

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 are likely to 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 through 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 $S1$ through $S4$, for each of the five portfolios. The outcomes, means, and risks are laid out in the tabular portfolio of Exhibit 8.1. 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 pretend that you are analyzing historical data: scenario $S1$ happened at time 1, $S2$ at time 2, and so forth. This is not entirely correct, but it is often a helpful metaphor.) 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?

If you like visuals, Exhibit 8.1 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 together (they have less risk). Visual inspection confirms that investment F has outcomes perfectly bunched at the same spot, so it is

We work with five assets that have four equally likely outcomes.

Historical samples can be viewed as scenarios.

► [Why this is not entirely correct.](#)
Pg.183.

Graphics version of the table.

In a histogram, bars to the right mean higher returns. Bars that are more spread out indicate higher risk.

► [Random variables are histograms.](#)
Pg.105.

	In S1 (♣)	In S2 (♦)	In S3 (♥)	In S4 (♠)	Reward $E(r)$	Variance ^a $Var(r)$	Risk $Sdv(r)$
Investment M	-3%	3%	5%	11%	4%	25%%	5%
Investment A	3%	11%	-3%	5%	4%	25%%	5%
Investment B	5%	-1%	7%	13%	6%	25%%	5%
Investment C	17%	3%	11%	-7%	6%	81%%	9%
Investment F	1%	1%	1%	1%	1%	0%%	0%

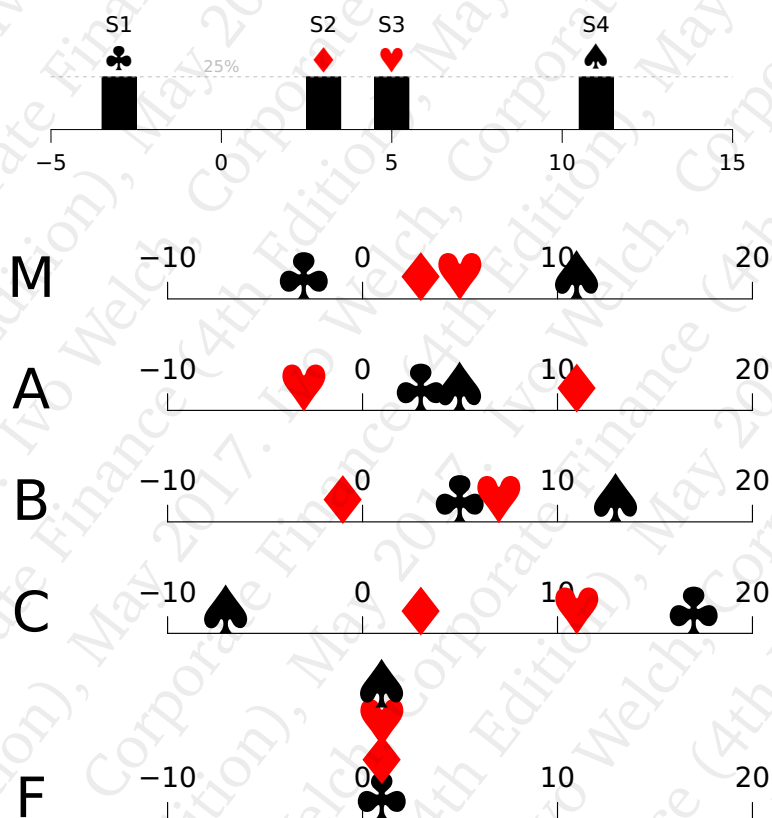


Exhibit 8.1: Rates of Return on Five Investment Assets. There are only four possible future scenarios, S1 through S4, each equally likely and indicated with a card suit. There are only 5 available investments (M, A, B, C, and F). (These could themselves be portfolios, of course.) The variance (Var) and standard deviation (Sdv) were explained in Section 6.1. The middle figure is a “traditional” histogram of M. The bottom figure contains the “condensed” histograms for all 5 assets.

Table note [a]: We use the ‘%%’ notation only for variance computations. Just like ‘%’ means ‘divide by 100’, ‘%%’ means ‘divide by 100 and then divide by 100 again’, i.e., ‘divide by 10,000’. This makes it easy to see that $\sqrt{(5\%)^2 + (5\%)^2 + (5\%)^2 + (5\%)^2}$ is $\sqrt{(25\% + 25\% + 25\% + 25\%) = 100\%} = 10\%$. If you find it easier to read $\sqrt{(0.0025 + 0.0025 + 0.0025 + 0.0025 = 0.01)} = 10\%$, then be my guest and use this notation instead. The answers are always the same.

not just least risky but in fact completely risk-free. It is followed by the risky M and A, then B, and finally, the riskiest 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

$$\begin{aligned} E(r_M) &= (1/4) \cdot (-3\%) + (1/4) \cdot (+3\%) + (1/4) \cdot (+5\%) + (1/4) \cdot (+11\%) = +4\% \\ &= \text{Sum of (each probability times its outcome)} \end{aligned}$$

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,

$$E(r_M) = [(-3\%) + (+3\%) + (+5\%) + (+11\%)]/4 = 4\%$$

Measure reward with the expected rate of return.

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 6.1. 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.

Measure risk with the standard deviation of the rate of return.

► [The standard deviation \(measure of risk\),](#)
Sect. 6.1, Pg.108.

	In S1 (♣)	In S2 (♦)	In S3 (♥)	In S4 (♠)
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. This is the **variance**:

The average deviation from the mean is always 0. It cannot measure risk.

$$\begin{aligned} \text{Var}(r_M) &= 1/4 \cdot (-3\% - 4\%)^2 + 1/4 \cdot (3\% - 4\%)^2 + 1/4 \cdot (5\% - 4\%)^2 + 1/4 \cdot (11\% - 4\%)^2 \\ &= [(-7\%)^2 + (-1\%)^2 + (+1\%)^2 + (+7\%)^2]/4 = 25\% \quad (8.1) \\ &= \text{Sum of (each probability times its squared-deviation-from-the-mean)} \end{aligned}$$

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

The standard deviation of the portfolio's rate of return is a common measure of risk.

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

$$\text{Sdv}(r_M) = \sqrt{\text{Var}(r_M)} = \sqrt{25\%} = 5\% \quad (8.2)$$

The standard deviation of the portfolio's rate of return is the most common measure of overall **portfolio risk**. Now look again at Exhibit 8.1. 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 Exhibit 8.1 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 a whopping 9%.

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 will not measure the investment risk *contributions* of individual assets *inside* a portfolio via their standard deviations. This will be explained in Section 8.3.)

A preview: Smart investors eliminate unnecessary risk. After they have done so, more reward requires taking more risk.

At this point, you should begin to wonder how risk and reward are related in a reasonable world. This will be the subject of much of the next chapter. The brief answer for now is that you can speculate in dumb ways that give you high investment risk with low reward—as anyone who has gambled 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 want to earn higher rewards.

Q 8.1. What happens if you compute the average deviation from the mean, rather than the average squared deviation from the mean?

Q 8.2. Asset M from Exhibit 8.1 offers -3% , $+3\%$, $+5\%$, and $+11\%$ with equal probabilities. Now add 5% to each of these returns. This new asset offers $+2\%$, $+8\%$, $+10\%$, and $+16\%$. Compute the expected rate of return, the variance, and the standard deviation of this new asset. How does it compare to the original M?

Q 8.3. Confirm the risk and reward of C in Exhibit 8.1.

Nerdnote: It would be really convenient if we could quote all gambles on the same terms. 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

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.

Many assets at the same time.

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:

Portfolios are bundles of multiple assets. Their returns can be averaged.

	In S1 (♣)	In S2 (♦)	In S3 (♥)	In S4 (♠)	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 S2, you could have computed the portfolio rate of return from M’s 3% rate of return and A’s 11% rate as

$$r_{MA} = r_{MA=50\% \text{ in M}, 50\% \text{ in A (all in S2)}} = 50\% \cdot 3\% + 50\% \cdot 11\% = 7\%$$

$$r_{MA=(w_M, w_A) \text{ in S2}} = w_M \cdot r_{M \text{ in S2}} + w_A \cdot r_{A \text{ in S2}}$$

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 Exhibit 8.2 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 lower range (0% to 8%) and a much lower standard deviation:

Visually, the M and A combination portfolio called MA has lower variability (risk and range) than either M or A.

$$\text{Var}_{50\% \text{ in M}, 50\% \text{ in A}} = \frac{(0\% - 4\%)^2 + (7\% - 4\%)^2 + (1\% - 4\%)^2 + (8\% - 4\%)^2}{4} = 12.5\%$$

$$= \frac{[r_{S1} - E(r)]^2 + [r_{S2} - E(r)]^2 + [r_{S3} - E(r)]^2 + [r_{S4} - E(r)]^2}{N}$$

$$\Rightarrow \text{Sdv}_{50\% \text{ in M}, 50\% \text{ in A}} = \sqrt{\text{Var}} = \sqrt{12.5\%} \approx 3.54\%$$

MA is simply 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.

This is caused by diversification.

	In S1 (♣)	In S2 (♦)	In S3 (♥)	In S4 (♠)	Reward E(r)	Variance Var(r)	Risk Sdv(r)
Investment M	-3%	3%	5%	11%	4%	25%	5%
Investment A	3%	11%	-3%	5%	4%	25%	5%
Portfolio MA	0%	7%	1%	8%	4%	12.5%	3.54%




Exhibit 8.2: Rate of return outcomes for M, A, and the (50%, 50%) combination portfolio MA. Because each half-M/half-A point is halfway between M and A, MA has lower spread (risk) than either of its components, M and A, by themselves.

Q 8.4. The combination portfolio named MA invests 90% in M and 10% in A.

1. Compute its risk and reward.
2. In a plot similar to those in Exhibit 8.1, would this new MA portfolio look less spread out than the MA = (50%, 50%) portfolio that was worked out in the table in Exhibit 8.2?

How Risk Grows With Time

Brief important diversions.

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. We will look only at them in passing.

If two variables are uncorrelated, the variance of the sum is the sum of the variances.

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 variables has a variance that is the sum of the two variances.

$$\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y) \quad \text{if } X \text{ and } Y \text{ are uncorrelated}$$

(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. (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 $\sqrt{2} \approx 1.4$.

Stocks have uncorrelated returns. Thus, with time, the risk grows more slowly than the reward.

Risk grows approximately with the squareroot of time.

IMPORTANT

Q 8.5. Please ignore compounding in this question:

1. What is the risk and reward of the "10% mean, 20% risk" investment that we just discussed in the text over 4 years? What is your reward-risk ratio? (This ratio is called the **Sharpe ratio** and often confusingly called a risk-reward ratio.)
2. What is the risk and reward of the same 10% mean, 20% risk investment over 9 years? What is the Sharpe ratio?
3. Can you guess what the risk and reward of a stock with annual mean E and risk S_{dv} are over T years? What is the Sharpe ratio?

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.

Finding the best choice.

Exhibit 8.3 plots the investment performance (mean and standard deviation) of various portfolio combinations. Each portfolio has a unique spot in this coordinate system. This plot 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

What is the typical mean-variance plot?

In the top-left plot, you can invest only in M and A. They are both at the same spot in the plot. 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 top-right plot allows you to invest not only in M and A, but in B, too. 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 the lowest-risk highest-reward set of portfolios. (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—the S&P 500 is reasonably close, but not exactly on the efficient frontier. The bottom-left plot allows you to invest in C, too. You can see how this expands the efficient frontier even further. In fact, 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 0.377 in M, 0.261 in A, 0.091 in B, and 0.272 in C.) But even if you do not want to play it safe, you can always do at least as well with more assets

More assets expand your opportunity set. The best investment choices are on the "efficient frontier."

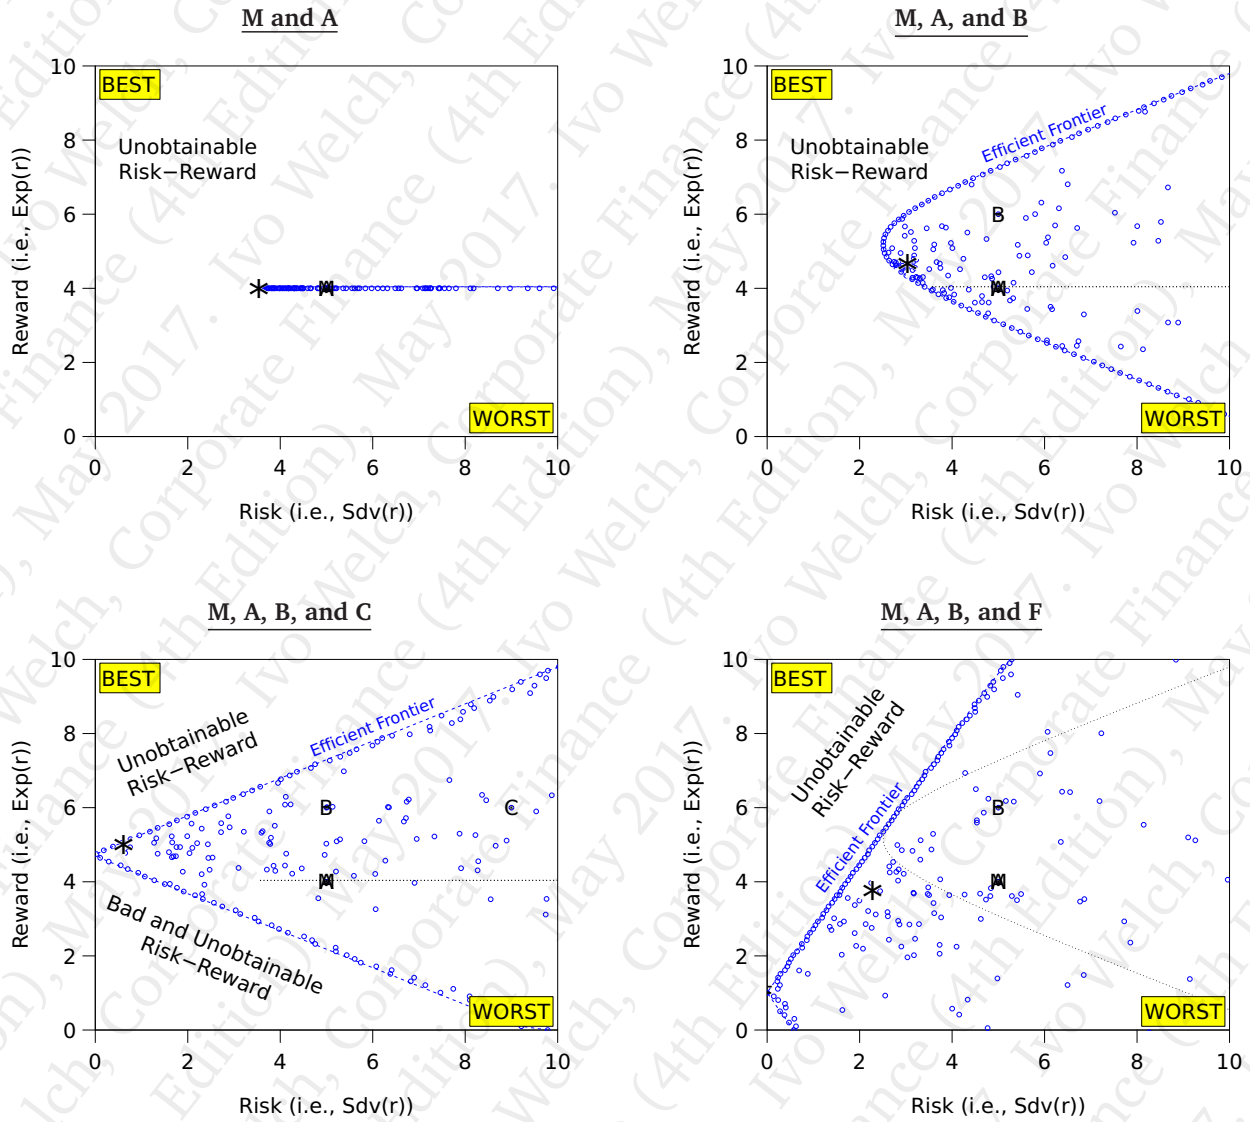


Exhibit 8.3: *The Efficient Frontier.* These plots show the mean and standard deviation of returns of portfolios composed of the stocks that are indicated in the header. (They appeared earlier, e.g., in Exhibit 8.1.) The “*” in the plots is the equal-weighted portfolio. The north-west border is the efficient frontier. (Note that the two lower plots even allow you to invest money and earn a risk-free rate. The lowest-risk portfolio is also called the **minimum-variance portfolio**.) The dotted line is the frontier from a previous plot.

than with fewer, so your efficient frontier has been pushed out further. The bottom-right plot shows your possible investments if instead of access to C, you had access to F. In both bottom figures, in which there is a risk-free asset, the efficient frontier is a line (the limit case for the hyperbola).

Q 8.6. How would the efficient frontier look if you were allowed to invest in all 5 assets, M, A, B, C, and F?

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 to invest in on their behalves.

The main question.

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 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.

Investors love diversification: the more the better. They could like the market portfolio because it is highly diversified.

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 what each and every one of your investors 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:

If your investors like high reward and low risk and hold the market portfolio, you can work out how your projects affect them.

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 going for 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?

- 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 believe that your investors tend to hold diversified portfolios. They could even hold portfolios as heavily diversified as the "entire market portfolio."

IMPORTANT

- 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 in itself.)

Assume that investors hold the overall market. Now what?

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?

Comovement determines risk contribution.

Idiosyncratic Asset Risk (Sdv) 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 \$50 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 is a buffer for the other.

Pretend M is not just "My portfolio," but the market.

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

Is B or C a better addition to your M portfolio?

Assume that B and C are two projects that your firm could invest in, but you cannot choose both. Both offer 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?

The combination MC has almost the same risk as M.

The answer is no. Let's assume that your investors start out with the market portfolio, M. Exhibit 8.4 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 S1 (♣)	in S2 (♦)	in S3 (♥)	in S4 (♠)	Reward	Risk
MB	1%	1%	6%	12%	5%	4.5%

The upper graph in Exhibit 8.4 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%.

The combination MC has much lower risk than M.

Now consider the combination of MC, which is the lower graph in Exhibit 8.4. 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 S1 (♣)	in S2 (♦)	in S3 (♥)	in S4 (♠)	Reward	Risk
MC	7%	3%	8%	2%	5%	2.6%

The risk is much lower! Look again at the 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:

	In S1 (♣)	In S2 (♦)	In S3 (♥)	In S4 (♠)	Reward $E(r)$	Variance $Var(r)$	Risk $Sdv(r)$
Investment M	-3%	3%	5%	11%	4%	25%%	5%
Investment B	5%	-1%	7%	13%	6%	25%%	5%
Investment C	17%	3%	11%	-7%	6%	81%%	9%
Portfolio MB	1%	1%	6%	12%	5%	20.5%%	4.5%
Portfolio MC	7%	3%	8%	2%	5%	6.5%%	2.6%




Exhibit 8.4: Combining the Market M with Either B or C. Although C is riskier than B by itself (look at C's one disaster outcome!), C is much better than B in reducing risk when it is added to the market portfolio M. This is because C tends to move opposite to M, especially if M turns in its worst outcome (-3%).

Portfolio	Reward	Risk	Note
M (=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%	4.2%	Portfolio risk decreases less if B is added to M than when C is added to M!
MC: half M, half C	5%	2.6%	

The implication for your project choices as a corporate manager: Everything else equal, C could better reduce portfolio risk for your investors despite its higher idiosyncratic risk.

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. 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 assets other than C, they would not care about C's diversification benefits and only about its own risk. Thus, as a manager, you cannot determine whether your investors would prefer you to invest in B or C unless you know their entire portfolios. (Moreover, it could also depend on how your investors would like you to trade off more overall reward against more overall risk.)

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.

Q 8.7. Confirm the risk and reward calculations for the MB and MC portfolios in the table in Exhibit 8.4.

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

C reduces M's risk because it tends to move in the opposite direction.

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.

Comovement can be measured by a line slope (beta). The market beta has the asset's rate of return on the y-axis and the market's rate of return on the x-axis.

Exhibit 8.5 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 8.3, I will teach you again how to compute the beta. For now, take my word that the two best-fitting lines are

$$r_B \approx 3.4\% + (+0.64) \cdot r_M \tag{8.3}$$

$$r_C \approx 12.4\% + (-1.60) \cdot r_M$$

$$r_i = \alpha_{i,M} + \beta_{i,M} \cdot r_M$$

This formula is sometimes called the **market model**. 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

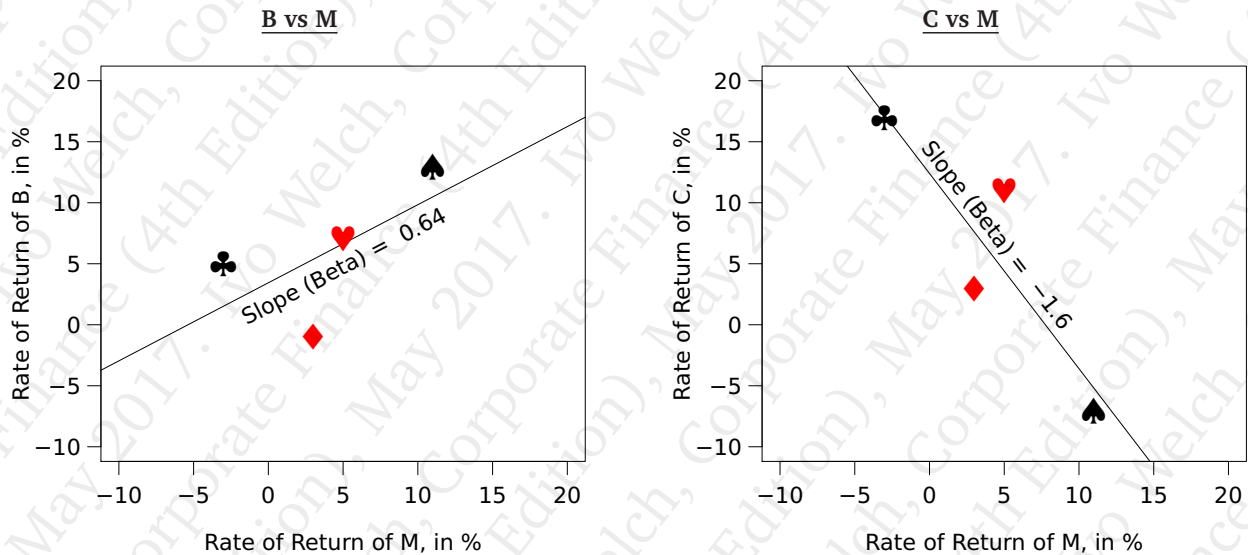


Exhibit 8.5: Possible Outcomes: Rates of Return of C and D versus Rate of Return of M. The four data points in each plot are taken from Exhibit 8.1 on page 166. They are the rates of return on the portfolios M, B, and C, quoted in percent. In the example, you know that these are the four true possible outcomes. In the real world, if the four points were not the true known outcomes, but just the historical outcomes (sample points), then the slope would not be the true unknown beta, but only the “estimated” beta.

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 7.1.) 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. A higher slope means more comovement and less diversification; a lower, or even negative, slope means less comovement and more diversification.

Market beta is a big deal in finance. It measures how your project covaries with the market.

Market beta of Intel,
Sect. 7.1, Pg.147.

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 “toxicity.” 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.

Warning: All of this beta-related risk measuring is interesting only if your investors are holding (portfolios close to) the overall market.

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.

- 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 to get into the market portfolio, so that investors willingly buy your shares, then you should care about your own firm's market beta.
- 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.)

IMPORTANT

Portfolio Alpha

Although we shall not use it further in this book, the alpha intercept in Formula 8.3 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 8.3 tells you that you would expect the rate of return on C to be

$$E(r_C | \text{if } r_M = 10\%) \approx 12.4\% + (-1.60) \cdot 10\% \approx -3.6\%$$

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 first both r_C and r_M in their regressions—and this usually does not make much difference.)

Alpha has meaning, too, even though you won't use it just yet.

Computing Market Betas from Historical Rates of Return

So how can you actually compute beta? Let's return to the assets in Exhibit 8.1. 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. That is, for M and C, subtract their own means from every realization.

	In S1 (♣)	In S2 (♦)	In S3 (♥)	In S4 (♠)
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. Compute 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 8.1: $\text{Var}(r_M) = 25\%$.

► [Base Investment Assets, Exhibit 8.1, Pg.166.](#)

You can compute the best-fit beta via a 4-step procedure. First, de-mean each rate of return. (How demeaning!)

► [Variance calculations, Sect. 6.1, Pg.108.](#)

► [Variance of M, Formula 8.1, Pg.167.](#)

For covariances, multiply net-of-mean returns, then average.

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

$$\begin{aligned}\text{Cov}(r_M, r_C) &= \frac{1}{4} \cdot (-7\%) \cdot (+11\%) + \frac{1}{4} \cdot (-1\%) \cdot (-3\%) \\ &\quad + \frac{1}{4} \cdot (+1\%) \cdot (+5\%) + \frac{1}{4} \cdot (+7\%) \cdot (-13\%) = -40\% \\ &= \text{Sum of (each probability times the returns' products)}\end{aligned}$$

This statistic is called the **covariance**, here between the rates of return on M and C.

The beta is the covariance divided by the variance.

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,

$$\begin{aligned}\beta_C &= \beta_{C,M} = \frac{-40\%}{25\%} \approx -1.6 \quad (8.4) \\ &= \frac{\text{Cov}(r_M, r_C)}{\text{Var}(r_M)}\end{aligned}$$

You can confirm our calculations using a spreadsheet.

Think of market beta as the characteristic of an asset.

This slope of -1.6 is exactly the market beta we drew in Exhibit 8.5. 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 an underlying base 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? Equivalently, what is the beta of the market portfolio? Replace C in Formula 8.4 with M:

$$\beta_M = \frac{\text{Cov}(r_M, r_M)}{\text{Var}(r_M)}$$

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.

IMPORTANT

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

Why torture you with computations? So you can play with scenarios.

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% . Knowing how to compute a market beta therefore makes it useful to think of such scenarios. (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, you will sometimes think in terms of such scenarios. However, you will more often have to compute a market beta from historical rates of return, using overall stock market returns and your own project (or similar project) returns. 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:

Practical advice to help you estimate market beta in the real world: Use 3-5 years of daily observations and then adjust.

1. Should you use daily, weekly, monthly, or annual rates of return? The answer is that the best market beta estimates come from daily (or weekly) data. Annual 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 simply 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. This seemed like a good idea but was bad advice. In practice, this approach has 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 \[finance.yahoo.com\]](http://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.

Q 8.8. Return to your computation of market beta of -1.6 in Formula 8.4. We called it $\beta_{C,M}$, or β_C for short. Is the order of the subscripts important? That is, is $\beta_{M,C}$ also -1.6 ?

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 $+100\%$:

Covariance and beta (and correlation) always have the same sign.

- 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 $+100\%$. This means that two investments, the second being a million times bigger than the first (all project rates of return multiplied by a million), have the same correlation with the stock market. Yet the second investment goes up or down with any slight tremor in the market by a million times more, which would of course mean that it would contribute much more risk. The correlation ignores this fact, which disqualifies it as a serious candidate for a project risk measure. Fortunately, beta takes care of scale—indeed, the beta for the second project would be a million times larger. This is why we prefer beta over correlation as a measure of risk contribution to a portfolio.

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.

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.

covar computes the population covariance between two series. (If Excel was consistent, this function should be called covarp rather than covar.) Unlike the earlier 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 Minor Statistical Nuances (Nuisances)

In this chapter, we have continued to presume (just as we did in Section 7.1) that historical data gives us a good guide to the future when it comes to means, variances, covariances, and betas (assuming you calculate them well—2 years of daily data, appropriately shrunk). Of course, this is a simplification—and remember that it can be a problematic one. I already noted that this is less of a problem for covariances, variances, and betas than it is for means. 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. 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 unbiased covariance estimate. 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,

In real life, you can do calculations faster with a spreadsheet.

Is history a good guide?

► Will history repeat itself?, Sect. 7.1, Pg.153.

When working with a sample, the (co)variance formula divides by $N-1$. When working with the population, the (co)variance formula divides by N .

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 due to lack of space, 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 the stock market portfolio and considers adding *just a little* of your firm’s project. From your perspective as a manager 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.

You can easily look up the market betas of publicly traded stocks on many financial websites. Exhibit 8.6 lists the betas of some randomly chosen companies in June 2016 from [YAHOO! FINANCE](#) and from Google’s finance site. Most company betas are in the range of around 0 to about 2. (American Airlines’ *historical* market beta was so high that Google even refused to admit to its own estimate.) A beta above 1 is considered risk-increasing for an investor holding the overall stock market (it is riskier than the stock market itself), while a beta below 1 is considered risk-reducing. Negative betas are rare and usually temporary. Gold is an asset that sometimes did and sometimes did not have a negative market beta (see [Barrick Gold](#) here). In this decade, long-term Treasury bonds had negative betas; but in past decades, it was positive. In almost all cases, it is better to estimate future market betas with firms’ own historical market betas (though shrunk) rather than their industry market betas.

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, AMD’s market beta of approximately 2 says that if the stock market will return an extra 10% next year (above and beyond its expectations), AMD’s stock will likely return an extra $2 \cdot 10\% = 20\%$ (above and beyond AMD’s expectations). For now, let’s say that the expected rate of return on the market is 6% and the expected rate of return on AMD is 9%. Then, if the market were to turn in -4% (10% less than its expected return), you would expect AMD to turn in $9\% + 2 \cdot (-10\%) = -11\%$. Conversely, if the market were to turn in 16% (10% more than its expected return), you would expect AMD to turn in $9\% + 2 \cdot (10\%) = 29\%$. AMD’s high market beta is useful because it

This is important to keep in mind if you use a spreadsheet to check your work.

For market beta, the divisor cancels out and does not matter.

My fault: Our notation should have distinguished between true population and estimated sample statistics.

Market beta works well when investors are holding the market and adding only a little of your project.

Most financial websites publish market beta estimates.

Beta can be viewed as the marginal change of your project with respect to the market.

Company	Stock Ticker	Mkt Cap	Market Beta		Company	Stock Ticker	Mkt Cap	Market Beta	
			Yahoo	Google				Yahoo	Google
Intel	INTC	151	1.06	0.96	IBM	IBM	146	0.81	0.71
Coca-Cola	KO	199	0.82	0.52	PepsiCo	PEP	149	0.69	0.49
AMD	AMD	3.4	2.13	2.29	NVIDIA	NVDA	25	1.19	1.31
Ford	F	52	1.26	1.37	General Motors	GM	45	1.59	1.71
Apple Inc	AAPL	541	1.49	1.03	Google (Alphabet)	GOOG	493	1.03	N/A
Citigroup	C	129	1.54	1.97	Morgan Stanley	MS	49	1.45	1.63
Goldman Sachs	GS	66	1.32	1.69	J.P. Morgan	JPM	233	1.18	1.14
Volkswagen	VLKAY	75	1.89	1.83	Sony	SNE	35	1.49	1.83
Philip Morris	PM	157	1.02	0.95	Procter&Gamble	PG	221	0.65	0.49
American Airlines	AAL	19	3.95	N/A	Southwest	LUV	28	1.29	1.03
Boeing	BA	84	1.23	1.08	Airbus	AIR	41	N/A	N/A
Hewlett-Packard	HPQ	23	1.54	1.70	Yahoo	YHOO	35	2.16	1.48
Exxon Mobil	XOM	373	0.90	0.83	Barrick Gold	ABX	23	-0.20	-0.01

Exhibit 8.6: Some Market Betas and Market Capitalizations in June 2016. “MktCap” is the equity market value in billions of dollars. **YAHOO! FINANCE** explained its betas as follows: *The Beta used is Beta of Equity. Beta is the monthly price change of a particular company relative to the monthly price change of the S&P500. The time period for Beta is 3 years (36 months) when available. YAHOO! FINANCE ignores dividends, but this usually makes little difference. I could not find an explanation for Google’s market betas. I hope it’s not a secret.*

informs you that if you hold the stock market, adding AMD stock would not help you much with diversifying your market risk. Holding AMD would amplify any market swings.

Beta is not alpha.

But in any case, AMD’s market beta does not tell you whether AMD 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 AMD is. (This would be the aforementioned alpha [which can be interpreted as an expected rate of return]. In the next chapter, you will learn that the CAPM formula relates market beta to the expected rate of return, giving you a commonly used benchmark for alpha.)

► Alpha,
Pg.179.

Beta is also a “hedge ratio.”

Betas have another common important use. Let’s say that you want to speculate that AMD will go up, but you do not want to be exposed to market risk. The AMD beta of 2 tells you that if you buy long \$100 of AMD stock and go short \$200 in the stock market (which you can do easily, e.g., by shorting the SPDR ETF), 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, AMD goes up (down) on average by \$2. Thus, the market-beta of 2 is also the **hedge ratio** that tells you how you can “immunize” a speculative stock position against market-wide changes.

Q 8.9. You estimate your project x to return -5% if the stock market returns -10%, and +5% if the stock market returns +10%. What would you use as the market beta estimate for your project?

Q 8.10. You estimate your project y to return +5% if the stock market returns -10%, and -5% if the stock market returns +10%. What would you use as the market beta estimate for your project?

8.5 Market Betas for Portfolios and Conglomerate Firms

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 million and division C is worth \$2 million, then a firm consisting of B and C is worth \$3 million. 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 Exhibit 8.1 on Page 166, 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 S1 (♣)	In S2 (♦)	In S3 (♥)	In S4 (♠)	Reward E(r)	Variance ^a Var(r)	Risk Sdv(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. Check this.

Unfortunately, you cannot compute value-weighted averages for all statistics. As the table shows, variances and standard deviations cannot be averaged (1/3 · 25%% + 2/3 · 81%% ≈ 62.3%%, which is not the variance of 30%%; and 1/3 · 5% + 2/3 · 9% ≈ 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 8.3, you already computed the market-betas for as +0.64 and -1.60. So, their value-weighted average is

$$\beta_{BCC} = \frac{1}{3} \cdot (+0.64) + \frac{2}{3} \cdot (-1.60) \approx -0.8533 \quad (8.5)$$

$$\beta_{BCC} = w_B \cdot \beta_B + w_C \cdot \beta_C$$

You will be asked to confirm this conclusion in Q8.11.

Portfolios consist of multiple assets (themselves possibly portfolios). Definitions of value-weighted and equal-weighted portfolios.

What are the expected rate of return and market beta of a portfolio?

You can average actual rates of return.

You can average expected rates of return.

(But you cannot average variances or standard deviations!)

News flash: You can also average market betas.

► Market betas of B and C, Formula 8.3, Pg.176.

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

$$r_{ab} = w_a \cdot r_a + w_b \cdot r_b$$

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,

$$E(r_{ab}) = w_a \cdot E(r_a) + w_b \cdot E(r_b)$$

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,

$$\beta_{ab} = w_a \cdot \beta_a + w_b \cdot \beta_b$$

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.

A firm is a portfolio of debt and equity. Thus, the portfolio formulas apply to the firm (with debt and equity as its components), too!

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 and your \$300 million of equity has a market beta of, say, 2.0, then your firm has a market beta of

$$\frac{1}{4} \cdot (0.4) + \frac{3}{4} \cdot (2.0) = 1.6 \quad (8.6)$$

$$\left(\frac{\text{Debt value}}{\text{Firm value}} \right) \cdot \beta_{\text{Debt}} + \left(\frac{\text{Equity value}}{\text{Firm value}} \right) \cdot \beta_{\text{Equity}} = \beta_{\text{Firm}}$$

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), 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.

Q 8.11. Let’s check that the beta combination formula (Formula 8.5 on page 185) is correct. Start with the BCC line in the table on Page 185

- Write down a table with the demeaned market rate of return and demeaned BCC rate of return in each of the four possible states.

2. Multiply the demeaned rates of return in each scenario. This gives you four cross-products, each having units of %.
3. Compute the average of these cross-products. This is the covariance between BCC and M.
4. Divide the covariance between BCC and M by the variance of the market.
5. Which is faster—this route or Formula 8.5? Which is faster if there are a hundred possible scenarios?

Q 8.12. Confirm that you cannot take a value-weighted average of component variances (and thus of standard deviations) the same way that you can take value-weighted average expected rates of return and value-weighted average market betas.

1. What is the value-weighted average variance of BCC?
2. What is the actual variance of BCC?

Q 8.13. Consider an investment of $\frac{2}{3}$ in B and $\frac{1}{3}$ in C. Call this new portfolio BBC. Compute the variance, standard deviation, and market beta of BBC. Do this two ways: first from the four individual scenario rates of return of BBC, and then from the statistical properties of B and C itself.

Q 8.14. Assume that a firm will always have enough money to pay off its bonds, so the beta of its bonds is 0. (Being risk free, the rate of return on the bonds is obviously independent of the rate of return on the stock market.) Assume that the beta of the underlying assets is 2. What would financial websites report for the beta of the firm's equity if it changes its current capital structure from all equity to half debt and half equity? To 90% debt and 10% equity?

Q 8.15. (Advanced) Does maintaining a value-weighted or an equal-weighted portfolio require more trading? (Hint: Make up a simple example.)

Summary

This chapter covered the following major points:

- The expected rate of return is a measure of expected reward:

$$E(r_p) = \frac{\text{Sum over Scenarios } [r_p \text{ in Scenario}]}{N}$$

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

$$\text{Var}(r_p) = \frac{\text{Sum over Scenarios} \{ [r_p \text{ in Scenario}] - E(r_p) \}^2}{N \quad (\text{or } N - 1)}$$

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.

$$\text{Sdv}(r_p) = \sqrt{\text{Var}(r_p)}$$

- 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 higher expected rates of return at lower risks than undiversified ones.
- 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 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.

Keywords

Asset beta, 186. Covariance, 180. Diversification, 169. Efficient frontier, 171. Equal-weighted portfolio, 185. Equity beta, 186. Expected rate of return, 167. Hedge ratio, 184. Linear regression, 180. MVE, 188. Market beta, 176. Market model, 176. Market portfolio, 173. Mean-variance efficiency, 188. Minimum-variance portfolio, 172. Portfolio risk, 168. Reward, 167. Sharpe ratio, 171. Standard deviation, 168. Value-weighted portfolio, 185. Variance, 167.

Answers

Q 8.1 The average deviation from the mean is always 0.

Q 8.2 The mean of portfolio M was 4%. Adding 5% to each return will give you a mean of 9%, which is 5% higher. The variance and standard deviation remain at the same level, the latter being 5%. If you think of 5% as a constant $c = 5\%$, then you have just shown that $E(r + c) = E(r) + c$ and $Sdv(r + c) = Sdv(r)$.

Q 8.3 The reward of portfolio C is its expected rate of return, i.e., $[(17\%) + 3\% + 11\% + (-7\%)]/4 = 6\%$. (We can just divide by 4, rather than multiply each term by 1/4, because all outcomes are equally likely.) The variance of C is $[(11\%)^2 + (3\%)^2 + (5\%)^2 + (-13\%)^2]/4 = 81\%$. The standard deviation, which is our measure of risk, is $\sqrt{81\%} \approx 9\%$.

Q 8.4 The combination portfolio MA of 90% in M and 10% in A has rates of return of -2.4%, 3.8%, 4.2%, and 10.4%.

- Thus, its mean rate of return is 4%. Its variance is 20.5%. Its standard deviation is approximately 4.528%.
- It would look more spread out, because it has higher standard deviation.

Q 8.5 1. The reward is $4 \cdot 10\% = 40\%$. The variance is $4 \cdot 400\% = 1,600\%$. Thus, the standard deviation (risk) is $\sqrt{1,600\%} = 40\%$. The Sharpe ratio is 1.

- The reward is 90%. The risk is $\sqrt{9 \cdot 400\%} = 3 \cdot 20\% = 60\%$. The Sharpe ratio is 1.5
- The reward is $T \cdot E$. The standard deviation is $\sqrt{T} \cdot Sdv$. The Sharpe ratio is $(\sqrt{T} \cdot E) / Sdv$.

Q 8.6 Exhibit 8.3 shows that by combining M, A, B, and C, you get a risk-free rate of 3.6%; and investing in F alone gets you a risk-free rate of 1%. This means that you could borrow at 1% and invest at 3.67%, both risk-free—an arbitrage. The efficient frontier would be a vertical line at 0. Obviously, this could never be the case in the real world.

Q 8.7 For the MB portfolio, the portfolio combination rates of return in the four scenarios were on the bottom of Exhibit 8.4 on Page 175. Confirm them first:

$$\begin{aligned} \text{In S1 (♣): } & 0.5 \cdot (-3\%) + 0.5 \cdot (5\%) = 1\% \\ \text{In S2 (♦): } & 0.5 \cdot (3\%) + 0.5 \cdot (-1\%) = 1\% \\ \text{In S3 (♥): } & 0.5 \cdot (5\%) + 0.5 \cdot (7\%) = 6\% \\ \text{In S4 (♠): } & 0.5 \cdot (11\%) + 0.5 \cdot (13\%) = 12\% \end{aligned}$$

The expected rate of return is

$$E(r_{MB}) = \frac{1\% + 1\% + 6\% + 12\%}{4} = 5\%$$

The portfolio variance is

$$\begin{aligned} \text{Var}(r_{MB}) = & [(1\% - 5\%)^2 + (1\% - 5\%)^2 \\ & + (6\% - 5\%)^2 + (12\% - 5\%)^2]/4 \end{aligned}$$

Therefore, $Sdv(MC) = \sqrt{20.5\%} \approx 4.52\%$.

For the MC portfolio,

$$\begin{aligned} \text{In S1 (♣): } & 0.5 \cdot (-3\%) + 0.5 \cdot (17\%) = 7\% \\ \text{In S2 (♦): } & 0.5 \cdot (3\%) + 0.5 \cdot (3\%) = 3\% \\ \text{In S3 (♥): } & 0.5 \cdot (5\%) + 0.5 \cdot (11\%) = 8\% \\ \text{In S4 (♠): } & 0.5 \cdot (11\%) + 0.5 \cdot (-7\%) = 2\% \end{aligned}$$

The expected rate of return is

$$E(r_{MC}) = \frac{7\% + 3\% + 8\% + 2\%}{4} = 5\%$$

The variance is $\text{Var}(MC) = [(7\% - 5\%)^2 + (3\% - 5\%)^2 + (8\% - 5\%)^2 + (2\% - 5\%)^2]/4 = 26\%$. Therefore, $Sdv(MC) = \sqrt{26\%} \approx 5.1\%$.

Q 8.8 The order of subscripts on market beta is important. Algebraically, $\beta_{C,M} = [\text{cov}(r_C, r_M)] / [\text{var}(r_M)]$, while $\beta_{M,C} = [\text{cov}(r_C, r_M)] / [\text{var}(r_C)]$. The denominator is different. If you work this out, $\beta_{M,C} \approx -0.49$. Fortunately, you will never ever need to compute $\beta_{M,C}$. I only asked you to do this computation so that you realize that the subscript order is important.

Q 8.9 The market beta of this project is

$$\beta_{x,M} = \frac{r_{x,2} - r_{x,1}}{r_{M,2} - r_{M,1}} = \frac{(-5\%) - (+5\%)}{(-10\%) - (+10\%)} = +0.5$$

(This is not “half as volatile” because market beta is not a measure of volatility.)

Q 8.10 Using the same formula, the market beta of y is $[(+5\%) - (-5\%)] / [(-10\%) - (+10\%)] = -0.5$.

Q 8.11 1. Start with our standard table:

	♣	♦	♥	♠	E(r)	Var(r)	Sdv(r)
BCC	13%	1.67%	9.67%	-0.33%	6%	30%	5.5%
... in dev	7%	-4.33%	3.67%	-6.33%			
M	-3%	+3%	+5%	+11%	4%	25%	5%
... in dev	-7%	-1%	1%	+7%			

(Variances and standard deviations are rounded.)

- The four cross-products are -49%, 4.33%, 3.67%, and -44.33%.
- The average (covariance) is -21.33%.
- The beta is $-21.33/25 \approx -0.8533$.
- This is the more painful route—and it is more painful when there are more possible scenarios.

Q 8.12 Actually, this was already in the text. BCC has a variance of about 30%%, while the value-weighted average of the variances is about 62.3%%.

Q 8.13 The equivalent table is

	♣	♦	♥	♠	$E(r)$	$\text{Var}(r)$	$\text{Sdv}(r)$
B	5%	-1%	7%	13%	6%	25%%	5%
C	17%	3%	11%	-7%	6%	81%%	9%
BBC	9%	0.33%	8.33%	6.33%	6%	11.67%%	3.4%

The market beta is easiest to compute as $\frac{2}{3} \cdot \beta_B + \frac{1}{3} \cdot \beta_C \approx \frac{2}{3} \cdot (0.64) + \frac{1}{3} \cdot (-1.60) \approx -0.11$.

Q 8.14 For a firm whose debt is risk free, the overall firm beta is $\beta_{\text{Firm}} = 0.5 \cdot \beta_{\text{Equity}} + 0.5 \cdot \beta_{\text{Debt}}$. Thus, $0.5 \cdot \beta_{\text{Equity}} + 0.5 \cdot 0 = 2$. Solve for $\beta_{\text{Equity}} = \beta_{\text{Firm}}/0.5 = 4$. For the (90%, 10%) case, the equity beta jumps to $\beta_{\text{Equity}} = 2/0.1 = 20$.

Q 8.15 Value-weighted portfolios usually require no trading (unless there is a payout, like a dividend). For example, using the numbers from this section, if B triples from \$1 million to \$3 million and C halves from \$2 million to \$1 million, your original value-weighted portfolio or firm would become $\$3 + \$1 = \$4$ million. You would still be exactly value-weighted. B would now constitute 75% of the firm and C 25% of the firm. In contrast, in an originally equal-weighted portfolio, your \$1.5 million in B would become \$4.5 million, your \$1.5 million in C would become \$0.75 million, and your portfolio would be worth \$5.25 million. This means you would want to have \$2.625 million invested in each. To maintain an equal-weighted portfolio, you would have to sell some stock in your past winner to buy some stock in your loser. Only a value-weighted portfolio requires no trading. Another interesting aspect is that if you do not trade, in the very long run, any portfolio will look more and more value-weighted, because those stocks that have had large returns will automatically garner a larger weight both in your portfolio and the economy.

End of Chapter Problems

Q 8.16. Multiply each rate of return for M by 2.0. This portfolio offers $-6%$, $+6%$, $+10%$, and $+22%$. Compute the expected rate of return and standard deviation of this new portfolio. How do they compare to those of the original portfolio M?

Q 8.17. The following table contains the closing year-end prices of the Japanese stock market index, the Nikkei-225. Assume that each historical rate of return was exactly one representative scenario (independent sample draw) that you can use to estimate the future. If a Japanese investor had purchased a mutual fund that imitated the Nikkei-225, what would her annual rates of return, compounded rate of return (from the end of 1984 to the end of 2010), average rate of return, and risk have been?

Year	N-225	Year	N-225	Year	N-225
1984	11,474	1993	17,417	2002	8,579
1985	13,011	1994	19,723	2003	10,677
1986	18,821	1995	19,868	2004	11,489
1987	22,957	1996	19,361	2005	16,111
1988	29,698	1997	15,259	2006	17,225
1989	38,916	1998	13,842	2007	15,308
1990	24,120	1999	18,934	2008	8,860
1991	22,984	2000	13,786	2009	10,546
1992	16,925	2001	10,335	2010	10,229

Q 8.18. Compute the value-weighted average of $1/3$ of the standard deviation of B and $2/3$ of the standard deviation of C. Is it the same as the standard deviation of a BCC portfolio of $1/3$ B and $2/3$ C, in which your investment rate of return would be $1/3 \cdot r_B + 2/3 \cdot r_C$?

Q 8.19. Why is it so common to use historical financial data to estimate future market betas?

Q 8.20. What are the risk and reward of a combination portfolio that invests 40% in M and 60% in B?

Q 8.21. Consider the following five assets, which have rates of return in six equally likely scenarios:

	Awful	Poor	Med.	Okay	Good	Great
Asset P1	$-2%$	$0%$	$2%$	$4%$	$6%$	$10%$
Asset P2	$-1%$	$2%$	$2%$	$2%$	$3%$	$3%$
Asset P3	$-6%$	$2%$	$2%$	$3%$	$3%$	$1%$
Asset P4	$-4%$	$2%$	$2%$	$2%$	$2%$	$20%$
Asset P5	$10%$	$6%$	$4%$	$2%$	$0%$	$-2%$

1. Assume that you can only buy one of these assets. What are their risks and rewards?
2. Supplement your previous risk-reward rankings of assets P1–P5 with those of combination portfolios that consist of half P1 and half of each of the other 4 portfolios, P2–P5. What are the risks and rewards of these four portfolios?
3. Assume that P1 is the market. Plot the rates of return for P1 on the x-axis and the return for each of the other stocks on their own y-axes. Then draw lines that you think best fit the points. Do not try to compute the beta—just use the force (and your eyes), Luke. If you had to buy just a little bit of one of these P2–P5 assets, and you wanted to lower your risk, which would be best?

Q 8.22. Assume that you have invested half of your wealth in a risk-free asset and half in a risky portfolio P. Is it theoretically possible to lower your portfolio risk if you move your risk-free asset holdings into another risky portfolio Q? In other words, can you ever reduce your risk more by buying a risky security than by buying a risk-free asset?

Q 8.23. Is it wise to rely on historical statistical distributions as your guide to the future?

Q 8.24. Look up the market betas of the companies in Exhibit 8.6. Have they changed dramatically since June 2016, or have they remained reasonably stable?

Q 8.25. You estimate your project to return $-20%$ if the stock market returns $-10%$, and $+5%$ if the stock market returns $+10%$. What would you use as the market beta estimate for your project?

Q 8.26. Go to [YAHOO! FINANCE](http://YAHOO!FINANCE). Obtain two years' worth of daily stock rates of return for PepsiCo, Coca Cola, and for the S&P 500 index. Use a spreadsheet to compute PepsiCo's and Coca-Cola's *historical* market betas. (Note: For future market betas, you should further shrink towards 1.)

Q 8.27. Consider the following assets:

	Bad	Okay	Good
Market M	-5%	5%	15%
Asset X	-2%	-3%	25%
Asset Y	-4%	-6%	30%

1. Compute the market betas for assets X and Y.
2. Compute the correlations of X and Y with M.
3. Assume you were holding only M. You now are selling off 10% of your M portfolio to replace it with 10% of either X or Y. Would an M&X portfolio or an M&Y portfolio be riskier?
4. Is the correlation indicative of which of these two portfolios ended up riskier? Is the market beta indicative?

Q 8.28. Compute the expected rates of return and the portfolio betas for many possible portfolio combinations (i.e., different weights) of M and F from Exhibit 8.1 on Page 166. (Your weight in M is 1 minus your weight in F.) Plot the two against one another. What does your plot look like?

Q 8.29. Are historical covariances or means more trustworthy as estimators of the future?

Q 8.30. Are geometric average rates of return usually higher or lower than arithmetic average rates of return?

Q 8.31. The following represents the probability distribution for the rates of return for next month:

Probability	Pfio P	Market M
$1/6$	-20%	-5%
$2/6$	-5%	+5%
$2/6$	+10%	0%
$1/6$	+50%	+10%

Compute by hand (and show your work) for all the following questions.

1. What are the risks and rewards of P and M?
2. What is the correlation of M and P?
3. What is the market beta of P?
4. If you were to hold $1/3$ of your portfolio in the risk-free asset, and $2/3$ in portfolio P, what would its market beta be?

Q 8.32. Download the historical daily stock prices for the S&P 500 index and for VPACX (the *Vanguard Pacific Stock Index* mutual fund) from [YAHOO! FINANCE](#), beginning January 1 three years ago and ending December 31 of last year. Load them into a spreadsheet and position them next to one another. Compute the risk and reward. Compute VPACX's market beta, i.e., with respect to the S&P 500 index. How do your historical estimates compare to the Fund Risk reported by [YAHOO! FINANCE](#) and other financial websites? If you were interested not in the *historical* but *future* market beta, would this be a good estimate?

Q 8.33. Download 3 years of historical daily (dividend-adjusted) prices for Intel (INTC) and the S&P 500 from [YAHOO! FINANCE](#).

1. Compute the daily rates of return.
2. Compute the average rates of return and risk of portfolios that combine INTC and the S&P 500 in the following proportions: (0.0, 1.0), (0.2, 0.8), (0.4, 0.6), (0.6, 0.4), (0.8, 0.2), (1.0, 0.0). Then plot them against one another. What does the plot look like?
3. Compute the historical market beta of Intel.

Q 8.34. Why do some statistical packages estimate covariances differently (and different from those we computed in this chapter)? Does the same problem also apply to expected rates of return (means) and betas?