



How To Measure and Manage Risk

Professor Alex Edmans

24th February 2022

What is Risk?

“What is risk?” seems to be an obvious question. Wikipedia defines risk as “the possibility of something bad happening”. For example, I have £10 to buy dinner, but I have a 50% chance of losing it.

This isn’t actually what we mean by risk. The 50% probability of losing your money means that your expected (= average, mean) cash flow is £5. You’re worse off, because your *expected cash flow* is lower. Instead, when we speak about risk, we hold constant the expected cash flow – we’re asking whether you prefer £5 for sure, or a 50-50 chance of either £10 or £0. A *risk-neutral* person is indifferent. A *risk-averse* person prefers £5 for sure.

Most people are risk-averse, because money exhibits *diminishing marginal utility*. Each marginal (i.e. extra) £ that you receive gives you utility (i.e. happiness). But the more money you already have, the smaller the extra utility from each additional £. Moving from £0 to £5 gives you significant happiness, because it means that you’ll be able to have at least some dinner and won’t go hungry. Moving from £5 to £10 makes you even happier as you can afford a better dinner – but doesn’t make as much difference as moving from £0 to £5.

Why do we care about risk? Recall from Lecture 3, [How to Make Financial Decisions](#), that we defined an *investment* as a decision that costs money today in the hope of making more money in the future, such as building a widget factory. We gave two reasons why we can’t evaluate an investment by simply comparing the expected future cash benefits with today’s cash cost:

- The costs are borne *today* but the benefits arise in the *future*. £1 today is worth more than £1 tomorrow, because you can do something with money today – you can put it in the bank and earn interest. This is known as the *time value of money*, and was the focus of Lecture 3.
- The costs are *certain* but the benefits are *risky*. You know how much the factory costs you, but you don’t know how much you’ll be able to sell the widgets for. It will depend on the state of the economy, customer reviews of your widgets versus a competitor’s, and so on. This is the focus of today’s lecture.

Intuitively, it seems that we should *discount* risky cash flows to reflect the fact that we don’t like risk, just like Lecture 3 showed that we should discount future cash flows to reflect the time value of money. In the gameshow *Deal or No Deal*, if the contestant is down to the last two boxes and they’re £1 or £1 million (meaning the contestant will get either £1 or £1 million with 50-50 probability), the expected prize is £500,000.50. But, if the banker offers them £300,000, most contestants will take it – they discount the prize down from £500,000.50 to £300,000 because of the risk.

But it seems that the amount we should discount them by is subjective – it depends on your attitude to risk. Some people (we’ll call them “chickens”¹) don’t like risk, but others (we’ll call them “lions”) do. This subjectivity seems to be a big problem for companies owned by many shareholders. A lion wants the company to invest in the widget factory but the chicken prefers it to pay out the cash as a dividend, because it’s safe. But it turns out that there’s an objective way to calculate the risk discount, meaning that both the chicken and lion will agree on whether to build the factory.


¹ This isn’t meant to be a pejorative term, just short-hand.

How Do We Measure Risk?

To calculate how much the risk discount should be, we first need to measure risk. The common measure is the *standard deviation* (denoted σ), which captures how “spread out” a distribution is. (A *distribution* gives the range of possible outcomes, e.g. possible cash flows for a widget factory).

The following chart shows three distributions all with an *expected value* or *mean* of 50, but different standard deviations:

- The blue curve, with the highest peak, has the smallest standard deviation – it’s very likely that the actual cash flow will be close to 50.
- The green curve, with the lowest peak, has the largest standard deviation – there’s a reasonable chance that the actual cash flow will be below 25 or above 75.



However, it turns out that standard deviation is actually the *wrong* measure of risk. To see why, consider the following two companies, Izy’s Ice Cream and Carola’s Coffee, where their cash flows depend on the *state* of the world – whether the weather is Sun or Rain:

	Sun	Rain	Mean	Standard Deviation
Izy’s Ice Cream	14	6	10	5.7
Carola’s Coffee	9	9	9	0

Which would you prefer? Izy’s Ice Cream offers a higher mean, but comes with a higher standard deviation. It seems that your choice of company depends on your preferences for risk vs. return. A lion might prefer Izy’s Ice Cream due to its higher mean, but a chicken might opt for Carola’s Coffee due to its lower standard deviation. Carola’s Coffee is immune to weather conditions, but Izy’s Ice Cream is highly sensitive to it.

This argument turns out to be wrong, because standard deviation isn’t the relevant measure of risk. Let’s introduce a third company, Ursula’s Umbrellas, which is *negatively correlated* with Izy’s Ice Cream – it does well in Rain (when Izy does badly) and badly in Sun (when Izy does well).

	Sun	Rain	Mean	Standard Deviation
Izy’s Ice Cream	14	6	10	5.7
Carola’s Coffee	9	9	9	0
Ursula’s Umbrellas	8	12	10	2.8

If the investor buys 1/3 of Izy and 2/3 of Ursula, his portfolio gives him 10 in each state, with no risk:

	Sun	Rain	Mean	Standard Deviation

1/3 of Izy and 2/3 of Ursula	10	10	10	0
------------------------------	----	----	----	---

1/3 of Izy and 2/3 of Ursula is unambiguously better than Carola's Coffee. It provides a higher return in each state. Thus, no-one holds Carola's Coffee, no matter how risk-averse they are. Even though Carola has a standard deviation of zero, and Izy and Ursula both have positive standard deviations, both the chicken and lion prefer Izy and Ursula to Carola. The "trick" is that investors don't have to hold just one stock, they can hold *portfolios* of stocks. By holding both Izy and Ursula, investors can *diversify* away the exposure of each individual stock to the weather, because weather affects Izy and Ursula in different directions. So, standard deviation is not the relevant measure of risk, because investors may be able to diversify away some of a company's standard deviation.

But diversification has its limits. While investors can diversify away some risk, they usually can't diversify away all risk. Now assume that the two relevant states are economic conditions (Boom and Bust), rather than weather conditions (Sun and Rain):

	Boom	Bust	Mean	Standard Deviation
Izy's Ice Cream	14	6	10	5.7
Carola's Coffee	13	7	9	4.2
Ursula's Umbrellas	12	8	10	2.8
1/3 of Izy and 2/3 of Ursula	7.3	12.7	10	3.8

By moving from holding only Izy to holding 1/3 of Izy and 2/3 of Ursula, the investor reduces his standard deviation from 5.7 to 3.8. (Indeed, if he moves towards holding Ursula only, he reduces his standard deviation even further to 2.8). But, he can't get his risk down to zero. This is because, unlike weather conditions, which affect Izy and Ursula in different directions and thus are diversifiable, economic conditions affect them in the same direction and so are not diversifiable. Both Izy and Ursula do badly in a Bust. Thus, moving from Izy to Ursula doesn't completely reduce the investor's vulnerability to economic conditions.

This means that there are two types of risk:

- *Idiosyncratic / unique* risk is specific to one particular company and not shared by all other companies. For Izy's Ice Cream, this is the risk of rain. For a football club, an example is the risk your star player gets injured. For a drug company, an example is the risk that you suffer a product recall.
 - Idiosyncratic risk is *diversifiable*. Since it's not shared by other stocks, you can reduce its effect by diversifying your portfolio and holding other stocks.
- *Systematic / market* risk is shared by the rest of the market. For Izy's Ice Cream, this is the risk of an economic downturn. A downturn also hurts the football club and drug company.
 - Systematic risk is *non-diversifiable*. Since it's shared by other stocks, you can't entirely eliminate its effect by diversifying your portfolio, because other stocks share this risk too. They all suffer in an economic downturn.

Even if a risk is diversifiable, it's important to stress that you're still affected by it. Consider a drug company that pays a single cash flow of £10 next year, but there's a 50% chance the cash flow is £0. Even if the risk is fully diversifiable, the expected cash flow falls to £5. All diversifiability affects is whether we need to discount the value of the drug company even further, for the fact that it doesn't give you £5 for certain, but a 50-50 chance of £10 or £0.

- If the risk of getting £0 is idiosyncratic, because it arises from a product recall unrelated to the rest of the market, we don't need to apply an additional discount. Even if the product gets recalled and the drug company is worth £0, you'll still be able to buy dinner as none of your other stocks are affected. Thus, you don't really mind that the stock is worth £10 or £0 rather than £5 for certain.

- If the risk of getting £0 is systematic, because it arises from an economic downturn, we do need to apply an additional discount. When the company is worth £0, all your other companies are worth less also because they suffer from the recession, so you might not be able to afford dinner. Thus, you do mind that the stock is worth £10 or £0 rather than £5 for certain – because the times in which the stock pays £0 are bad times when the rest of your portfolio is doing badly. The stock provides negative insurance, and so you do need to apply a risk discount. So even though the expected cash flow is £5, it might only be worth £4 to you.


How Do We Measure Systematic Risk?

So we've seen that the risk discount should be applied only to *systematic risk* – the non-diversifiable portion of risk, rather than the *total risk* σ . But how do we measure systematic risk?

Systematic risk is the risk that's shared with, i.e. *correlated* with, the rest of the market. The *correlation coefficient* measures the correlation between two stocks (or one stock and one portfolio such as the overall market). It's denoted ρ (the Greek letter "rho") and ranges from -1 to +1:

- $\rho = +1$ is *perfect positive correlation*, where the stocks move in lockstep.
- $\rho = -1$ is *perfect negative correlation*, where the stocks move in opposite directions - like Izy and Ursula when weather was the only source of uncertainty.
- $\rho = 0$ is *zero correlation*, where the stock returns bear no relation to each other.

The following graphs illustrate the relationship between the returns of two stocks for different levels of correlation:



The *correlation coefficient* is standardised to always be between -1 and 1. It only takes into account the correlation *between* two stocks; it ignores these stocks' actual level of risk. $\rho = 1$ means that, when stock 1 rises, stock 2 rises as well, but doesn't say anything about how much stocks 1 and 2 will typically rise by.

The *covariance* takes into account the stocks' actual level of risk (given by their standard deviation) and thus isn't constrained to lie between -1 and +1. The covariance between stock 1 and the market m is given by:

$$\sigma_{1m} = \rho_{1m}\sigma_1\sigma_m$$

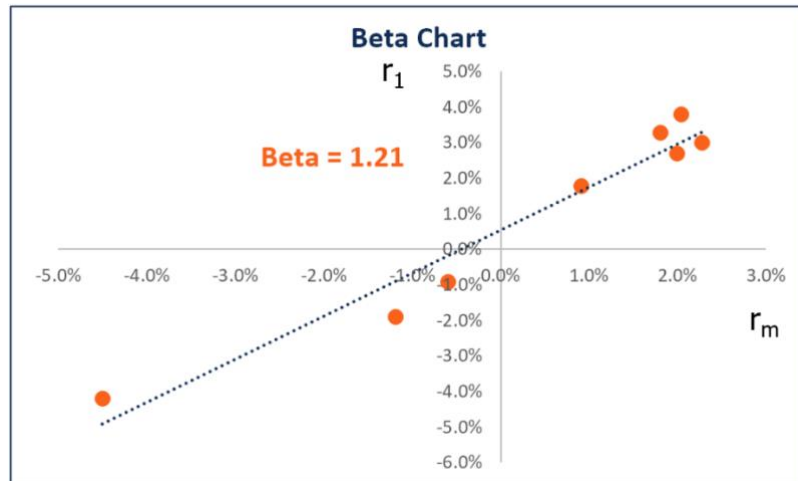
i.e. the correlation coefficient multiplied by the individual standard deviations.

We can finally show how to measure stock 1's systematic risk. You take the covariance between stock 1 and the market, and divide it by the variance of the market (the square of its standard deviation):

$$\beta_1 = \frac{\sigma_{1m}}{\sigma_m^2} = \frac{Cov(r_1, r_m)}{Var(r_m)}$$

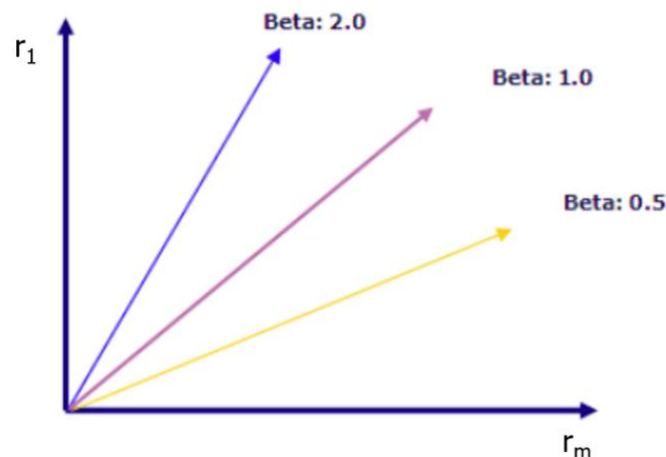
What Does Beta Mean?

Systematic risk is denoted β , a stock's *beta*, and is one of the most famous variables in finance. Mathematically, it's the covariance divided by the variance. But let's step back from the maths and see what it actually means. If you plot a graph of the return on stock 1 versus the return on the market, beta is the slope of the *best-fit line*, otherwise known as the *regression coefficient*. It measures how much, on average, stock 1 rises by when the market rises by 1%. In the diagram below, the slope of the best-fit line is 1.21. This means that, on average, when the market goes up by 1%, the return on stock 1 goes up by 1.21%.



While this seems a bit like the correlation coefficient ρ (which we also illustrated by plotting graphs of the returns of two stocks), it's distinct:

- ρ measures how close the dots are to the best fit line; the closer they are to -1 or +1, the better the fit.
- β measures how steep the best-fit line is (its slope), and is not bounded between -1 and +1. The following chart shows how the slope of the best-fit line varies with β :



What determines beta? One factor is *business risk*, the cyclicality of the company's business – i.e. how sensitive the company's business is to movements in the market. For example:

- If $\beta > 1$, the stock increases more than 1% when the market increases by 1%. Examples are *luxury goods* such as jewellery. Only in good times will people buy them.
- If $\beta = 1$, the stock has the same systematic risk as the market – it increases by 1% when the market increases by 1%.
- If $0 < \beta < 1$, the stock increases less than 1% when the market increases 1%. Examples are *consumer goods* such as most food items. In good times, you can buy more food, but there's a limit to how much extra food you can eat.

- If $\beta = 0$, the stock is uncorrelated with the market.
- If $\beta < 0$: the stock decreases when the market increases. An example is an insolvency company that helps restructure businesses when they go bankrupt.

The second factor is *financial risk*, how much debt the company has. We'll consider this in Lecture 6, [How to Finance a Company](#).

The Relationship Between Risk and Return

Let's take stock (sorry) of where we've got to so far.

- Investors are typically risk-averse, so they'll discount cash flows that are risky.
- But they won't discount the entire amount of a company's risk. Idiosyncratic risk is diversifiable, and so shouldn't lead to a risk discount. Only non-diversifiable risk should.
- Non-diversifiable risk is measured by beta – how much a stock's return rises when the market rises by 1%.
- So the risk discount of a stock depends on its beta.

There's one final step – to show *how much* the risk discount of a stock depends on its beta. It's given by the following equation:

$$r_1 = r_f + \beta_1(r_m - r_f)$$

where r_f is the risk-free rate. r_1 is known as the *cost of equity*, or discount rate for a stock. This equation is the Nobel Prize-winning *Capital Asset Pricing Model (CAPM)*. It shows how much extra return a stock should offer to compensate for its risk.

Cast your mind back to Lecture 3 ([How to Make Financial Decisions](#)), where we assumed away risk and considered only certain cash flows. Back then, we discounted all cash flows by the risk-free rate. For example, the present value of £C in year 1 is $\text{£}C \times \frac{1}{1+r_f}$ today. However, if the cash flows from stock 1 are risky, we discount them not by the risk-free rate r_f , but by the cost of equity $r_1 (= r_f + \beta(r_m - r_f))$. This is higher than r_f due to two factors:

- β_1 , the *amount of risk* of stock 1.
- $r_m - r_f$, the *price of risk*. This shows how much extra return a stock needs to offer to compensate for its systematic risk. This extra return is given by the *market risk premium*, the excess return the market offers over and above the risk-free rate.
 - Recall that a stock with $\beta = 1$ is as risky as the market. Thus, it should offer the same excess return as the market. The greater the systematic risk, the greater the extra return the stock should offer, which is why these two terms are multiplied.

An Example

Vodafone is considering a project that pays a single cash flow next year, which is either £10 or 0 with 50-50 probability. The project has the same systematic risk as Vodafone. Vodafone's beta is 0.88, the market risk premium is 3.93%, and the risk-free rate is 0.897%. What is the present value of the project?

Using the CAPM, we have $r_1 = r_f + \beta(r_m - r_f) = 0.897\% + 0.88 \times 3.93\% = 4.36\%$.

The expected cash flow is £5, and so the present value of the project is $5/1.0436 = \text{£}4.79$.

Thus, if the cost of the project is less than £4.79, the *Net Present Value* is positive and you should take the project.

Note that if the risk were idiosyncratic ($\beta = 0$), we would discount by the risk-free rate, and so the present value would be $5/1.0087 = \text{£}4.96$. It's less than £5 due to the *time value of money*. Thus, with $\beta = 0.88$, the project is worth only £4.79 rather than £5 for two reasons:

- The £5 is received in the future, reducing its present value to £4.96.

- The £5 bears systematic risk, reducing its present value further to £4.79. Thus, the project suffers an additional risk discount because of it offers a 50-50 chance of £10, rather than £5 for certain.

Typically, people speak about *risk premiums*, not risk discounts – the CAPM gives you the *risk premium* on stock 1, or the extra return that stock 1 has to offer (above the risk-free rate) to compensate for its risk. This extra return leads to a higher discount rate, so the risk premium becomes a risk discount when we use it to discount cash flows.

© Professor Edmans 2022

Some of this summary is adapted from *Principles of Corporate Finance* by Richard Brealey, Stewart Myers, Franklin Allen, and Alex Edmans (14th edition, 2022). Please see the lecture for additional details, such as how to estimate betas and market risk premiums in practice.