51.51. Instead, the $51.51 was really just the average of two prices: the bid price of $51.50, at which the market maker was willing to buy shares and the ask price of $51.52, at which the market maker was willing to sell shares. (These prices are themselves determined both by other investors and the inventory of the market maker.) For a small transaction, up to 100 shares, the bid and ask prices are even guaranteed (until the next update). For larger lots, the prices may also move a little. However, for a company as liquidly traded company as Intel, even $50,000 would be unlikely to move prices. Therefore, you could (probably) purchase shares at $51.50 and sell them at $51.52, still a loss of "only" 2 cents. This amounts to round-trip transaction costs of ($0.02)/$51.52 ≈ 0.04%.

(Omitted fig)

Figure **??** shows the typical quoted bid-ask spread for publicly-traded stocks at the end of 2021. By rule, the minimum bid-ask spread is 1 cent per share. When not constrained by this floor, the typical (median) bid-ask spread was about 0.05 percent.

If you want to trade a larger lot, say $1,000,000 worth of shares, in Intel, the cost may be more akin to 0.06-0.08% than 0.04%, because prices tend to move a little against you while your trade is executed. In any case, if you

---

[*]With more aggressive regulations on minimum tick prices, retail transaction costs could probably come down further. Yet few retail investors really care about 0.04% vs. 0.03%. A Robinhood investor trading one $50 Ford (`F`) share would incur the 1 cent minimum trading cost. Getting it down to 0.5 cents is probably not high on her agenda.

traded these kinds of larger lots regularly, there are some cheaper institutional-trading venues than retail brokers. Nevertheless, even 0.08% is still only $800.

(Omitted anecdote)

You may sometimes read about high-frequency traders (HFT), who run algorithms to strategically pick off fractions of pennies because they have a nano-second earlier access to trading. By smart trading, they can even sometimes earn the spread, in effect becoming market makers themselves. Whether this is a problem or not can be debated, but if it ever was, the problem seems to be taking care of itself. There are now dozens of HFTs competing against one another for the business of buying and selling shares from the rest of us. They have almost surely competed away much of their possible excess profits. Moreover, other venues with better market structures are also appearing. Even if this game was (slightly) rigged a few years ago, it's no longer a major concern today.

Indirect costs such as opportunity costs: Investors do not need to spend a lot of time to find out the latest price of the stock: It is instantly available from many sources (e.g., from *YAHOO!*FINANCE). The information research costs are very low: Unlike a house, the market value of a stock is immediately known. Finally, buyers can be found practically instantaneously, so searching and waiting costs are also very low. In contrast, count on many anxiety-ridden waiting months when you want to sell your house.

Bond Transaction Costs

(Omitted fig)

Bonds are much less liquid than stocks and much more expensive to buy and sell. They also trade only in larger denominations. Figure **??** plots typical bond bid-ask spreads in early 2022. The plot shows that credit rating makes relatively little difference. However, trading longer-term maturities is more expensive. Most importantly, trading corporate bonds is about one order of magnitude more expensive than trading stocks. Most retail investors are better served buying a corporate bond mutual fund (like a Vanguard Short-Term or Long-Term Investment Grade bond fund (VFSTX or VWESX) or a Speculative High-Yield Corporate Bond fund (VWEHX). 2]fig:byrating-feb22Corp Yields on 2022/02/22 If absolutely need be, they should only purchase corporate bonds at issue and plan to hold them in their account until maturity. Trading them regularly would quickly erode their yield.

Comparing Financial Transaction Costs To Housing Transaction Costs

Let's compare the transaction costs in buying and selling financial securities to those of a house. Aside from the direct real estate broker fees of 6% (for the $100,000 equity investment in the $500,000 house, this comes to $30,000 for a round-trip transaction), you must add the other fees and waiting time. Chances are that you will be in for other transaction costs — say, another $10,000.

		Real Estate	Financial	
Cost Type	Explanation	(House)	Bond	Stock
Direct	Typical round-trip commission, etc.	≥6%	1%	0.1%
Search/Research	Time to determine fair price	High	Zero	Zero
Search/Liquidity	Time waiting to find buyer	Variable	Zero	Zero

And houses are just one example: Many transactions of physical goods or labor services can incur similarly high transaction costs.

In contrast, if you want to buy or sell 100 shares in, say, Microsoft stock, your transaction costs are relatively low. Because there are many buyers and many sellers, financial transaction costs are comparably tiny. Even for a $100,000 equity investment in a medium-sized firm's stock, the transaction costs are typically only about $300–$500. It may not be a perfectly correct assumption that the market for trading large stocks is perfect, but it is not far off. It certainly is convenient to assume that financial transaction costs are zero. For an individual buying and selling ordinary stocks only rarely (a buy-and-hold investor), a zero-transaction-cost assumption is often quite reasonable.

It is only for day trader[day traders] — someone who buys and sells stocks daily — that our perfect-market assumption for stocks would be inappropriate. For bonds, the best characterization would be "it depends." If you do it through a good mutual fund, the costs are more like those on stocks. If you do it yourself, they are far more troublesome.

(Omitted solvenow)

Transaction Costs in Returns and Net Present Values

As an investor, you usually care about rates of return *after* all transaction costs have been taken into account, not about pre-transaction-cost rates of return from quoted prices. Let's work out how you should take these transaction costs on both sides (buy and sell) into account.

Return to our housing example. If you purchase a house for $1,000,000 and you sell it to the next buyer at $1,100,000 through a broker, your rate of return is not 10%. At selling time, the broker charges you a 6% commission. There are also some other costs that reduce the amount of money you receive, not to mention your many opportunity costs. Say these costs amount to $70,000 in total. In addition, even when you purchased the house, you most likely had to pay some extra costs (such as an escrow transfer fee) above and beyond the $1,000,000 — say, $5,000. Your rate of return would therefore not be $1,100,000/$1,000,000 – 1 = 10%, but only

(Omitted eq)

Note how the $5,000 must be added to, not subtracted from, the price you originally paid. The price you paid was ultimately higher than $1,000,000. The $5,000 works against you. Incidentally, in order to make their returns look more appealing, many professional fund managers quote their investors' rates of return before taking their own fees (transaction costs) into account. They add a footnote at the bottom that satisfies the lawyers so that you cannot sue the fund for having been misled — you are supposed to know how to adjust the returns to take these transaction costs into account.

How do you take care of transaction costs in present value calculations? This is relatively straightforward. In the example, you put in $1,005,000 and receive $1,030,000 — say, after one year:

(Omitted eq)

The only thing you must still take care of is to quote your opportunity cost of capital also in after-transaction cost terms. You may not be able to get a 10% rate of return in comparable investments either, because you may also be required to pay a transaction cost on them. In this case, assume that your alternative investment with similar characteristics in the financial markets (not the housing markets) would earn an 8% per year rate of return, but with a 50-basis-point transaction cost. Your project would then have an appropriate NPV of

(Omitted eq)

(Omitted solvenow)

The Value of Liquidity

When *future* transaction costs influence your upfront willingness to buy an asset, proper pricing gets even more interesting and complex. You might not want to purchase a house even if you *expect* to recoup your transaction costs, because you dislike the fact that you do not know whether it will be easy or hard to resell. After all, if you purchase a stock or bond instead, you know you can resell without much of a transaction cost whenever you want.

What would make you want to take the risk of sitting on a house for months without being able to sell it? To get you to buy a house would require the seller to compensate you. The seller would have to offer you a liquidity premium — an extra expected rate of return to compensate you for your willingness to hold an asset that you may find difficult to convert into cash if a need were to arise. The liquidity analogy comes from physics. In the same way that physical movement is impeded by physical friction, economic transactions are impeded by transaction costs.

Housing may be an extreme example, but liquidity effects appear to be important everywhere, even in financial markets with their low transaction costs. (Some financial markets are generally considered low-friction, or even close to frictionless.) Even finance professors and the best fund managers do not yet fully understand liquidity premiums, but we do know that they can be very important. In financial crises, like 2008, liquidity seems to have been the only thing that was really important. Let's look at some examples of where liquidity premiums seem to play important roles.

➤ Treasury Bonds

Believe it or not, even Treasuries have differences in liquidity! The most recently issued Treasury of a particular maturity is called on-the-run. These bonds account for more than half of the total daily trading volume, yet less than 5% of the outstanding market cap. Every bond trader who wants to trade a bond with roughly this maturity focuses on this particular bond. This makes it easier to buy and sell the on-the-run bond compared to a similar but not identical off-the-run bond. In 2021, the typical on-the-run bond traded for about 5 basis points less than the equivalent off-the-run Treasury. In other words, you would have been able to buy the off-the-run bond at a lower price than the on-the-run bond, but resale would have been more difficult.

The reason why you might want to buy the on-the-run bond, even though it had a higher price, would be that you could resell it much more quickly and easily than the equivalent off-the-run bond. Of course, as the date approaches when this 10-year bond is about to lose its on-the-run designation and another bond is about to become the on-the-run 10-year bond, the previous on-the-run premium drops in value.

In a perfect world, there should be no difference between these two types of bonds. Yet when a two-year bond is on-the-run, its bid-ask spread is on average about 1 basis point lower, and it offers on average 0.6 basis points less in yield. For a ten-year bond, both the bid-ask spread and the yield difference between the on-the-run and off-the-run Treasury are usually about 3 basis points. This can only be explained by an investor preference for the immediate liquidity of the current on-the-run bond.

➤ Corporate Bonds

(Omitted fig)

Figure **??** shows how liquidity and other bond characteristics have affected corporate bond yields and returns over U.S. Treasury over the last two decades. Let's just look at the means — the medians are a bit smaller (not just in font size), so you can fill in the gaps.

First focus on the columns. The highest-rated bonds offered a yield spread of about 1% over Treasury, and ultimately paid about 0.2% per year over Treasury (the rest being lost in defaults). Mid-rate investment grade bonds offered a yield spread of 1.5%, and ultimately paid about 1.0%. Non-investment grade (junk) bonds offered 4.5%, and ultimately paid about 3.5%. Many institutional investors are only allowed to hold investment-grade bonds. This can explain the increase in the liquidity premium from BBB to BB-F. The next columns show that there was no duration premium for longer-term corporate bonds over U.S. Treasuries — the yield spread was a little above 2% regardless of length, of which about 1.5% paid out. This means that the corporate bond yield curve had the same slope as the Treasury yield curve. The last three columns show that the least volatile bonds offered about 1.3% above Treasury, medium volatile bonds about 2%, and highly volatile bonds about 3%. This resulted in returns of about 1%, 1.5%, and 2%, respectively.

Now focus on liquidity premia, which are the differences between the high and low rows. Less liquid bonds offered more yield only for low-rated bonds (470 vs. 448 bps), low-volatility bonds (143 vs. 121 bps), and medium-volatility bonds (234 vs. 228 bps). Taking on this extra liquidity risk interacted with default risk in a way that resulted in non-obvious rewards. Junk bonds, low- and mid-duration, and low-volatility bonds all yielded an extra rate of return of about 0.4% per annum. However, long-duration bonds with low liquidity offered only a 7 bps spread above high liquidity — which was so low that investors ended up earning *less* for having taken on this liquidity risk. (Risk means that one sometimes does worse than expected! Otherwise, it wouldn't be risk.)

The liquidity spread is however quite time-varying. For example, in the Great Recession, the average liquidity spread shot up from the 10-20 bps here to about 100-200 bps. In a sense, this is much of the point of liquidity — the ability to sell quickly when the overall going gets rough.

➤ Liquidity Provision As a Business: Market Making

You can think of a market maker on an exchange as someone who is providing liquidity. As a retail investor, you can sell your securities to the market maker in an instant, and it is up to the market maker to find some other investor who wants to hold it long term. To provide this liquidity, the market maker earns the bid-ask spread — a part of the liquidity premium.

The provision of liquidity in markets of any kind is a common business. For example, you can think of antique stores or used car dealerships as liquidity providers that try to buy cheap (being a standby buyer) and sell expensive (being a standby seller). Being a liquidity provider can require big risks and capital outlays. If it were easy, everyone could do it — and then competition would ensure that there would be no more great profits in liquidity provision!

➤ Liquidity Runs

The most remarkable empirical regularity about liquidity, however, is that every few years, investors in all markets suddenly seem to prefer only the most liquid securities. This is called a flight to quality or run on liquidity. In such situations, the spreads on almost all bonds — regardless of whether they are Latin American, European, corporate, mortgage-related, and so on — relative to Treasuries tend to widen all at the same time.

In early 2008, with the Great Recession, the U.S. economy was facing just such a run on liquidity. It started in the mortgage sector, then spread to many other bonds. Every fund and bank was afraid that its investors would pull their lines of credit. Thus, they themselves were pulling back all lines of credit that they had extended to their clients (often other banks and funds). Many were selling even highly rated securities for low prices (sometimes fire-sale prices), just to avoid being caught themselves in an even worse liquidity run. There were many extremely curious pricing oddities during the 2008 liquidity run, but they were difficult to exploit by arbitrageurs (because no one would trust lending them the money to execute these arbitrages). For example, two-year bonds issued by a federal government agency, GNMA, and thus always fully backed by the federal government, traded at a full 200 basis points quoted yield above the equivalent Treasuries (i.e., at a dramatically lower price).

Selling liquidity in order to collect the liquidity premium is also a very common method for Wall Street firms and hedge funds to make money — perhaps even *the* most common. If you know you will not need liquidity at sudden notice or that you want to hold bonds to maturity, it can make sense to buy less-liquid securities to earn the liquidity premium. A sample strategy might be to buy illiquid corporate bonds, financed with cheaper borrowed money. Most of the time, this strategy makes modest amounts of money consistently — except when a flight to liquidity occurs and liquidity spreads widen.

Exactly such a situation led to the collapse of a well-known hedge fund named Long-Term Capital Management (LTCM) in 1998. After Russia defaulted on its debt, the spreads on almost every bond widened — the average corporate bond spread in the United States rose from about 4% to about 8% in one week! LTCM simply could not find any buyers for its large holdings of non-Treasury bonds. On the other hand, those funds that could hold onto their positions throughout the crisis or that provided extra liquidity (buying securities that were now very cheap) did extremely well when liquidity returned to normal and their illiquid securities went back up in price. The same fate probably befell many financial firms in the Great Recession — first and foremost Lehman Bros and AIG, both of which were giants that collapsed. Their own financiers demanded their money back quickly, but there was no liquid market for them to unwind their positions quickly.

(Omitted solvenow)

11.4 Taxes

04-06 *The art of taxation consists in so plucking the goose as to get the most feathers with the least hissing.*

Jean-Baptiste Colbert

Code: 04-10

Certainty? In this world nothing is certain but death and taxes. Benjamin Franklin

Our fourth violation of market perfection is taxes. They are pervasive and are often an economically large component of project returns. The actual tax code itself is very complex, and its details change every year, but the basics have remained in place for a long time and are similar in most countries. Let me summarize briefly what you need to know for this book.

The Basics of (Federal) Income Taxes

Congress makes tax laws and the Internal Revenue Service (IRS) enforces them. The U.S. taxes individuals and corporations similarly. (There are some differences, but discussing them would be going down Alice in Wonderland's rabbit hole.) Gross income is adjusted by a set of allowable deductions into taxable income, and a (progressive) tax rate is applied. Before-tax expense[Before-tax expenses] (deductions) are better for taxpayers than after-tax expense[after-tax expenses]. For example, if you earn $100,000 and there was only one 40% bracket, a $50,000 before-tax expense would leave you

(Omitted eq)

while the same $50,000 as an after-tax expense would leave you with only

(Omitted eq)

Among the most important deductible items for both corporations and individuals are interest payments, although individuals can deduct them only for mortgages. In addition, there are some other deductions such as pension contributions. There are also some nonprofit investors (such as pension funds) that are entirely tax-exempt.

The tax code categorizes income into four different classes: ordinary income, interest income, dividend income, and capital gains. The tax rates on these classes differ, as does the ability to apply deductions on them to reduce the income tax burden.

Ordinary income applies to most income that is not derived from financial investments (such as wages). Individuals are allowed only very few deductions on ordinary income, and the tax rate is the highest. The highest marginal federal income tax rate in 2022 was 37%. Most U.S. states also have an income tax, which can add up to another 10-15% on top of the federal rate.

Interest income is basically treated like ordinary income.

Dividend income from shares in qualifying corporations (either U.S. or from a country with a tax treaty) are taxed at a lower rate, often about half that of ordinary income.

Capital gains on assets owned for one year or more are taxed at the lower rates, just like dividends. (Assets held for less than one year are taxed essentially at the same rate as ordinary income.) In addition, your capital losses are deductible against your capital gains. And unlike any other income, which is taxed every year, both short-term and long-term capital gains are taxed only when realized. Moreover, if you have moved for one year to a state with no income taxes, then you can realize your capital gain without paying state income tax — even if the appreciation itself has occurred mostly while you were living in a high-income tax state. (It is no accident that many senior citizens have been moving to Florida to avoid state income tax on their accumulated capital gains.)

From the perspective of an investor, capital gains are preferable to dividend income, and both are preferable to interest and ordinary income.

The average tax rate (the ratio of paid taxes to taxable income) is lower than the marginal tax rate (the rate on the last dollar of income), because lower marginal tax rates are applied to your first few dollars of income in the progressive U.S. tax system. For example, in 2022, the first $10,275 were taxed at 10%, the next $31,500 at 12%, and so on. Thus, ignoring a variety of subsequent adjustments, if you earned $30,000, you would have paid taxes of

(Omitted eq)

Therefore, your marginal tax rate — the one applicable to your last dollar of income — was 12%, while your average tax rate was about 11.3%. Economists almost always work only with marginal tax rates, because they are relevant to your choice of working just a little more or less. For large corporations, the distinction is often minor, because the federal corporate income tax rate reaches its maximum rate at around $100,000 of income. A corporation that earns or loses $10 million has an average tax rate that is, for all practical purposes, the same as its marginal tax rate. In 2022, it stands at its post-World War II lowest rate of of 21%, although the Biden administration is trying to raise it back to 28%.

Of course, there are also other important taxes, such as state income taxes, Social Security and Medicare taxes, property taxes, sales taxes, and so on. From the 1980s to 2017, an alternative tax system, the alternative minimum tax (AMT), was quite important, applying eventually to as many as 5 million tax payers. It differed, e.g., in not allowing tax payers to deduct state income taxes from their income. The tax reform reduced this to about 200,000 payers, although it did so largely by removing the state-income tax deduction also from the regular income tax calculation.

Because the AMT categorizes most income the same way, we won't distinguish between the standard income tax and the alternative minimum tax. If you have to file in multiple states, the details can become hair-raisingly complex. Professional athletes have to pay taxes in every state in which they have played a game, for example. Some retailers have to handle hundreds of (sales) tax authorities in the United States alone. It gets worse if multiple countries are involved.

Begin Important

- Remember that there are some tax-exempt investors, such as pension funds.
- You must understand how income taxes are computed (the principles, not the details), how to find the marginal tax rate, how to compute the average tax rate, and why the average tax rate is usually lower than the marginal tax rate.
- Expenses that can be paid from before-tax income are better than expenses that must be paid from after-tax income. Specifically, interest expenses are tax-deductible and thus better for the taxpayer.
- Capital gains and secondarily dividend income enjoy preferential tax treatment for the recipient, relative to interest and ordinary income.

End Important

(Omitted solvenow)

The Effect of Taxes on Rates of Return

How does finance work if there are income taxes? Mechanically, taxes are similar to transaction costs — they take a "cut," which makes investments less profitable. One difference between them is that income taxes are higher on more profitable transactions, whereas plain transaction costs are the same whether you made or lost money. And, of course, taxes often have many more nuances. A second and perhaps more important difference is that taxes are often orders of magnitude bigger and thus more important than ordinary transaction costs — except in illustrative textbook examples. For many investors and corporations, tax planning is an issue of first-order importance.

In the end, we assume that investors care about after-tax returns, not about before-tax returns. If they do, then it should not matter whether one receives \$100 that has to be taxed at 50% or whether one receives \$50 that does not have to be taxed. This leads to a recommendation analogous to that for transaction costs — *work only in after-tax money*. For example, say you invest \$100,000 in after-tax money to earn a return of \$160,000. Your marginal tax rate is 25%. Taxes are on the net return of \$60,000, so your after-tax net return is

(Omitted eq)

(The tax rate is commonly abbreviated with the Greek letter τ, tau.) In addition, you will receive your original investment back, so your after-tax rate of return is

(Omitted eq)

➤ Tax-Exempt Bonds and the Marginal Investor

In the United States, interest paid on bonds issued by smaller governmental entities is legally tax-exempt. (The Constitution's authors did not want to have the federal government burden states' or local governments' efforts to raise money.) If you own one of these bonds, you do not need to declare the interest on your federal income tax forms, and sometimes not even on your state's income tax form, either. (The arrangement differs from bond to bond.) The most prominent tax-exempt bonds are often just called municipal bond[municipal bonds], or muni[munis] for short. As their name suggests, many are issued by municipalities such as the City of Los Angeles (CA) or the City of Canton (OH). State bonds are also categorized as muni bonds, because they are also exempt from federal income tax. Unfortunately, unlike the U.S. Treasury, municipalities can and have gone bankrupt, so their bonds may not fully repay. (For example, Orange County California prominently defaulted in December 1994.) Still, many muni bonds are fairly safe AAA credit. Tax-exempt bonds are often best compared to taxable corporate bonds with similar bond ratings. The difference between the prevailing interest rates on equally risky taxable and tax-exempt bonds allows us to determine the effective tax rate in the economy.

For example, on Dec 31, 2021, bond yields were

	Characteristics		
	Rating	Duration	YTM
Tax-Exempt (Muni), VTEAX	AA	4.6 years	1.36%
Short-Term Inv-Grade, VFSUX	A-BBB	2.8 years	1.81%
(interpolated)	AA	4.6 years	≈ 2.20%
Intermed -Term Inv-Grade, VFSUX	BBB	6.5 years	2.64%

I had to interpolate what a modestly higher investment-grade bond fund with a 4.6 year duration would have provided — about 2.20% seems right. Would tax-exempt municipal or taxable bonds have been better for *you*? Well, it depends. If you had invested \$100 into munis at a 1.36% interest rate, you would have receiveed \$1.36 of interest at year's end and Uncle Sam would have gotten none of it. If you had invested \$100 in short-term investment grade corporate bonds of similar maturity and credit quality, you would have received (about) \$2.20. If your federal income tax rate was 0%, you would have preferred the \$2.20 over the \$1.36. However, if your marginal tax rate was 40%, Uncle Sam would have collected $\$2.20 \cdot 40\% \approx \$0.84$ in interest taxes and left you with $\$2.20 \cdot (1 - 40\%) \approx \$1.26$. Because \$1.36 is more than \$1.26, you would have preferred the tax-exempt muni bonds.

In economics, almost everything that is important is "on the margin." Thus, economists like to think about a hypothetical marginal investor. This is an investor whose marginal income tax rate is such that she would be exactly indifferent between buying the tax-exempt bond and the taxable bond. Using the same calculations, this marginal investor has a tax rate of

(Omitted eq)

Any investor with a marginal income tax rate above 38% (such as a high-income retail investor in a high-tax state) should have preferred the tax-exempt bond. Any investor with a marginal income tax rate below 38% (such as a tax-exempt endowment or pension plan) should have preferred the taxable bond.

(Omitted solvenow)

Taxes in Net Present Values

Again, as with transaction costs, you should take care to work only with cash in the same units — here, this means cash that you can use for consumption. Again, it should not matter whether you receive $100 that has to be taxed at 50% or whether you receive $50 that does not have to be taxed. As far as NPV is concerned, you should compute everything in after-tax dollars. This includes all cash flows, whether they occur today or tomorrow, and whether they are inflows or outflows.

Begin Important

Perform all NPV calculations in *after-tax* money. This applies both to the expected cash flows and to the opportunity cost of capital. **End Important**

Unfortunately, you cannot simply discount before-tax cash flows with the before-tax cost of capital (wrong!) and expect to come up with the same result as when you discount after-tax cash flows with the after-tax cost of capital (right!).

For example, consider a project that costs $10,000 and returns $13,000 next year. Your tax rate is 40%, and 1-year equivalently risky bonds return 25% if their income is taxable and 10% if their income is not taxable. First, you must decide what your opportunity cost of capital is. Section **??** showed that if you invest $100 into taxables, you will receive $125 but the IRS will confiscate ($125 – $100) · 40% = $10. You will thus own $115 in after-tax wealth. Tax-exempts grow only to $110, so you prefer the taxable bond — it is the taxable equally risky bond that determines your opportunity cost of capital. Your equivalent after-tax rate of return is therefore 15%. This 15% is your after-tax "opportunity" cost of capital — it is your best alternative use of capital elsewhere.

Return to your $10,000 project now. You know that your taxable project returns 30% taxable ($3,000), while taxable bonds return 25% ($2,500), so NPV should tell you to take this project. Uncle Sam will confiscate 40% · $3,000 = $1,200, leaving you with $11,800. Therefore, the NPV of your project is

(Omitted eq)

It makes intuitive sense: If you had invested money into the bonds, you would have ended up with $11,500. Instead, you will end up with $11,800 — the $300 difference occurring next year. Discounted, the $261 seems intuitively correct. Of course, there are an infinite number of ways of getting *incorrect* solutions, but let me point out a few. None of the following calculations that use the before-tax expected cash flows (and try different discount rates) give the same correct result of $260.87:

(Omitted eq)

You have no choice: To get the correct answer of $260.87, *you cannot work with before-tax expected cash flows*. Instead, you need to go through the exercise of carefully computing after-tax cash flows and discounting with your after-tax opportunity cost of capital.

You know that computing after-tax cash flows is a pain. Can you at least compare two *equally* taxable projects in terms of their before-tax NPV? If one project is better than the other in before-tax terms, is it also better in after-tax terms? If yes, then you could at least do relative capital budgeting with before-tax project cash flows. This may or may not work, and here is why. Compare project SAFE, which costs $1,000 and will provide $1,500 this evening; and project UNSAFE, which costs $1,000 and will provide either $500 or $2,500 this evening with equal probability. The expected payout is the same, and the cost of capital is practically 0% for 1 day. If you are in the 20% marginal tax bracket, project SAFE will leave you with $500 in *taxable* earnings. The IRS will collect $20\% \cdot (\$1{,}500 - \$1{,}000) = \$100$, leaving you with +$400 in after-tax net return. Project UNSAFE will either give you $1,500 or –$500 in *taxable* earnings.

- If the project succeeds, you would send $\$1{,}500 \cdot 20\% = \$300$ to the IRS. If the project fails, and if you can use the losses to offset gains from projects elsewhere, you would send $\$500 \cdot 20\% = \$100$ *less* to the IRS (because your taxable profits elsewhere would be reduced). In this case, projects SAFE and UNSAFE would have the same expected tax costs and after-tax cash flows: $1/2 \cdot \$300 + 1/2 \cdot (-\$100) = \$100$.
- If you drop into a different tax bracket, say, 25%, when your (additional) net income is $1,000 higher, then project UNSAFE becomes less desirable than project SAFE. For the $1,500 income, the first $500 would still cost you $100 in tax, but the remaining $1,000 would cost you $250. Thus, your project's marginal tax obligation would be either $350 or –$100, for an expected tax burden of $125. (The same logic applies if your losses would make you fall into a lower tax bracket — the UNSAFE project would become less desirable, because the tax reduction would be worth less.)
- If you have no capital gains elsewhere that you can reduce with the UNSAFE project capital loss, then the UNSAFE project would again be worth less. Corporations can ask for a tax refund on old gains, so the unrealized tax loss factor is less binding than it is for individuals, who may have to carry the capital loss forward until they have sufficient income again to use it — if ever.

Thus, whether you can compare projects on a before-tax basis depends on whether you have perfect symmetry in the applicable marginal tax rates across projects. If you do, then the project that is more profitable in after-tax terms is also more profitable in before-tax terms. This would allow you to simply compare projects by their before-tax NPVs. If gains and losses face different taxation — either because of tax bracket changes or because of your inability to use the tax losses elsewhere — then you cannot simply choose the project with the higher before-tax NPV. You will have to go through the entire after-tax NPV calculations and compare them.

Begin Important

You can only compare projects on a before-tax NPV basis if the tax treatment is absolutely symmetric. This requires consideration of your overall tax situation. **End Important**

You now know how to discount projects in the presence of income taxes. However, you do not yet know how to compute the proper discount rate for projects that are financed by debt and equity, because debt and equity face different tax consequences. Unfortunately, you will have to wait until Chapter before we can do a good job discussing the two suitable methods — called adjusted present value (APV) and the weighted average cost of capital (WACC) — to handle differential taxation for different corporate securities.

(Omitted solvenow)

Tax Timing

In many situations, the IRS does not allow reinvestment of funds generated by a project without an interim tax penalty. This can be important when you compare one long-term investment to multiple short-term investments that are otherwise identical. For example, consider a farmer in the 40% tax bracket who buys grain (seed) that costs \$300 and that triples in value every year.

- If the IRS considers this farm to be *one long-term two-year project*, the farmer can use the first harvest to reseed, so \$300 seed turns into \$900 in one year and then into a \$2,700 harvest in two years. Uncle Sam considers the profit to be \$2,400 and so collects taxes of \$960. The farmer is left with an after-tax cash flow of $\$2{,}700 - \$960 = \$1{,}740$.
- If the IRS considers this production to be *two consecutive 1-year projects*, then the farmer's after-tax profits are lower. He ends up with \$900 at the end of the first year. Uncle Sam collects $40\% \cdot (\$900 - \$300) = \$240$, leaving the farmer with \$660. Replanted, the \$660 grows to \$1,980, of which the IRS collects another $40\% \cdot (\$1{,}980 - \$660) = \$528$. The farmer is left with an after-tax cash flow of $\$1{,}980 - \$528 = \$1{,}452$.

The discrepancy between \$1,740 and \$1,452 is due to the fact that the long-term project can avoid the interim taxation. Similar issues arise whenever an expense can be reclassified from "reinvested profits" (taxed, if not with some credit at reinvestment time) into "necessary maintenance."

(Omitted solvenow)

11.5 Entrepreneurial Finance

Now that you understand how to work with market imperfections, for what types of firms do they matter most? Market imperfections are probably mild for large, publicly traded corporations. These types of firms typically face only modest interest rate spreads between their (risky) borrowing and lending rates. Of course, their *promised* borrowing interest rates are a little higher than what they can receive investing their money in Treasury bonds. Yet, given that they still have some possibility of going bankrupt, large firms' required *expected* borrowing costs of capital are probably fairly close to the *expected* rates of return they could earn if they invested in bonds with characteristics similar to the bonds that they themselves have issued. Thus, large public corporations can often pretend to live in a reasonably perfect market. This also means that they have the luxury of separating their project choices from their financial needs.2]subsect:creditratingsAltman study of bond default rates

In the world of individuals, entrepreneurs, and small companies, however, it is quite plausible that the costs of capital are often higher than equivalent expected savings interest rates. In fact, the most important difference between "ordinary corporate finance" and "entrepreneurial finance" is the degree to which their capital markets are perfect. Almost all entrepreneurs find it very difficult to convey credibly their intent and ability to pay back loans. And any credit that entrepreneurs receive is usually also very illiquid: Lenders cannot easily convert it into cash, should the need arise. Therefore, they demand a high liquidity spread, too. Many entrepreneurs even end up having to resort to financing projects with credit cards, which may charge 1,000 basis points or more above Treasury.

In sum, small firms often face extraordinarily high differentials between expected borrowing and lending rates. Entrepreneurs' high borrowing costs can thus prevent them from taking many projects that they would have undertaken if they had the money already on hand. Cash-on-hand can become a prime determinant of all their decisions.]subsect:separation-investment-consumptionSeparation of Decisions More established firms or wealthier entrepreneurs should optimally take more projects than poorer entrepreneurs. Yes, the world is not fair. Fortunately, micro-credit programs in many poorer countries have become more competitive over the years, helping the poorest of the poor to make investments that they could otherwise not and that they often can pay back fairly quickly. Making financial markets less imperfect for them is often not only a profitably activity but also a social blessing.

However, be careful in the real world before you believe the claims of entrepreneurs. Entrepreneurs also tend to have notoriously overoptimistic views of their prospects. Even venture capital — the financing vehicle for many high-tech entrepreneurial ventures — may advertise rates of return of 30% per year or more, but they seem to have managed to return only a couple of percentage points above the risk-free rate over the last 30 years *on average*. Adjusting for the correct default rates may actually mean that entrepreneurs face only high *promised* borrowing costs, not high *expected* borrowing costs. Thus, the large quoted spread between entrepreneurs' borrowing and lending rates, which is really all that you can easily observe, likely has a large component that is due not to information disagreements but simply to credit risk.

This issue of how to deal with market imperfections for small firms also arises frequently in the courts, where a cost-of-capital estimate is necessary to compute the value for an entrepreneurial enterprise — for example, for purposes of assessing the inheritance tax or resolving disputes among former business partners. (Such valuation services are an important revenue business for many finance professors and consulting firms.) It has become customary and court-sanctioned to compute first the value of an equivalent publicly traded business or company as if it faced a perfect market, and then to apply a private discount of around 10% to 30% to this hypothetical private firm value in order to reflect its limited access to capital. The amount of this discount is ad hoc, but it is better than no attempt at all.

(Omitted solvenow)

11.6 Deconstructing Quoted Rates of Return

In Sections]subsect:deconstructdefaultpremiumDefault Premium Deconstruction and]subsect:deconstructriskpremiumRisk Premium Deconstruction, you learned that you could decompose quoted rates of return into a term premium, a default premium, and a risk premium. Market imperfections can create additional premiums.

(Omitted eq)

and as always

(Omitted eq)

Quantifying imperfect market premiums is not easy, but we will try anyway. Unfortunately, there is not much that can be said about one of the imperfect market premiums — the premium compensating for differences in opinions. The nature of information disagreements is that they are idiosyncratic. But this does not mean that they are unimportant. As noted earlier, imperfections can be so large, even in financial markets, that they may destroy a financial market's viability altogether. Fortunately, the other three imperfections — taxes, transaction costs, and shallow markets — create premiums that are sometimes a little easier to quantify than the premium associated with information disagreements.

Tax differences are often modest across assets in the *same* class. However, when there are assets that are treated differently from a tax perspective, the one with the worse treatment has to offer a higher rate of return. For example, U.S. Treasury bonds are exempted from state and local income taxes(just as state and local bonds are exempt from U.S. income tax). This means that corporate bonds need to pay a higher yield than Treasury bonds to make up for state income taxes — a tax premium.

Transaction costs and deep competitive markets also play important perfect capital market roles. The resulting premiums when they are not satisfied are often lumped under the more general name "liquidity premiums." The idea is that when given a choice between a very liquid security (that you can resell in an instant to many different investors in case you need money) and a very illiquid security, you will demand an extra rate of return to buy the less liquid one. We can thus extend our earlier premiums analysis to the following:

(Omitted eq)

Again, there could be other premiums that should go into this formula, such as information premiums or bond contract feature premiums. I omit them because I don't have empirical evidence to show you. In addition, our concept of a clean decomposition is a little problematic in itself, because these premiums overlap. For example, it is quite possible that there are covariance-risk aspects to liquidity. (In other words, it could be that liquidity spreads increase when the market goes down, which could demand an extra risk premium.) Thus, a part of the quoted spread could be considered either as a risk premium or as a liquidity premium. Nevertheless, the basic decomposition in the above formulas is useful.

Let's go back to corporate bonds. You already learned in Section 2]subsect:creditratingsCredit Ratings that many corporate bonds have significant default risk, which means that they have to offer a default premium (relative to Treasuries, of course). Let me now tell you that, depending on credit rating, they have market betas between about

0.1 (investment-grade bonds) and 0.5 (junk bonds). This means that junk bonds may have to offer meaningfully large premiums to compensate investors for market risk, but for investment-grade bonds, even if one applied the CAPM, any inferred beta premium would be trivial.

However, many corporate bonds are difficult to resell quickly — most have to be traded over-the-counter, and not on an organized exchange. Therefore, they have to offer their buyers a liquidity premium. Finally, unlike U.S. Treasuries, corporate bonds are subject to state income taxes. This means that they have to offer a tax premium.

(Omitted solvenow)

11.7 How to Work Novel Composite Problems

Of course, in the messy real world, you can suffer many problems (such as inflation, transaction costs, disagreements, sole potential buyers, and taxes) all at once, not just in isolation. In fact, there are so many possible real-world problems that no one can possibly give you a formula for each one. 0.5]sect:apples-nominal-inflationInflation Thus, it is important that you approach the real world keeping a multitude of issues in mind.

1. Ask yourself in a given situation whether the assumption of a perfect market is reasonably appropriate. For example, in the case of large and possibly tax-exempt companies, you may consider it reasonable to get away with assuming a perfect market, and just work out the "perfect market" answer — a simple NPV, for example. Then think about the direction in which market imperfections would push you, judge the magnitude, and make an intuitive adjustment. You can thereby often work out a good answer without the enormous complications that the perfectly correct answer would require.
2. If you conclude that you are a long way from home (i.e., from a perfect market), then you must first determine which market imperfections are most important. Then you must work out a good solution by yourself. If you had hoped for the one magic bullet that tells you how to solve every different kind of problem you might encounter, I have to disappoint you. There are just too many possibilities, and the task is often hard. Probably the best way to answer such new and thorny questions is to internalize the method of "thinking by numerical example." You really must be able to work out formulas for yourself when you need them.

Solving a Problem with Inflation and Taxes

For example, let's see how you could approach a situation with both taxes and inflation. Always start by making up some numbers you find easy to work with. Let's say you are considering an investment of \$100. Further, assume that you will earn a 10% rate of return on your \$100 investment and Uncle Sam will take $\tau = 40\%$ (or \$4 on your \$10 return). Therefore, you get \$110 before taxes but end up with only \$106 in nominal terms. What you have just calculated is

(Omitted eq)

Now you need to determine what your \$106 is really worth, so you must introduce inflation. Pick some round number, say, a rate of $\pi = 5\%$ per annum. Consequently, in purchasing power, the \$106 is worth:

(Omitted eq)

Your after-tax, post-inflation, real rate of return is $\$100.95/\$100 - 1 = 0.95\%$. Knowing the numerical result, you need to translate your numbers into a formula. You computed

(Omitted eq)

This is, of course, not a formula that anyone remembers. However, it is a useful illustration of how you should approach and simplify complex questions — numerical example first, formula second.

➤ Taxes on Nominal Returns?

Here is an interesting question: If the real rate of return remains constant, does it help or hurt an investor if inflation goes up? Let's assume that the real rate of return is a constant 20%. If inflation is 50%, then the nominal rate of return is 80% (because $(1 + 50\%) \cdot (1 + 20\%) = 1 + 80\%$): You get \$180 for a \$100 investment. Now add income taxes to the tune of 40%. The IRS sees \$80 in interest, taxes \$32, and leaves you with \$48. Your \$148 will thus be worth $\$148/(1 + 50\%) \approx \$98.67$ in real value. Instead of a 20% increase in real purchasing power when you save money, you now suffer a $\$98.67/\$100 - 1 \approx -1.3\%$ change in real purchasing power. Despite a high real interest rate, Uncle Sam ended up with more, and you ended up with less purchasing power than you started with. The reason is that although Uncle Sam claims to tax only interest gains, you can actually lose in *real* terms because the interest tax is on *nominal* interest payments. Contrast this with the same scenario without inflation. In this case, if the real rate of return were still 20%, you would have earned \$20, Uncle Sam would have taxed you \$8, and you could have kept \$112 in real value.

Begin Important

If real before-tax interest rates remain constant, because the IRS taxes nominal returns, not real returns, you get the following results:

- Higher inflation and interest rates hurt *taxable* savers.
- Higher inflation and interest rates help *taxable* borrowers.

(Economic forces of demand and supply for capital may therefore have to adjust, so that real rates of return increase when inflation increases.) **End Important**

For much of postwar U.S. history, real rates of return on short-term government bonds have indeed been *negative* for taxed investors.

(Omitted solvenow)

Summary

This chapter covered the following major points:

- If markets are perfect, there are infinitely many buyers and sellers, no disagreements (opinions), no transaction costs, and no taxes.
- In perfect markets, *promised* borrowing and lending rates can be different, but *expected* borrowing and lending rates cannot. In imperfect markets, even *expected* borrowing and lending rates can be different.
- If markets are not perfect, capital budgeting decisions can then depend on the cash position of the project owner. NPV and interest rate computations can still be used, although you have to exert special care in working with correct and meaningful inputs (especially for the cost of capital). This is usually best done by thinking in terms of concrete examples first, then translating them into formulas later.
- Transaction costs can be direct (such as commissions) or indirect (such as searching or waiting costs). It is often useful to think of round-trip transaction costs.
- Financial assets' transaction costs tend to be very low, so that it is reasonable in many (but not all) circumstances just to ignore them.
- In the real world, buyers often prefer more liquid investments. To induce them to purchase a less liquid investment may require offering them some additional expected rate of return.
- Many financial markets have such low transaction costs and are often so liquid that they are believed to be close to perfect — there are so many buyers and so many sellers that it is unlikely that you would pay too much or too little for an asset. Such assets are likely to be worth what you pay for them.
- The tax code is complex. For the most part, individuals and corporations are taxed similarly. You must understand the following:
  - How income taxes are computed (the principles, not the details)
  - The fact that expenses that can be paid from before-tax income are better than expenses that must be paid from after-tax income
  - How to compute the average tax rate
  - How to obtain the marginal tax rate
  - That capital gains enjoy preferential tax treatment
  - Why the average and marginal tax rates differ, and why the marginal tax rate is usually higher than the average tax rate
- Taxable interest rates can be converted into equivalent tax-exempt interest rates, given the appropriate marginal tax rate.
- Tax-exempt bonds are usually advantageous for investors in high-income tax brackets. You can compute the critical tax rate for the investor who is indifferent between the two.
- You should do all NPV calculations with after-transaction-cost and after-tax cash flows and costs of capital.
- Long-term projects often suffer less interim taxation than short-term ones.
- Entrepreneurial finance can be viewed as the finance of imperfect markets. Small and startup firms suffer market imperfections more than large and established firms.

- Market imperfections are often responsible for large differences in required costs of capital. Limited diversification, liquidity, tax premia, etc., can be responsible for higher costs of capital for many projects. Their magnitude can be much higher than the CAPM-type risk premia that compensate investors for cash-flow covariance with the stock market.
- Quoted rates of return on financial instruments contain not only the term premium, default premium, and risk premium, but also many imperfect market premiums (such as tax premiums and liquidity premiums). For many bonds, the CAPM-style risk premium is very small compared to other premiums.
- The IRS taxes nominal returns, not real returns. This means that higher inflation rates are bad for savers and good for borrowers.
