8

Investor Choice: Risk and Reward

We are still after the same prize: a good estimate of the corporate cost of capital $E(r)$ in the NPV formula. But before you can understand the opportunity costs of capital for your firm's own projects, you have to understand your investors' other opportunities. This means that you must understand better what investors like (reward) and what they dislike (risk), how they measure their risks and rewards, how diversification works, what portfolios smart investors are likely to hold, and why it matters that "market beta" is a good measure of an investment asset's contribution to the market portfolio's risk.

8.1 Measuring Risk and Reward

Put yourself into the shoes of an investor and start with the most basic questions: How should you measure the risk and reward of your portfolio? As always, we first cook up a simple example and then generalize our insights into a broader real-world context. Say you are currently investing in an asset named *M*, short for "My Portfolio," but there are also other assets you could buy, named *A*, *B*, and *C*, plus a risk-free asset named *F*. These assets could themselves be portfolios, themselves consisting of many individual assets and/or yet other portfolios. (This is essentially what a mutual fund is.) So, let's just call M, A, B, C, and F themselves portfolios, too.

We will work with four equally likely scenarios, named S-1 through S-4, for each of the five portfolios. The outcomes, means, and risks are laid out in the table at the top of Figure **??**. Each scenario gets a card deck suit to remind you that it is a random draw. (If you find it easier to think in terms of historical outcomes, you can also pretend that you are analyzing historical data: scenario S-1 happened at time 1, S-2 at time 2, and so forth. This is not entirely correct, but it is often a helpful metaphor.)1]pg:N-1vsNWhy this is not entirely correct Which investment strategies do you deem better or worse, safer or riskier? If you can buy only these portfolios, what trade-offs of risk and reward are you facing?

(Omitted fig)

If you like visuals, Figure **??** also shows these returns in graphic form. The middle figure is the standard histogram, which you have seen many times elsewhere. However, each scenario is equally likely (the bars are equally tall), so it's more visually obvious to just put the card suit symbols where the bar is. This is what we do in the lower figure. It makes it easier to compare many different investments.

In this plot, you prefer assets that have scenario outcomes farther to the right (they have higher returns), outcomes that are *on average* farther to the right (they have higher *expected* rates of return), and outcomes that are more bunched]pg:what-is-a-histogramRandom variables *are* histograms together (they have less risk). Visual inspection confirms that investment F has outcomes perfectly bunched at the same spot, so it is not just least risky but in fact completely risk-free. It is followed by the risky M and A and B, and one much more risky investment, C.

Measuring Reward: The Expected Rate of Return

Although graphical measures are helpful, we really need formulas to give us numerical measures. A good measure for the reward is easy: You can use the expected rate of return, which is the probability-weighted average of all possible returns. For example, the mean rate of return for your portfolio M is

(Omitted eq)

If you invest in M, you would expect to earn a rate of return of 4%. Because each outcome is equally likely, you can compute this faster as a simple average,

(Omitted eq)

Measuring Risk: The Standard Deviation of the Rate of Return

A good measure of risk is less obvious than a good measure of reward, but fortunately you already learned a good measure — the standard deviation — in Section]subsect:uncertainty-variance-firsttimeThe standard deviation (measure of risk). Let's compute it in the context of our assets. We first write down how far away each point is from the center (average). You just saw that the average for M was +4%. An outcome of +3% would be closer to the mean than an outcome of –3%. The former is only 1 unit away from the mean. The latter is 7 units away from the mean.

	In S-1 (♣)	**In S-2 (♦)**	**In S-3 (♥)**	**In S-4 (♠)**
Asset M Rate of Return	–3%	+3%	+5%	+11%
…in deviation from its 4% mean	–7%	–1%	1%	+7%

Unfortunately, you cannot compute risk as the average deviation from the mean, which is always zero ($[-7 + (-1) + 1 + 7]/4 = 0$). You must first "neutralize" the sign, so that negative deviations count the same as positive deviations. The "fix" is to compute the average *squared* deviation from the mean (which turns out to be easier to work with than the absolute value). This is the variance:

(Omitted eq)

The variance has units that are intrinsically impossible to interpret by humans (% *squared* $= 0.01 \cdot 0.01$, written as x%%). Therefore, the variance carries very little intuition, except that *more variance means more risk*.

The measure that has more humanly meaningful (humane?) units is the standard deviation, which is just the square root of the variance:

(Omitted eq)

The standard deviation of the portfolio's rate of return is the most common measure of overall portfolio risk. Now look again at Figure **??**. You can see that this standard deviation of 5% seems like a reasonable measure of how far the typical outcome of M is away from the overall mean of M. (However, 5% is more than the average absolute deviation from the mean, which in this case would be 4%; the standard deviation puts more weight on far-away outcomes than the average absolute deviation.) The last column in Figure **??** lists the standard deviations of all investments. As the visuals indicate, F is risk-free; M, A, and B are equally risky at 5%; and C is riskiest at 9%.

Begin Important

- You can measure investment portfolio reward by the expected rate of return on the *overall* portfolio.
- You can measure investment portfolio risk by the standard deviation of the rate of return on the *overall* portfolio.

(Warning: You should not measure the investment risk *contributions* of individual assets *inside* a portfolio via their standard deviations. This will be explained in Section **??**.) **End Important**

At this point, you should begin to wonder how risk and reward are related in a reasonable world. This will be the topic of the next chapters, but the brief answer for now is that you can invest and speculate in dumb ways that give you high investment risk with low reward — as anyone who has gambled in a casino knows. However, if you are smart, after eliminating all investment mistakes (the low-hanging fruit), you have no choice but to take on more risk if you still want to earn yet greater expected rates of return (higher rewards).

(Omitted solvenow)

Nerdnote: It would be really convenient if we could compare gambles using a single measure. We could then easily compare them, like apples to apples. Fortunately, such a measure exists. It is called the "certainty equivalence." Unfortunately, it depends on a more complex model of the world, and it is notoriously difficult to get used to. Thus, we will cover it only in the companion Web chapter.

8.2 Diversification and Diversifiable Risk

In the real world, you are usually not constrained to buy assets in isolation — you can buy a little bit of many assets. This ability to buy many assets has the important consequence of allowing you to reduce your overall portfolio risk. Let's see why.

An Example Mixing Portfolio

Start again with your portfolio M. Now let's consider adding some of portfolio A. Why would you? It has the same risk and reward as M. However, although A has the same list of possible returns, *it offers them in different scenarios*. This rearrangement will make a lot of difference. So, let's say you have $100 in M, but you now sell half of these holdings to buy A. You will have $50 in M and $50 in A. Let's call this investment portfolio MA. In this case, your $100 investment would look like this:

	In S-1 (♣)	**In S-2 (♦)**	**In S-3 (♥)**	**In S-4 (♠)**	**Average**
Return on $50 in M:	$48.50	$51.50	$52.50	$55.50	$52.00
Return on $50 in A:	$51.50	$55.50	$48.50	$52.50	$52.00
⇒ Total return in MA:	$100.00	$107.00	$101.00	$108.00	$104.00
Rate of return in MA:	0%	7%	1%	8%	4%

You could have computed this more quickly by using the returns on M and A themselves. Your portfolio MA invests portfolio weight $w_M = 50\%$ into M and $w_A = 50\%$ in A. For example, to obtain the 7% in scenario S-2, you could have computed the portfolio rate of return from M's 3% rate of return and A's 11% rate as

(Omitted eq)

(Omitted fig)

Now let's look at these three portfolios (M, A, and MA) in a histogram. Even better, because our histogram bars are all equally tall, we can omit the bars and plot just the symbols. As Figure **??** shows, the range of M is from –3% to +11%; the standard deviation is 5%. The range of A is also from –3% to +11%; the standard deviation is also 5%. Yet the average of M and A has a much narrower range (0% to 8%) and a much lower standard deviation:

(Omitted eq)

MA is less risky than either of its ingredients.

The reason for this reduction in risk is diversification — the mixing of different investments within a portfolio that reduces the impact of each one on the overall portfolio performance. More simply put, diversification means that not all of your eggs are in the same basket. If one investment component goes down, the other investment component sometimes happens to go up, or vice-versa. The imperfect correlation ("non-synchronicity") reduces the overall portfolio risk.

(Omitted solvenow)

Random Diversification in 2021

Let's run a little experiment using real data. What would your realized risk have been if you had randomly invested into one vs. two (or four or more; call it "N") stocks at the start of 2020, all equal-weighted, for one year? I ran this experiment a million times and recorded the average realized return and risk each time.

What return would you have ended up with? Not surprisingly, the average of the average is the average — by randomly picking any stocks, you would have ended up with the equal-weighted average rate of return in the stock market, regardless of the number of stocks in your portfolio.

Random N:	1	2	4	8	16	32	64	128	256
Pfio Geo Return:	19%	19%	19%	19%	19%	19%	19%	19%	19%

The risk — measured as the time-series annualized standard deviation of these portfolios' monthly rates of return — was

Random N:	1	2	4	8	16	32	64	128	256
Pfio SD:	51%	41%	33%	27%	23%	20%	18%	17%	17%

Incidentally, even if you had invested in all ≈4,500 stocks, your portfolio time-series standard deviation would still not have been zero. It would have been lower — but still around 16-17%. Real-world diversification reduces risk, but it does not push it all the way down to zero.

The heterogeneity in the range of return outcomes you would have "enjoyed" would also have been higher the fewer stocks you would have invested in:

Random N:	1	2	4	8	16	32	64	128	256	All
Pfio Ret 25%	–20%	–14%	–6%	0%	5%	9%	12%	14%	16%	19%
Pfio Ret 75%	42%	40%	37%	32%	29%	27%	25%	24%	23%	19%

(This is sometimes called the "interquartile range," and your chance of lying inside or outside this range is the same 50-50.) You could have lost your shirt, even in this otherwise-great year for stocks!

So, if you are risk-averse, make sure to invest not just in one or two individual stocks but at least in a few dozen.

How Risk Grows With Time

Before we continue, I need to cover two aspects that fit more into the subfield of investments than into the subfield of corporate finance. But both are important for a general competence in finance (corporate financiers need to understand how investors think, and vice versa). We will look only at them in passing.

The first diversion is about how risk grows with time. Trust me on the following: If two random draws are independent, then the sum of these two random draws has a variance that is the sum of the two variances.

(Omitted eq)

(This is not true if the two variables move together!) Why should you care? Well, the rates of return of any one asset in a perfect market should be uncorrelated over time — if not, you could earn an extra rate of return by trading this

asset based on its own lagged return. (For example, if the correlation were positive, you would get rich quick by buying the asset *after* it has gone up and selling it *after* it has gone down.)

Now let's use an approximation: Ignore compounding. This means that the total return is approximately the sum of the consecutive returns. Now, if you expect a stock to earn a 10% mean expected rate of return with a standard deviation of 20% over one year ($20\% \cdot 20\% = 400\%\%$ variance), then over two years, you expect the same stock to earn 20% with a variance of $400\%\% + 400\%\% = 800\%\%$. Thus, this stock's risk (standard deviation) is $\sqrt{800\%\%} \approx 28.28\%$. In other words, its mean goes up by a factor of 2, but its risk goes up only by a factor of $\sqrt{2} \approx 1.4$.

Begin Important

Risk grows approximately with the square root of time. **End Important**

(Omitted solvenow)

The Best Mixing Portfolio — The Efficient Frontier

The second diversion is not just about how you calculate the risk and reward of a given portfolio, but how the *best* possible portfolios looks like. And how well can your best portfolios do? The details of this question are covered better in this chapter's appendix (in the companion Web chapter), but this section gives you a good though basic flavor.

(Omitted fig)

Figure **??** plots the investment performance (mean and standard deviation) of various portfolio combinations. Each portfolio has a unique spot in this coordinate system. This mean-variance graph is very common and familiar to all financiers. In such a plot, you want a portfolio that is higher on the Y-axis (has a higher expected rate of return) and lower on the X-axis (has a lower standard deviation). That is, you would always want to slide towards the north-west (up-left) if you can. One says a portfolio is *inside* the efficient frontier if it is south-east of another achievable portfolio, and *on* the frontier if there is no portfolio north-west

In the top-left plot, you can invest only in M and A. They are both at the same spot in the plot, though they are uncorrelated. Because both have a 4% mean rate of return, any combination of them does, too. The best (lowest risk for given mean) portfolio is the left-most one, which happens to be the equal-weighted combination.

The next plot is more interesting — B has a 64% correlation with M and a higher average rate of return than M. The bottom left plot is even more interesting. It allows you to invest in all three assets, M and A and B. You can see that B helps greatly, but not because you would buy it by itself. In fact, B itself is far inside the north-west boundary — the efficient frontier — which is obtained from the set of portfolios with the lowest-risk for any given level of reward, equivalently the highest reward for any given level of risk. (Its shape is always a hyperbola.) Presumably, smart investors would buy only portfolios on this efficient frontier. Anything inside (south-east) of the frontier is worse. Anything north-west of it is not obtainable. The equal-weighted portfolio is close to, but not on the efficient frontier. This is often the case for large diversified portfolios — in real life, the S&P 500 is reasonably close, but not exactly on the efficient frontier. The bottom-left plot then allows you to invest in C, too. You can see how this expands the efficient frontier even further. In fact, because C is so negatively correlated with the other three assets, it is now possible to create a risk-free asset with a rate of return of about 4.5% by cleverly combining investments. (Not that clever — invest about 37.7% in M, 26.1% in A, 9.1% in B, and 27.2% in C.) But even if you do not want to play it safe, you can always do at least as well with more assets than with fewer, so your efficient frontier has been pushed out further.

(Omitted fig)

Figure **??** shows an interesting aspect about the efficient frontier: If there is a risk-free asset, it becomes a straight line (from the Y-axis to the tangency portfolio of the risky assets)! What does this mean? A smart investor would only ever purchase a combination of this risk-free asset and this one tangency portfolio. To achieve a much higher expected rate of return, the investor would borrow to invest more than 100% of her portfolio in the risky tangency portfolio. For example, the best an investor could do while maintaining a 5% standard deviation was an expected return of about 6.5%. By borrowing, she can now achieve one of about 9%! Furthermore, if an investor wants to take more risk, she would prefer a lower risk-free rate. If she wants to take less risk, she would prefer a higher one. In the real world, investors cannot borrow an infinite amount to achieve an infinite expected rate of return — and certainly not at the same interest rate as the U.S. Treasury, which is the perfect market assumption to which we are still clinging (for the moment). As fascinating as the math about the efficient frontier is, we need to move on to our next topic: investor preferences.

(Omitted solvenow)

8.3 Investor Preferences and Risk Measures

You now understand that diversification can reduce risk. You still need to understand what projects the investors in your corporation — remember, this is *corporate finance* — would like *you* (their manager) to invest in on their behalves.

If Investors Care Only about Risk and Reward

Your intuition should now tell you that well-diversified portfolios — portfolios that invest in many different assets — tend to have lower risk. As a corporate manager, it would be reasonable for you to assume that your investors are smart. Because diversification helps investors reduce risk, you can also reasonably believe that they are indeed holding well-diversified portfolios. The most well-diversified portfolio may contain a little bit of every possible investable asset under the sun. Therefore, like most corporate executives, you would probably assume that your investors' portfolios are typically the overall market portfolio, consisting of all available investment opportunities.

Why would you even want to make any assumptions about your investors' portfolios? The answer is that if you are willing to assume that your investors are holding the market (or something very similar to it), your job as a corporate manager becomes much easier. Instead of asking each and every one of your investors what they might possibly like, you can just ask, "When would my investors want to give me their money for investment into my firm's project, given that my investors are currently already holding the broad overall stock-market portfolio?" The answer will be as follows:

1. Your investors should like projects that offer more reward — this means higher expected rates of return.
2. Your investors should like projects that help them diversify away some of the risk in the market portfolio, so that their *overall* portfolios end up being less risky. Be careful, though. This does not mean always going for the lowest-risk projects. Instead, this may well be searching out projects that behave very differently from other projects — unusual ones.

In sum, your corporate managerial task is to take those projects that your investors would like to add to their current (market) portfolios. You should therefore search for projects that have high expected rates of return and high diversification benefits with respect to the market. Let's now turn toward measuring this second characteristic: How can your projects aid your investors' diversification, and how should you measure how good this diversification is?

Begin Important

- Diversification is based on imperfect correlation, or "non-synchronicity," among investments. It helps smart investors reduce the overall portfolio risk.
- Therefore, as a corporate manager, in the absence of contradictory intelligence, you should assume that your investors tend to hold diversified portfolios. They could even hold portfolios as heavily diversified as the "entire market portfolio" — perhaps reasonably represented by something like `VFIAX` (`S&P 500`).
- As a corporate manager, your task is to think about how a little of your project can aid your investors in terms of its contribution to the risk and reward of their heavily diversified overall portfolios. (You should not think about how risky your project is on its own.)

End Important

If we are willing to assume that our smart investors are holding all assets in the market, then what projects offer them the best diversification?

Idiosyncratic Asset Risk and Risk Reduction

Obviously, diversification does *not* help if two investment opportunities always move in the same direction. For example, if you try to diversify one $50 investment in M with another $x investment in M (which always has the same outcomes), then your risk does not decrease. On the other hand, if two investment opportunities always move in *opposite* directions, then diversification works extremely well: One counterbalances the other.

Let's formalize this intuition. For explanation's sake, assume that "My Portfolio" M is also the market portfolio.

Go back to our assets B and C. Assume now that they are two projects that your firm could invest in, but you cannot choose both. Each project offers the same expected rate of return (6%), but B has lower risk (5%) than C (9%). As a manager, would you therefore assume that project B is better for your investors than project C?

(Omitted fig)

The answer is no. Let's assume that your investors start out with the market portfolio, M. Figure **??** shows what happens if they sell half of their portfolios to invest in either B or C. You can call these two "(50,50)" portfolios MB and MC, respectively. Start with MB. If your investors reallocate half their money from M into B, their portfolios would have the following rates of return:

	in S-1 (♣)	in S-2 (♦)	in S-3 (♥)	in S-4 (♠)	Reward	Risk
MB	1%	1%	6%	12%	5%	4.5%

The upper graph in Figure **??** plots the MB rates of return, plus the rates of return for both M and B by themselves. The averages are all close to both original rates of return. There is not much change in the risk of your portfolio in moving from a pure M portfolio to the MB portfolio. The risk shrinks slightly, from 5.0% to 4.5%.

Now consider the combination of MC, which is the lower graph in Figure **??**. By itself, C is a very risky investment (9% risk). It also has the single-worst outcome of any investment you have seen so far. However, if your investors instead reallocate half of their wealth from M into C, their overall portfolio would have the following rates of return:

	in S-1 (♣)	in S-2 (♦)	in S-3 (♥)	in S-4 (♠)	Reward	Risk
MC	7%	3%	8%	2%	5%	2.6%

The risk is much lower! Look again at Exhibit **??** — the MC outcomes are bunched much more closely than either M or C alone. And MB, too, has a much wider range than the MC portfolio. The MC combination portfolio is simply much safer — even though C by itself is much riskier.

In sum:

Portfolio	Reward	Risk	Note
M alone	4%	5.0%	Your investors' (market) portfolios
B alone	6%	5.0%	
C alone	6%	9.0%	C is riskier than B, if purchased by itself.
MB: half M, half B	5%	4.5%	Portfolio risk decreases less if B is added
MC: half M, half C	5%	2.6%	to M than if C is added to M!

You now know that C's own high standard deviation compared to B's is not a good indication of whether C helps your investors reduce portfolio risk more or less than B. It depends on what else your investors are holding:

- If your investors are primarily holding M, then a very risky project like C can allow them to build lower-risk portfolios.

- However, if your investors are not holding any other assets, they would not care about C's diversification benefits and only about your projects' own risk. B would be less risky for them.

Thus, as a manager, you could not determine whether your investors would prefer you to invest in B or C *unless* you knew the rest of their portfolios. (Moreover, it could also depend on how your investors would like you to trade off more overall reward against more overall risk.) To figure out how you can best help your "minions," you would have to guess what portfolios they actually do hold. (We will discuss an often reasonable and most common such guess — the market portfolio — soon.)

Begin Important

A project's (own) standard deviation is not necessarily a good measure of how it influences the risk of your investors' portfolios. Indeed, it is possible that a project with a very high standard deviation by itself may actually help lower an investor's overall portfolio risk. **End Important**

(Omitted solvenow)

(Market-) Beta and (Market-) Portfolio Risk Contribution

(Omitted fig)

Why is portfolio C so much better than portfolio B in reducing the overall risk when held in combination with the M portfolio? The reason is that C tends to go up when M tends to go down, and vice-versa. The same cannot be said for B — it tends to move together with M. You could call this "synchronicity" or "comovement." It is why B does not help investors who are heavily invested in the overall market in their quests to reduce their portfolio risks.

Figure **??** shows the comovement graphically. The rate of return on the market is on the X-axis; the rate of return on the asset is on the Y-axis. Its line slope in the plot is called the market beta. (It is common to write the formula for a line as $y = \alpha + \beta \cdot x$, which is where the Greek letter beta comes from.) A beta of 1 is a 45° diagonal line; a beta of 0 is a horizontal line. A positive beta slopes up; a negative beta slopes down. In statistics, you should have learned that you can find the beta by running a linear regression. If you don't remember, no worries: In Section **??**, I will teach you again how to compute the beta. For now, take my word that the two best-fitting lines are

(Omitted eq)

This formula is sometimes called the market model, because it measures an individual asset's returns relative to those of the entire market. The subscripts on the betas remind you what the variables on the X-axis and the Y-axis are. The first subscript is always the variable on the Y-axis, and the second is the variable on the X-axis. Thus, $\beta_{B,M} \approx 0.64$ and $\beta_{C,M} \approx -1.60$. Market beta plays such an important role in finance that the name "beta" has itself become synonymous for "market beta," and the second subscript is usually omitted.

In finance, we care about the market model line. As a corporate manager, you want to know how the rate of return on your own project comoves with that of the market. This is because you typically posit that your smart investors are on average holding the market portfolio. The best-fitting line between M and B slopes up. (It is also the same kind of line that you already saw in Section]sect:line-as-slopeMarket beta of Intel.) The positive slope means that B tends to be higher when M is higher. In contrast, the best-fitting line between M and C slopes down. The negative slope means that C tends to be lower when M is higher (and vice-versa). Again, this market slope is a common measure of expected comovement or countermovement — how much diversification benefit an investor can obtain from adding a particular new project to a well-diversified market-like investment portfolio. A higher slope means more comovement and less diversification; a lower, or even negative, slope means less comovement and more diversification.

Begin Important

- Diversification works better if the new investment project tends to move in the opposite direction from the rest of the portfolio than if it tends to move in the same direction.
- It is often reasonable to assume that smart investors are already holding the market portfolio and are now considering investing into just a little of one additional asset — your firm's new project.
- If this new investment asset has a negative beta with respect to the market (its "market beta"), it means that it tends to go down when the market goes up, and vice-versa. If this new investment asset has a positive beta with respect to the market, it means that it tends to move together with the market. If this new investment asset has a zero beta with respect to the market, it means that it moves independently of the market for all practical purposes.
- The market beta is a good measure of an investment asset's risk contribution for an investor who holds the market portfolio. The lower (or negative) the market beta, the more this investment helps reduce your investor's risk.
- The market beta of an asset can be interpreted as a line slope, where the rate of return on the market is on the x-axis and the rate of return on the new asset is on the y-axis. The line states how you expect the new asset to perform as a function of how the market will perform.
- You can think of market beta as a measure of "noxiousness." In a reasonable equilibrium, holding everything else constant, risk-averse investors who are holding the market portfolio would agree to pay more for assets that have lower market betas. They would pay less for assets with higher market betas.

End Important

Before we conclude, some caveats are in order. From your perspective as the manager of a company, perhaps a publicly traded company, it is reasonable to assume that your investors are holding the market portfolio. It is also reasonable to assume that your new project is just a tiny new additional component of your investors' overall portfolios. We will staunchly maintain these assumptions, but you should be aware that they may not always be appropriate. If your investors are *not* holding something close to the market portfolio, then your project's market beta would *not* be a good measure of your projects' risk contributions. In the extreme, if your investors are holding *only* your project, market beta would not measure the project's risk contribution at all. This is often the case for entrepreneurs. They often have no choice but to put all their money into one basket. Such investors should care only about their project's standard deviation, and not about the project's market beta.

When Beta? When Standard Deviation?

Do you care about your portfolio's beta or your portfolio's standard deviation? As CFO, do you care about your firm's beta or your firm's standard deviation? Make sure you understand the answers to these questions.

Begin Important

- As an investor, you usually care only about your portfolio's standard deviation (risk), and not about the risk of its individual ingredients.
- Typically, you do not care about the overall market beta of your portfolio. (The individual market betas can help you design your overall portfolio.)
- If you are the CFO of a firm that wants its shares to be purchased by investors that in turn want to hold the market portfolio, then you should care about your own firm's market beta. The lower your shares' market beta, the more these investors will like your shares.
- If you act purely in the interest of your diversified investors, you should not care about your firm's own standard deviation. Your investors can diversify away your firm's idiosyncratic risk. (If you care about your job or bonus, you might, however, take a different attitude towards risk. Corporate governance is the subject of companion Web chapter.)

End Important

Portfolio Alpha

Although we shall not use it further in this book, the alpha intercept in Formula **??** also plays an important role. Together, alpha and beta help determine how attractive an investment is. For example, if the rate of return on the market will be 10%, Formula **??** tells you that you would expect the rate of return on C to be

(Omitted eq)

The higher the alpha, the better the average performance of your investment given any particular rate of return on the market. Just as investment professionals often call the market beta just beta, they often call this specific intercept (here 12.4%) just alpha. (There is one small complication: They usually first subtract the risk-free interest rate from both r_C and r_M in their regressions — and this usually does not make much difference. We already mentioned the more serious real-world problem in Chapter : alphas are very volatile and difficult to predict.)

Computing Market Beta

So how can you actually compute beta? Let's first return to the assets in Figure **??**, where we knew the true outcome probabilities.3]fig:baseBase Investment Assets What is the market beta of C? I have already told you that this slope is –1.6. To calculate it, I followed a tedious, but not mysterious, recipe. Here is what you have to do:

1. Just as you did for your variance calculations, first translate all returns into deviations from the mean.]subsect:uncertainty-variance-firsttimeVariance calculations That is, for M and C, subtract their own means from every realization.

	In S-1 (♣)	**In S-2 (♦)**	**In S-3 (♥)**	**In S-4 (♠)**
Asset M Rate of Return	–3%	+3%	+5%	+11%
…in deviation from 4% mean	–7%	–1%	+1%	+7%
Asset C Rate of Return	+17%	+3%	+11%	–7%
…in deviation from 6% mean	+11%	–3%	+5%	–13%

2.]eq:varMVariance of MCompute the variance of the series on the X-axis. This is the variance of the rates of return on M. You have already done this in Formula **??**: $\text{Var}(r_M) = 25\%\%$.

3. Now compute the probability-weighted average of the products of the two net-of-mean variables. In this case,

   (Omitted eq)

   where XR is the return net of its mean. This statistic is called the covariance, here between the rates of return on M and C.

4. The beta of C with respect to the market M, formally $\beta_{C,M}$ but often abbreviated as β_C, is the ratio of these two quantities,

   (Omitted eq)

This slope of –1.6 is exactly the market beta we drew in Figure **??**. Many spreadsheets and all statistical programs can compute it for you: They call the routine that does this a linear regression.

You should always think of an asset's beta with respect to a portfolio as a characteristic measure of your asset relative to that portfolio. The rate of return on portfolio P is on the X-axis; the rate of return on asset i is on the Y-axis. As we stated earlier, most often — but not always — the portfolio P is the market portfolio, M, so $\beta_{i,M}$ is often just called the market beta of i, or just the beta of i (and the second subscript is omitted).

Now think for a moment. What is the average beta of a stock in the economy? Without going into proofing details (trust me!), it turns out mathematically that the question turns into asking what the value-weighted beta of the market portfolio is. Call the market-cap weight of each stock i to be w_i, and replace C in Formula **??** with M:

(Omitted eq)

This is because, if you look at the definition of covariance, you can see that the covariance of a variable with itself *is* the variance. (The covariance is a generalization of the variance concept from one to two variables.) Therefore,

$\text{Cov}(r_M, r_M) = \text{Var}(r_M)$, and the market beta of the market itself is 1. Graphically, if both the X-axis and the Y-axis are plotting the same values, every point must lie on the diagonal. Economically, this should not be surprising, either: the market goes up one-to-one with the market.

Begin Important

The (value-weighted) average beta of all stocks in the market is 1 by definition. **End Important**

Now that you know how to compute betas and covariances, you can consider scenarios for your project. For example, you might have a new project for which you would guess that it will have a rate of return of –5% if the market returns –10%; a rate of return of +5% if the market returns +5%; and a rate of return of 30% if the market returns 10%. Once you have come up with such scenario analysis, can you then estimate a market-beta? Of course. (You can also use this technique to explore the relationship between your projects and some other factors. For example, you could determine how your projects covary with the price of oil to learn about your project's oil risk exposure.)

Real-World Market Beta Estimation

In the real world, what is the best way to obtain an estimate of market-beta? For some project, you will have to think in terms of such long-outcome scenarios. However, for others, you may have observed historical rates of return, so you can use the historical stock-market returns and those on your own project (or similar project). Fortunately, as we noted upfront, the beta computations themselves are exactly the same. In effect, when you use historical data, you simply assume that each time period was one representative scenario and proceed from there. Nevertheless, there are some real-world complications you should think about:

1. Should you use daily, weekly, monthly, or annual rates of return? The answer is that the best market beta estimates for publicly-traded companies come from daily data. Annual or even monthly data should be avoided (except in a textbook in which space is limited). Monthly data should be used only if need be.
2. How much data should you use? Most researchers tend to use three to five years of historical rate of return data. This reflects a trade-off between having enough data and not going too far back into ancient history, which may be less relevant. If you have daily data, 2-3 years works quite well. The minimum is 1 year, and more than 5 years is not useful.
3. Is the historical beta a good estimate of the future beta? It turns out that history can sometimes be deceptive, especially if your estimated historical beta is far away from the market's beta average of 1. You should run a regression with daily historical returns and "shrink" your historical beta toward the overall market beta of 1 (or below 1 if your firm is small). This is important. For example, in the simplest such shrinker, you would just compute an average of the overall market beta of 1 and your historical market beta estimate. If you computed a historical market beta of, say, 4 for your project, you should work with a prediction of future market beta of about $(4 + 1)/2 = 2.5$ for your project.

   Historical textbooks (including my own past editions) used to recommend averaging many industry projects instead of just your own. This seemed like a good idea at the time but empirical analysis shows that this was bad advice — it predicted very badly. For the most part, try to use your own historical daily returns, and not that of other firms in the industry.

Many executives start with a statistical beta estimated from historical data (or they just look up the statistical beta on a website, such as *YAHOO!* FINANCE [finance.yahoo.com]) and then use their intuitive judgment to adjust it. It is unlikely that such adjustments are any good. Even trained financial economists with years of experience calculating betas cannot do this well. The only modification which tends to work is shrinking towards 1.

(Omitted solvenow)

Why Not Correlation or Covariance?

There is a close family relationship between covariance, beta, and correlation. The beta is the covariance divided by one of the variances. The correlation is the covariance divided by both standard deviations. The denominators are always positive. Thus, if the covariance is positive, so are the beta and the correlation; if the covariance is negative, so are the beta and the correlation; and if the covariance is zero, so are the beta and the correlation. The nice thing about the correlation — which makes it useful in many contexts outside finance — is that it has no scale and is always between –100% and +100%:

- Two variables that always move perfectly in the same direction have a correlation of 100%.
- Two variables that always move perfectly in opposite directions have a correlation of –100%.
- Two variables that are independent have a correlation of 0%.

Such simplicity makes correlations very easy to interpret. The not-so-nice thing about correlation is that it has no scale and is always between –100% and +100%. Compare two investment assets:

G	–20%	+20%
G'	–0.2%	+0.2%
M	–10%	+10%

Say you have \$100 invested in M. If you replace \$1 of M with \$1 of G, your portfolio becomes more volatile — instead of \$90 or \$110, you now earn \$89.90 or \$110.10. If you do the same with G', it becomes less volatile — you earn \$90.10 or \$109.90. Market-beta appropriately reflects this: G has a market-beta of 2, while G' has a market-beta of 0.02. Correlation does not. All three asset returns have 100% correlation. Thus correlation ignores this scale difference between G and G', which disqualifies it as a serious candidate for a project risk contribution measure.

Spreadsheet Functions To Calculate Risk, Beta, and Reward

Doing all these calculations by hand is tedious. We computed these statistics within the context of just four scenarios, so that you would understand the meanings of the calculations better. However, you can do this a lot faster in the real world. Usually, you would download reams of real historical rates of return data into a computer spreadsheet, like Excel or OpenOffice. Spreadsheets have all the functions you need already built in — and you now understand what their functions actually calculate. In practice, you would use the following functions in Excel:

average computes the average (rate of return) over a range of cells. Also sometimes called the **mean**.

varp (or **var.p**) computes the (population) variance. If you worked with historical data instead of known scenarios, you would instead use the **var** (or **var.s**) function. (The latter divides by N – 1 rather than by N, which I will explain in a moment.)

stdevp (or **stdev.p**) computes the (population) standard deviation. If you used historical data instead of known scenarios, you would instead use the **stdev** (or **stdev.s**) function (but this makes little difference if the number of observations is large).

covar computes the population covariance between two series. (If Excel were consistent, this function should be called covarp rather than covar.) Unlike the three preceding functions, this and the next two functions require two data cell ranges, not one.

correl computes the correlation between two series.

slope computes a beta. If *range-Y* contains the rates of return of an investment and *range-X* contains the rates of return on the market, then this function computes the market beta.

Some More Minor Statistical Nu(is)ances

In this chapter, we have continued to presume (just as we did in Section]subsect:history-repeatsWill history repeat itself?) that historical data gives us a reasonably good guide to the future when it comes to variances, covariances, and betas (assuming you calculate them well — 2-3 years of daily or weekly data, appropriately shrunk). Of course, this is a simplification — and remember that it can be a problematic one. I already noted that an equivalent historical representativeness assumption for means and alphas is *really* problematic. Rely on historical means as predictors of future expected rates of return only at your own risk!

There is a second, minor statistical issue of which you should be aware. Statisticians often use a covariance formula that divides by N – 1, not N. Strictly speaking, dividing by N – 1 is appropriate if you work with historical data. With a finite number of historical realizations, these are just sample draws and not the full population of possible outcomes. With a sample, you do not really know the true mean when you de-mean your observations. The division by a smaller number, N – 1, gives a larger but less biased covariance estimate to compensate for uncertainty about the true mean. It is also often called the *sample covariance*. In contrast, dividing by N is appropriate if you work with "scenarios" that you know to be true and equally likely. In this case, the statistic is often called the *population covariance*. The difference rarely matters in finance, where you usually have a lot of observations — except in our book examples where you have only four scenarios. (For example, dividing by N = 1,000 and by N = 1,001 gives almost the same number.)

The only reason why you even needed to know this distinction is that if you use a program that has a built-in variance or standard deviation function, you should not be surprised if you get numbers different from those that you have computed in this chapter. In some programs, you can get both functions. In Excel, you can use the *varp* and *stdevp* population statistical functions to get the population statistics, not the *var* and *stdev* functions that would give you the sample statistics.

Beta is not affected by whether you divide the variance/covariance by N or N – 1, because both numerator (covariance) and denominator (variance) are divided by the same number.

Furthermore, statisticians distinguish between underlying unknown statistics and statistics estimated from the data. For example, they might call the unknown true mean μ and the sample mean m (or $\bar{x}$). They might call the unknown true beta β^T and the estimated sample beta a beta with a little hat ($\hat{\beta}$). And so on. Our book is casual about the difference to reduce clutter, but keep in mind that whenever you work with historical data, you are really just working with sample estimates.

8.4 Interpreting Some Typical Stock-Market Betas

The market beta is the best measure of "diversification help" for an investor who holds (primarily) the stock market as her portfolio and now considers adding *just a little* of some firm's project. From your perspective as a manager offering such projects and seeking to attract investors, this is not a perfect, necessarily true assumption — but it is a reasonable one. Recall that we assume that investors are smart, so presumably they are holding highly diversified portfolios. To convince your market investors to like your $10 million project, you just need the average investor to want to buy $10 million divided by about $20 trillion (the stock-market capitalization), which is 1/2,000,000 of their portfolios. For your investors, your corporate projects are just tiny additions to their (likely) market portfolios.

(Omitted tbl)

Table **??** lists the market-betas of some stocks in December 2021. As we discussed in the previous chapter, like expected returns, expected market-betas are not perfectly known. Unlike expected returns, market-betas are reasonably stable (on average), which makes their estimation more reliable.

The table shows that the large Tech companies have market betas between about 0.9 and 1.2. `AMD` is smaller and covaries more with the market. The financials similarly covary about 1-to-1 with the market. The auto manufacturers now vary *more* with the market, while the two retail giants (Walmart and Target) vary *less* with it. The last stocks in the table were unusually popular with Robinhood investors in 2020. `RIOT` and `GME` are the only stocks where the prediction of market-beta for 2021 were far off — they covaried far more and far less with the stock-market than predicted. Both were more likely errant realizations (i.e., not expected, either before 2022 or now looking forward).

Most company betas are in the range of around 0 to about 2. As we already explained in the previous chapter, a market-beta above 1 is considered risk-increasing for an investor holding the overall stock market (it covaries more with the stock market itself), while a beta below 1 is considered risk-reducing. Negative betas would be great, but they are rare and usually temporary. (You can see some for the non-stock assets in Table .)

Market beta has yet another nice intuitive interpretation: It is the degree to which the firm's value tends to change if the stock-market changes. For example, if Tesla's 2021 market beta of approximately 1.5 continues to hold, it says that if the stock market will return an extra 10% next year (above and beyond its expectations), then Tesla's stock will likely return an extra $1.5 \cdot 10\% = 15\%$ (above and beyond Tesla's expectations). (Of course, Tesla's non-market risk is even higher. If you are planning to hold it, and if you are risk-averse rather than Musk-risk-seeking, make sure you are well diversified, so that it is only a small part of your portfolio!)

For now, let's say that the expected rate of return on the market is 6% and the expected rate of return on Tesla is 9%. (I cannot vouch for the two preceding estimates. They are speculative assumptions for the sake of the example.) Then, if the market were to turn in –4% (10% less than its expected return), you would expect Tesla to turn in

9% + 1.5 · (–10%) = –6%. Conversely, if the market were to turn in 16% (10% more than its expected return), you would expect Tesla to turn in 9% + 1.5 · (10%) = 24%. Tesla's high market beta is useful because it informs you that adding Tesla stock would not help you much with diversifying your portfolios risk *if* you mostly hold the overall stock market. Holding a little more Tesla would amplify any market swings into your portfolio.

But in any case, Tesla's market beta does not tell you whether Tesla is priced too high or too low on average, so that you should buy or avoid it in the first place. Market beta is not a measure of how good an investment Tesla is. (This would be the aforementioned alpha [which can be interpreted as an expected rate of return]. It is only a measure of how, on average, adding a little will increase or decrease your overall portfolio risk. In the next chapters, you will learn that the CAPM formula tries to relate to and adjust for market beta in its expected rate of return, giving you a commonly used benchmark for alpha.)

Betas have yet another common and important use. Let's say that you want to speculate only on Musk (and that Tesla will do *better* than what stock investors already believe), but you do not want to be exposed to market risk. The Tesla beta of 1.5 tells you that if you buy long $100 of Tesla stock and go short $150 in the stock market (which you can do easily, e.g., by shorting the VFIAX (S&P 500)) to create a "market-neutral" position, then your overall portfolio is not likely to be subject to market-wide swings. After all, for every $1 of general decrease (increase) in the overall stock market, Tesla goes up (down) on average by $1.50. Thus, the market-beta of 1.50 is also the hedge ratio that tells you how you can "immunize" a speculative stock position against market-wide changes.

(Omitted solvenow)

8.5 Market Betas for Portfolios and Conglomerates

Let's go back to your managerial perspective of figuring out the risk and return of your corporate projects. Many small projects are bundled together, so it is very common for managers to consider multiple projects already packaged together as one portfolio. For example, you can think of your firm as a collection of divisions that have been packaged together. If division B is worth $1 billion and division C is worth $2 billion, then a firm consisting of B and C is worth $3 billion. B constitutes $1/3$ of the portfolio "Firm" and C constitutes $2/3$ of the portfolio "Firm." This kind of portfolio is called a value-weighted portfolio because the weights correspond to the market values of the components. (A portfolio that invests $100 in B and $200 in C would also be value-weighted. A portfolio that invests equal amounts in the constituents — for example, $500 in each — is called an equal-weighted portfolio.)

Thus, as a manager, you have to know how to work with a portfolio (firm) when you have all the information about all of its underlying component stocks (projects). If I tell you the expected rate of return and market beta of each project, can you tell me what the overall expected rate of return and overall market beta of your firm are? Let's try it. Use the B and C stocks from Figure **??** on Page **??**, and call BCC the portfolio (or firm) that consists of $1/3$ investment in division B and $2/3$ investment in division C.

Actually, you already know that you can compute the returns in each scenario, and then the risk and reward.

	In S-1 (♣)	In S-2 (♦)	In S-3 (♥)	In S-4 (♠)	Reward E(r)	Variance Var(r)	Risk SD(r)
Investment B	5%	–1%	7%	13%	6%	25%%	5%
Investment C	17%	3%	11%	–7%	6%	81%%	9%
Portfolio BCC	13%	1.67%	9.67%	–0.33%	6%	≈30%%	≈5.5%

It is also intuitive that *expected* rates of return can be averaged. In our example, B has an *expected* rate of return of 6%, and C has an *expected* rate of return of 6%. Consequently, your overall firm BCC has an expected rate of return of 6%, too.

Unfortunately, you cannot compute value-weighted averages for all statistics. As the table shows, variances and standard deviations cannot be averaged ($1/3 \cdot 25\%\% + 2/3 \cdot 81\%\% \approx 62.3\%\%$, which is not the variance of 30%%; and $1/3 \cdot 5\% + 2/3 \cdot 9\% \approx 7.67\%$, which is not the standard deviation of 5.5%.)

But here is a remarkable and less intuitive fact: Market betas — that is, the projects' risk contributions to your investors' market portfolios — *can* be averaged! That is, I claim that the beta of BCC is the weighted average of the betas of B and C. In Formula **??**, you already computed the market-betas for as +0.64 and –1.60. So, their value-weighted average is

(Omitted eq)

You will be asked to confirm this conclusion in Q**??**.

Begin Important

- You can think of the firm as a weighted investment portfolio of components, such as individual divisions or projects. For example, if a firm named ab consists of only two divisions, a and b, then its rate of return is always

  (Omitted eq)

  where the weights are the relative values of the two divisions. (You can also think of this one firm as a "subportfolio" within a larger overall portfolio, such as the market portfolio.)
- The expected rate of return ("reward") of a portfolio is the weighted average expected rate of return of its components,

  (Omitted eq)

  Therefore, the expected rate of return of a firm is the weighted average rate of return of its divisions.
- Like expected rates of return, market betas can be weighted and averaged. The beta of a firm — i.e., the firm's "risk contribution" to the overall market portfolio — is the weighted average of the betas of its components,

  (Omitted eq)

  The market beta of a firm is the weighted average market beta of its divisions.

- You cannot do analogous weighted averaging with variances or standard deviations.

End Important

You can think of the firm not only as consisting of divisions, but also as consisting of debt and equity. For example, say your $400 million firm is financed with debt worth $100 million and equity worth $300 million. If you own all debt and equity, you own the firm. What is the market beta of your firm's assets? Well, the beta of your overall firm must be the weighted average beta of its debt and equity. If your $100 million in debt has a market beta of, say, 0.4 (debt usually moves less), and your $300 million of equity has a market beta of, say, 2.0, then your firm has a market beta of

(Omitted eq)

This 1.6 is called the asset beta to distinguish it from the equity beta of 2.0 that financial websites report. Put differently, if your firm refinances itself to 100% equity (i.e., $400 million worth of equity and $0 of debt), then the reported market beta of your equity on YAHOO!FINANCE would fall to 1.6. The asset beta is the measure of your firm's projects' risk contribution to the portfolio of your investors. It determines the cost of capital that you should use as the hurdle rate for projects that are similar to the average project in your own firm.

(Omitted solvenow)

Summary

This chapter covered the following major points:

- The expected (or mean) rate of return is a measure of expected reward. If scenarios are equally likely, then

  (Omitted eq)

- The variance is (roughly) the average squared deviation from the mean.

  (Omitted eq)

  If you work with known scenario probabilities, divide by N. If you work with a limited number of historical observations that you use to guesstimate the future scenarios, then divide by N – 1. (With a lot of historical data, N is very large and it really makes no difference what you divide by.) The variance is an intermediate input to the more interesting statistic, the standard deviation.

- The standard deviation is the square root of the variance. The standard deviation of a portfolio's rate of return is the common measure of its risk.

  (Omitted eq)

- Diversification reduces the risk of a portfolio.
- Corporate executives typically assume that their investors are smart enough to hold widely diversified portfolios, which resemble the overall market portfolio. The reason is that diversified portfolios offer lower risk than undiversified ones given the same expected rate of return.

- An individual project's own risk *is not* a good measure of its risk contribution to a smart diversified investor's portfolio.
- Market beta *is* a good measure of an individual asset's risk contribution for an investor who holds the market portfolio.
- Market betas for typical stocks range between 0 and 2.5. It's very rare that one would predict a beta beyond this range.
- It requires straightforward plugging of data into formulas to compute beta, correlation, and covariance. These three measures of comovement are closely related and always share the same sign.
- Like expected rates of return, betas can be averaged (using proper value-weighting, of course). However, variances or standard deviations cannot be averaged.

Preview of the Chapter Appendix in the Companion

The appendix to this chapter explains

- how risk and reward vary for different combination portfolios.
- how one can use the "matrix" of variances and covariances to quickly recompute the overall portfolio risk of different combinations.
- what optimal combination portfolios are. This is the efficient frontier (mean-variance efficiency or MVE), which you have already briefly encountered in this chapter. It is the cornerstone of modern investment theory.
- how the availability of a risk-free asset makes the optimal portfolio always a combination of this risk-free asset and some tangency portfolio. Thus, every rational investor would buy only these two assets. The more risk-averse, the more an investor would allocate from the risk-free into the risky tangency asset.
- how market beta coincidentally affects idiosyncratic risk, and how it influences market-conditional realized rates of return.

Compiled: Thursday 3rd April, 2025