Topic 7: Portfolio Selection and CAPM

Eddie Shore
Corporate Finance – ECONS4280

July 20, 2022

In these notes, we'll synthesize our work on establishing portfolio variance and our new understanding of how stock returns work, to arrive at a theory of portfolio selection. We'll begin by outlining the *sources* of uncertainty in dividend payments. In particular we will emphasise that there are both *idiosyncratic* and *market* shocks that will affect a company's ability to earn money, and hence pay out dividends. We'll then reemphasise a point made previously that *diversification* can lower the variance of a portfolio. We will then argue that this *diversification* lowers the variance by *reducing/eliminating* idiosyncratic risk. *Market* risk cannot be eliminated entirely via diversification. We'll then turn to the notion of an *efficient* portfolio. An efficient portfolio is one that offers the *highest* return for a given level of risk. Finally, we will turn to models that identify these efficient portfolios, notably the famous CAPM model.

1 The risks of a given stock

Companies face random shocks all the time. Suppose a famous celebrity comes out and claims that your product gave them acne. This is bad news! People stop buying your product, your earnings go down, so your ability to pay dividends diminishes. This is an example of an *idiosyncratic* shock —the shock *only affects you.* The degree to which a stock has idiosyncratic risk is often labelled as the stock's α

Now consider an alternative shock: Russia invades Ukraine. It should go without saying that the real tragedy of this shock has nothing to do with corporate finance, but rather the unacceptable loss of life and livelihoods that this war has brought about. Nonetheless, this shock has also had financial consequences. Perhaps the

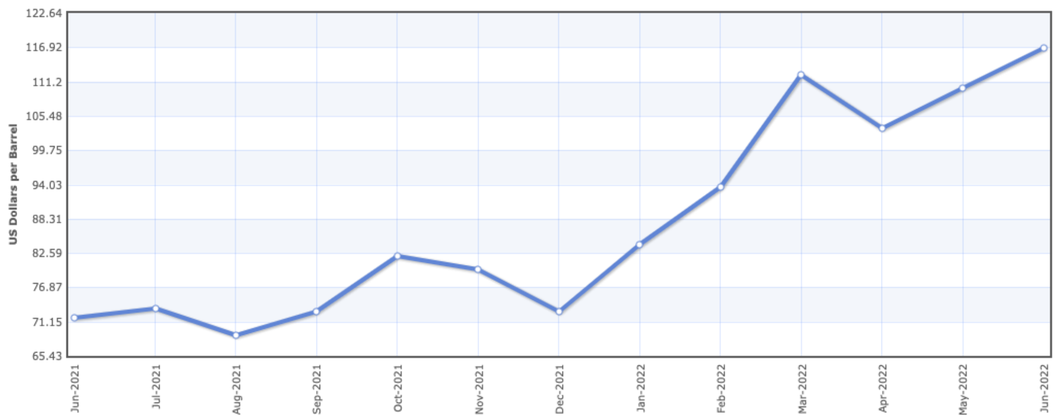

most obvious has been the impact on oil prices: Oil price shocks have far-reaching consequences for (almost) every firm. When oil prices are high, energy costs go up, and almost everyone requires energy to run their business. Similarly, the rise in energy costs tends to cause firms to raise their prices. Thus, the cost of inputs also increases. The Ukraine war is an example of a *market* shock —the whole market is affected.

Note that although this shock affects the entire market, *it will not affect everyone equally.* Some firms will rely heavily on oil products for their production. Likely these firms will suffer greatly. Others will rely very little on oil. Likely they will not be greatly affected. The degree to which a stock moves with market shocks is often labelled as the stock's β.

We can express the returns to a stock as being the sum of idiosyncratic shocks and responses to market shocks. Let market returns be denoted by r_t^m. Then:

$$r_{i,t} = \alpha_{i,t} + \beta_i r_t^m$$

Here, the $\alpha_{i,t}$ is the idiosyncratic shock to firm i at time t, and the β_i captures how much a firm is affected by market shocks.

2 What does this have to do with portfolios?

Recall the graph from the notes, 'Topic 5 - Risk and Return', where I randomly chose stocks for different sized portfolios, and then calculated the portfolio variance? What we saw was that, as the size of my portfolio increased, the portfolio variance fell.

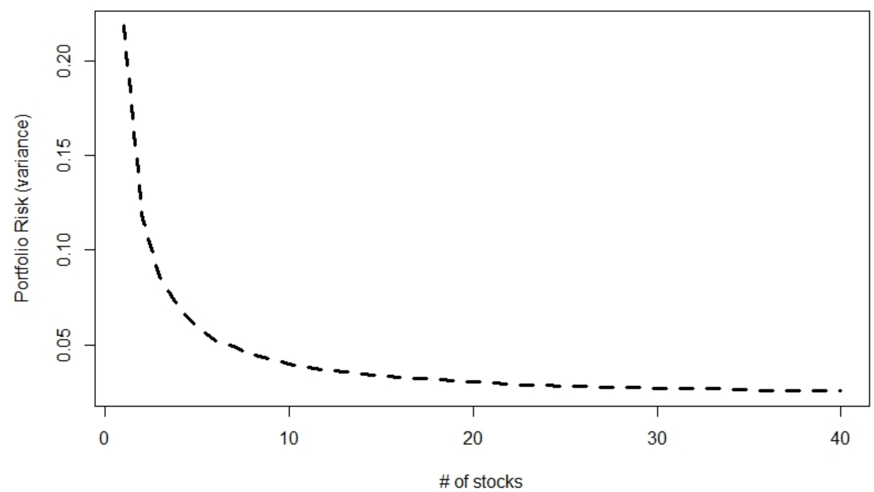

Why? Well, put simply, *as the size of my portfolio grows*, the impact of *idiosyncratic* shocks on my portfolio returns *goes to zero.* If the idiosyncratic shocks are *truly* idiosyncratic, *then on average they should offset each other*! Of course, this assumes that positive shocks are just as likely as negative shocks.

Note that there was a limit on how low I could make my portfolio variance just by growing its size. That's because, whilst the idiosyncratic shocks cancel out, *diversification will only eliminate market risk if firm responses to market shocks are symmetric around zero.* That is to say, only if my portfolio is made up of firms that *perfectly* offset each other's responses to market shocks.

2.1 How likely is it that my portfolio βs will offset market risk?

Not very. To illustrate, I took a random sample of 5,000 US firms and their monthly returns, and estimated their β's. I did this by first finding the average return of the entire market in each month, which I label r_t^m. What I want to find is the β_i for each of my 5,000 firms, so I then run a regression on each stock's data of the following form:

$$r_t^i = \beta_i r_t^m + \alpha_{i,t}$$

So to find β_i, I am *only* looking at data from firm i. Note that the $\alpha_{i,t}$ is the *idiosyncratic* shock for firm i. Once I've done this, I have 5,000 β's: $\{\beta_1, \beta_2, ..., \beta_{5,000}$.

I can then plot a histogram of those estimated β's. I can also plot the idiosyncratic shocks as well, $\alpha_{i,t}$. What do I find?

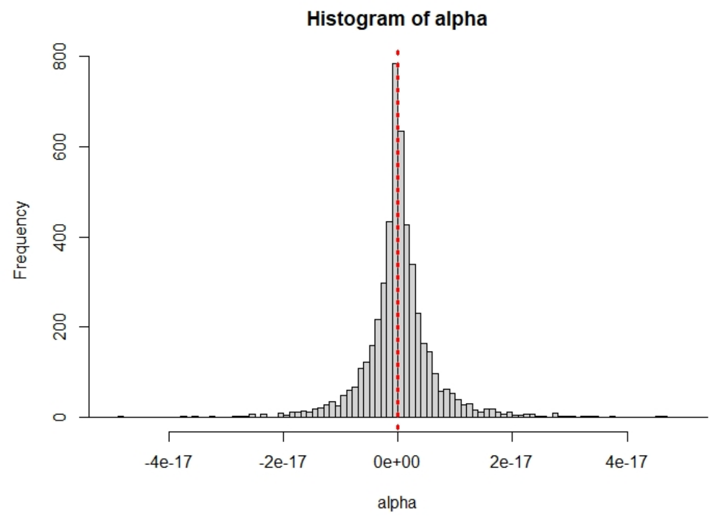

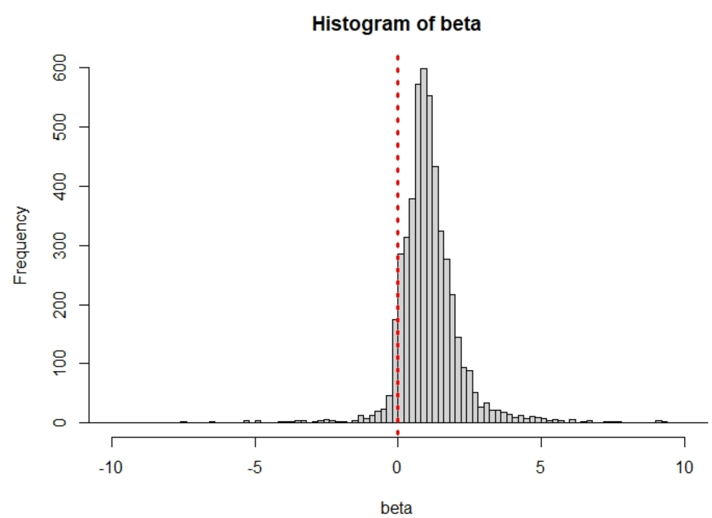

The first thing we can say is, idiosyncratic shocks do *indeed* appear to be just as likely to be positive and negative. So our claim that they should cancel each other out if we pick stocks randomly bears out! The second thing we can say is that this is definitely *not* true in the case of β. The *vast* majority (roughly 92%) of the stocks have β's that are *positive*. Thus, by picking stocks randomly, *we cannot expect to eliminate market risk*. If something shocks the market, *it will shock our portfolio*.

3 Efficient Portfolios

So, we know that two forces influence stock returns: idiosyncratic and market shocks. We know that with diversification, we can *eliminate* idiosyncratic risk in our portfolio. By contrast, we know that *we can't* eliminate market risk because *the majority* of stocks co-move with the market —when the market moves down, our portfolio moves down. We *also* know that risk and return are correlated: more risk = more reward. So, how do we *combine* what we now know about risk *and* reward, so that we arrive at a means of selecting a portfolio that captures our *preferences* for risk and reward? We'll start by thinking about the *range* of potential risk and reward pairings that we can achieve by mixing different stocks. We'll then think about whether there is an *optimal* mixing that we would prefer over any other.

3.1 A very simple example

Suppose we are choosing between two stocks, A and B. Suppose that these stocks have the following characteristics *ex-ante*; that is to say, these are what we *expect* their characteristics to be:

	Return, r	*Variance*, σ_i^2	*Covariance*, σ_{AB}
A	0.07	0.1	-0.05
B	0.02	0.03	-0.05

So A has higher returns *and* higher variance, and they co-move together *negatively*: when A moves up, B typically moves down. Suppose we are choosing how much of each stock to buy, i.e. the weights $\{\omega_A, \omega_B\}$. We know from previous notes that:

$$r_P = \omega_A r_A + \omega_B r_B$$

$$\sigma_P^2 = \omega_A^2 \sigma_A^2 + \omega_B^2 \sigma_B^2 + \omega_A \omega_B \sigma_{AB}$$

So, inputting our numbers:

$$r_P = 0.07\omega_A + 0.02\omega_B$$

$$\sigma_P^2 = 0.1\omega_A^2 + 0.03\omega_B^2 - 0.05\omega_A\omega_B$$

So, what range of *returns* and what range of *variance* can we achieve by varying the weights, $\{\omega_A, \omega_B\}$? Note that whatever I choose for ω_A, $\omega_B = 1 - \omega_A$, so really I'm only choosing one thing: ω_A. So, for $\omega_A = 1$, I'm only selecting asset A, and for $\omega_A = 0$, I'm only choosing asset B. Then, the returns I can get look like this:

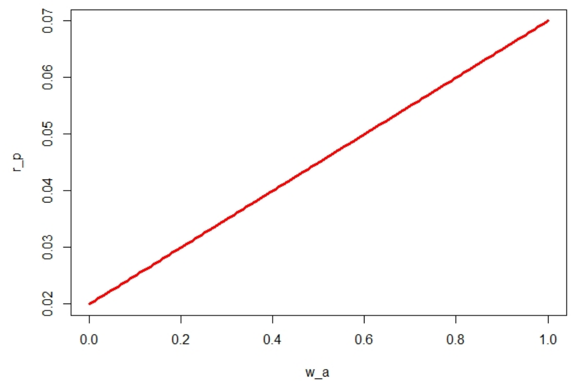

As discussed above, when $\omega_A = 0$, then I am *only* getting the expected returns from asset B, which is equal to 0.02. Similarly, if $\omega_A = 1$, I am *only* getting the expected returns from asset A, which is equal to 0.07. For any combination, I am somewhere along the straight line in the graph above.

Okay, so what about the variance? Note I will just use the formula above: $\sigma_P^2 = \omega_A^2\sigma_A^2 + \omega_B^2\sigma_B^2 + \omega_A\omega_B\sigma_{AB}$.

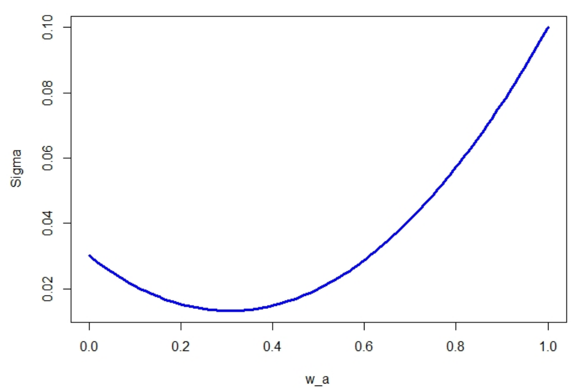

Now we have something interesting! Because the two portfolios have *negative* covariance, then some combination of the two assets gives me *strictly lower* portfolio variance than choosing either independently! Cool! The benefits of diversification in action!

Okay, so now we know how different weight combinations give us different risk and reward profiles. What we're going to do now is plot *both* of these profiles *on the same graph*. The idea is that, for each weight choice, we have a corresponding expected return *and* expected variance. So we can plot *both* of these *on their own axes* so that we can see what sort of risk-reward portfolios are available! I've done that below: for each weight pairing, there is an expected variance, which is on the x-axis, and an expected return, on the y-axis: Cool figure right! Very funky.

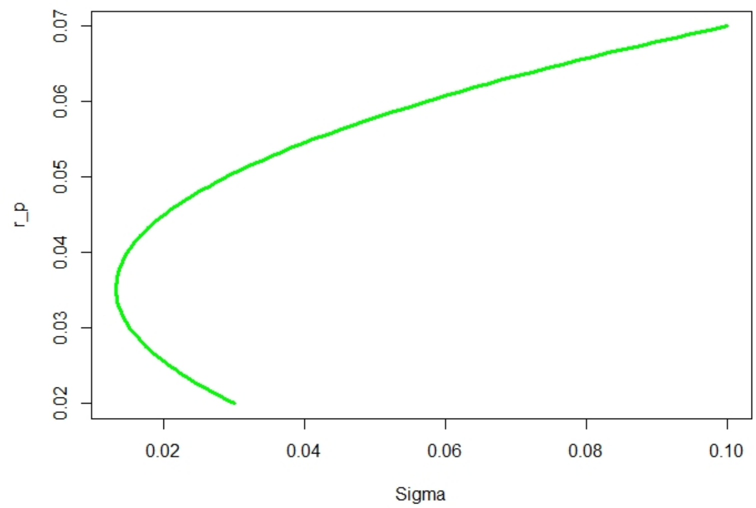

But what does it mean? Every possible portfolio, achieved by varying $\{\omega_A, \omega_B\}$, is represented by this line. Thus, this tells us *which combinations of risk and reward we can achieve*. What can we *learn* from this graph?

Well, we know that two portfolios we can always choose is just to have entirely either asset A or asset B. These two portfolios are represented by the top end of the curve and the bottom end of the curve respectively[1]. So, we know that travelling along this curve from the bottom up means *increasing* the weight of Asset A in our portfolio. Okay, cool! What else can we learn.

Suppose we wanted to achieve a portfolio that had a variance of around 0.02. How many such portfolios meet this criteria? Well, we can see that for an 'x-axis'

[1]If this isn't clear, note that where the curve ends at the bottom, the expected 'Sigma', or portfolio variance, is 0.03, which is just the variance of Asset B, and the expected *portfolio* return, r_P, is 0.02, which is just the expected return of Asset B. The same logic applies to where the curve ends at the top but for Asset A

value of 0.02, there are *two* combinations of Asset A and Asset B that deliver this level of portfolio risk. *But*, we learn something *stronger* than that. If given the choice between the portfolio indicated by the lower intersection of the curve and the x-axis value of 0.02, and the higher intersection of the curve and the x-axis value of 0.02, which would you pick? Of course the higher!! You get the *same* level of portfolio risk, *but strictly higher expected returns*! Huge!

So, one of the first things we learn from looking at this graph, is that there are certain portfolios *we should never pick*. We can maintain the *same* level of risk, *but achieve higher expected returns!*. So, every portfolio combination corresponding to the *downward sloping* component of the curve, should be avoided: it is *inefficient*.

What about the portfolios in the upward sloping portion of the curve? Whilst we can achieve higher returns by increasing our weight on asset A, *we will also have a riskier portfolio*. There are no 'Pareto gains' —I have to pay a cost to increase my expected returns, which is risk. *Which portfolio you will want on this upward slop* ***will depend on your preferences and attitudes towards risk***, and so there is no portfolio on this slope that is 'better' than any other, simply different.

3.2 The same exercise with real data

It's always cool to see these things applied with real data, so I went away and did some coding work :) I pulled data on Apple and Exxon Mobil. I think we all know the company Apple, but in case you're unfamiliar with Exxon, it is a major oil company. I picked these two stocks, because I anticipated that they would comove negatively. Why? Because Exxon benefits when oil prices go up, and Apple benefits when prices go down.

I start by calculating their average returns, and the variance of their returns, using data from 1999-2022. I then calculate the *covariance* of their returns. Here's what I found:

	Return, r	*Variance*, σ_i^2	*Covariance*, σ_{AB}
Apple	0.0143	0.0157	-0.0011
Exxon	0.0054	0.0026	-0.0011

No joke, these were the first two companies I tried, and it bloody worked!! Well pleased!! Okay, so, now let's repeat those figures that we constructed before:

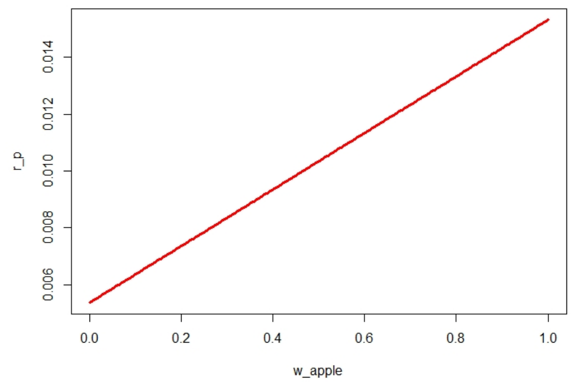
r_p
0.006
0.008
0.010
0.012
0.014
0.0
0.2
0.4
0.6
0.8
1.0
w_apple

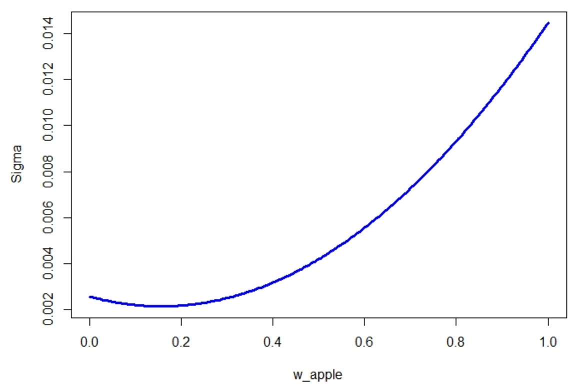
Sigma
0.002
0.004
0.006
0.008
0.010
0.012
0.014
0.0
0.2
0.4
0.6
0.8
1.0
w_apple

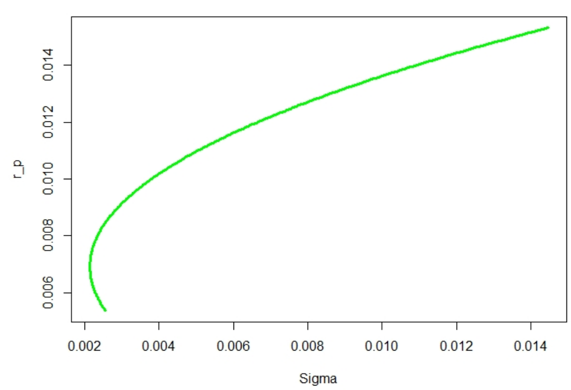
r_p
0.006
0.008
0.010
0.012
0.014
0.002
0.004
0.006
0.008
0.010
0.012
0.014
Sigma

How cool is that?!? It's amazing when you look at real data and get the same thing as when you make it up :) Okay, so again, what is this saying? If your choice is *only* between Apple and Exxon, *it is inefficient to only buy Exxon.* You can achieve the *same* portfolio risk, *with a strictly higher expected return* by diversifying your portfolio and buying Apple stock as well.

3.3 More than two stocks

Okay, we basically never have a portfolio of only two stocks, so how do we apply the same intuition to the case where we have several stocks? This is a little harder... In the previous case, we only really had a single choice: the value of ω_A. Now, if we have $\{\omega_1, \omega_2, ..., \omega_n\}$, we have *a lot* more flexibility. How do we proceed?

The first thing to note, is that the minute we move up to three stocks instead of two, we would have to plot 3-dimensional graphs to illustrate the potential returns. That's because we know have *two* free parameters rather than *just one* as before. Eww!! 3D?!? However, if we move from three stocks to four stocks we're really in trouble. Unfortunately we cannot construct graphs in four dimensions, our pitiful human brains just cannot comprehend them.

Anyway, I'll repeat the exercise for three stocks above so that we can get a feel of how moving up the number of stocks may make the calculation more difficult, *but doesn't really change the intuition*, and that's the key thing. Let's suppose that we add an extra stock to our portfolio. Their average returns and variance are as follows:

	Return, r	*Variance*, σ_i^2
A	0.07	0.1
B	0.02	0.03
C	0.03	0.04

And their covariance matrix:

$$\Sigma = \begin{array}{c|ccc} & A & B & C \\ A & 0.1 & -0.05 & 0 \\ B & -0.05 & 0.03 & -0.005 \\ C & 0 & -0.005 & 0.04 \end{array}$$

What expected returns can we achieve by varying our weights? Well note that whatever we choose for $\{\omega_A, \omega_B\}$ will determine what our choice of ω_C is. So we can choose $\{\omega_A, \omega_B\}$ freely, so long as they don't add up to more than 1, and then whatever we choose, $\omega_C = 1 - \omega_A - \omega_B$. So, we'll have two axes, one for ω_A, and one for ω_B, and then the returns will be the *3D surface* that these two choices generate. Cool right?

So it turns out that actually generating this plot is kind of hard... I tried for like an hour and have given up! I can't generate the 3D plot of Σ either. BUT! I *can* plot the various combinations of expected portfolio returns and expected portfolio variances *very easily.* **Plus, it looks incredibly awesome**. In the figure below, I plot all these combinations as before. I also include *overlaid on top* a dotted line in *yellow* that corresponds to the combinations we identified in the two-stock case.

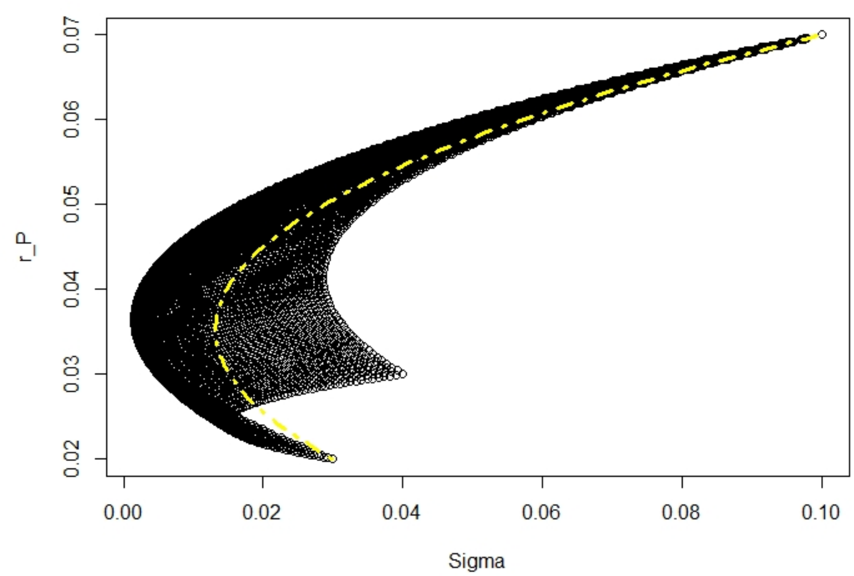

WOW! Isn't that cool??? Now that we have three stocks to choose from, rather than just a single *line* of potential portfolios, *there is a whole shaded area.* Any point in this area represents *a potential portfolio*: i.e. some combination of weights, $\{\omega_A, \omega_B, \omega_C\}$. What we can also notice here is that the 'pointy' parts of this area represent the simple portfolios where we just hold one of the assets on its own. By blending the assets together, we can access *any* combination of risk and return in the shaded area. Awesome right??

So which portfolios are *inefficient*? Well, **almost all of them**. The *only* portfolios we should accept are those *right at the top of the ascending section of the curve*. These portfolios guarantee the *highest* expected return *for a given level of portfolio risk*. This regions is known as the *efficient frontier*. Also note that adding the extra stock to my portfolio *extends* the efficient frontier: by constructing a portfolio that contains *all three* stocks, I can generate *strictly* higher expected return for the same level of expected risk. Winner!

It is worth noting that as you increase the number of stocks, you will just change the shape of the shaded area: moving from $n = 2$ to $n = 3$ is the big shift. Moving from $n = 3$ to $n = 100$ will result in a different *shape* of the shaded area, but it will still have the *same interpretation*.

Okay, so we've just seen how, when we consider *both* the risk *and* the return profile of portfolios, we can learn something about which portfolios we should *never* pick. We then established an *efficient frontier* of portfolios, and claimed that *any portfolio on this frontier is acceptable*, so your choice will depend on your risk attitudes. Hmm... can we do better than that?

3.4 Borrowing and lending

Suppose I am also able to borrow and lend at some risk-free rate: call it r_f. Then I can always combine any portfolio available to me with some holdings in this risk-free asset. Suppose that I've chosen a portfolio, P from the efficient frontier above. I want to combine it with some amount of the risk free asset, call it a_{rf}. Suppose the share that goes into the risky portfolio is ω. Then, my *total* portfolio, $\hat{P}$, would just be:

$$\hat{P} = \omega P + (1 - \omega)a_{rf}$$

What are my expected returns? Very simple:

$$r_{\hat{P}} = \omega r_P + (1 - \omega)r_{rf}$$

What about the expected variance? Recall:

$$\sigma^2_{\hat{P}} = \omega^2\sigma^2_P + (1 - \omega)^2\sigma^2_{rf} + \omega(1 - \omega)\sigma_{rf,P}$$

But hang on... this asset is risk free!! So it *can't* have *any* variance. If it has variance, then *it must have some risk*. So: $\sigma^2_{rf} = 0$. Similarly, if there is *no risk*, then it *cannot* be the case that the returns on the risky asset *could possibly* co-move with the returns of the *risky* asset. So, it also follows that: $\sigma_{rf,P} = 0$. Putting all of this together, we can show that:

$$\sigma^2_{\hat{P}} = \omega^2 \sigma^2_P$$

Recall that the standard deviation of the variance is just the square root of the variance. Thus, the standard deviation of our portfolio is just given by:

$$\sigma_{\hat{P}} = \omega \sigma_P$$

In other words, the standard deviation of our *total* portfolio, $\hat{P}$, is a *linear function* of how much weight we put in our *risky* portfolio, P. If we put nothing in the risky portfolio, then $\omega = 0$, and $\sigma_{\hat{P}} = 0$. If we put it all in the risky portfolio, then $\omega = 1$, and $\sigma_{\hat{P}} = \sigma_P$.

3.4.1 Why are you telling me all of this?

Suppose we were to modify slightly the figure above for three stocks that showed us every combination of expected portfolio *variance* and expected portfolio return, so that instead of variance, we have standard deviation. It would look something like this:

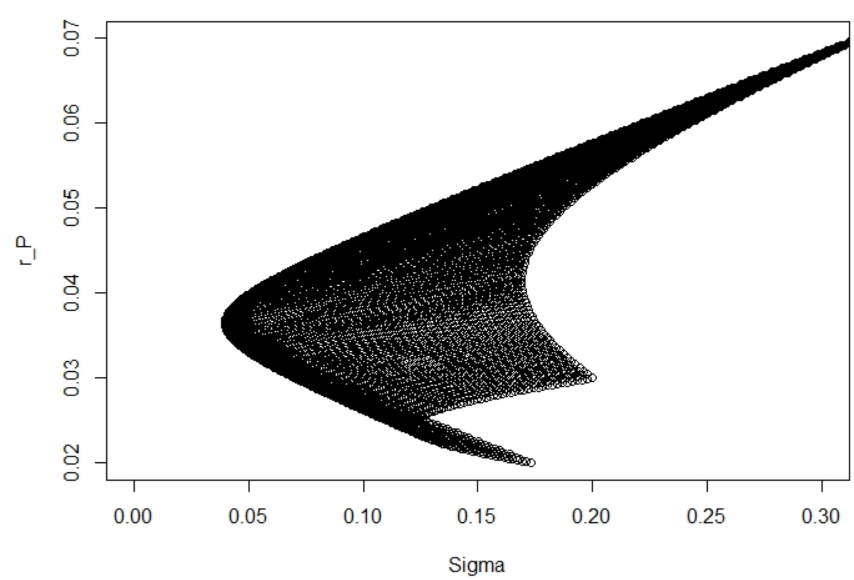

Crucially, note that nothing meaningful has changed about the efficient frontier, we have just rescaled the measure of 'risk' that we're using. This representation would, for the same attitudes towards risk, *produce the exact same rankings* of portfolios according to risk and return. What *has* changed, is that we can now introduce *graphically*, the idea of blending a portfolio *on the efficient frontier*, with the risk-free asset. If this isn't clear just yet, *bear with me*!

Suppose that the risk-free asset pays a rate of return equal to 3%. Then take a portfolio on the efficient frontier, and draw a straight line from that portfolio to the point where the return is 3%, and the σ is 0. I've done that for four potential portfolios here:

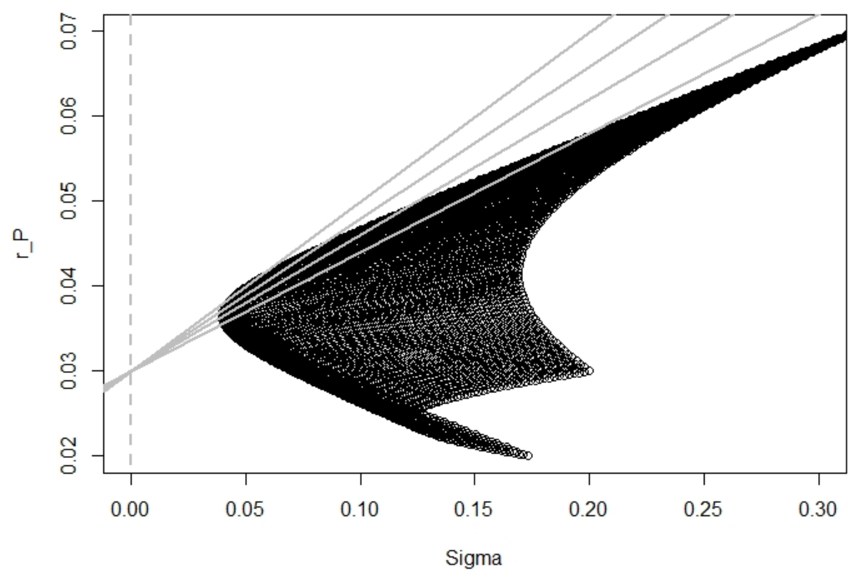

Why have I done that? Well, recall that we showed that the return of a blend of an efficient frontier and a risk free asset was given by:

$$r_{\hat{P}} = \omega r_P + (1 - \omega) r_{rf}$$

And the standard deviation was:

$$\sigma_{\hat{P}} = \omega \sigma_P$$

Note that the *rate* at which the expected *return* increases as I increase the value

ω *is exactly the same* as the rate at which $\sigma_{\hat{P}}$ increases???? Can you guess what this means??? **It means that every point on the grey line between the frontier and the y-axis is a feasible portfolio**. I just *lend money* at the risk free rate as part of my portfolio, and I can blend that with my risky portfolio!

The fun doesn't stop there. Suppose instead of *lending* at the risk free rate (equivalent to *buying* the risk-free asset), I instead *borrow* at the risk free rate. What this means is I *issue* a risk-free loan, and use that to *buy more of the Portfolio.* This would be the equivalent of setting $\omega > 0$. I borrow money at the risk free rate, *and invest it in* P. This allows me to access all the points on the grey lines *above* the efficient frontier! Amazing! But why would I want to do this?

3.5 The 'best' portfolio

Note that, for many of the portfolios in the example above, the grey line contains *inefficient portfolios*: I could get a better risk-reward profile by ignoring the risk-free asset and just readjusting the weights of the risky stocks. However, for one *very special* portfolio, this is not true. Consider the following figure:

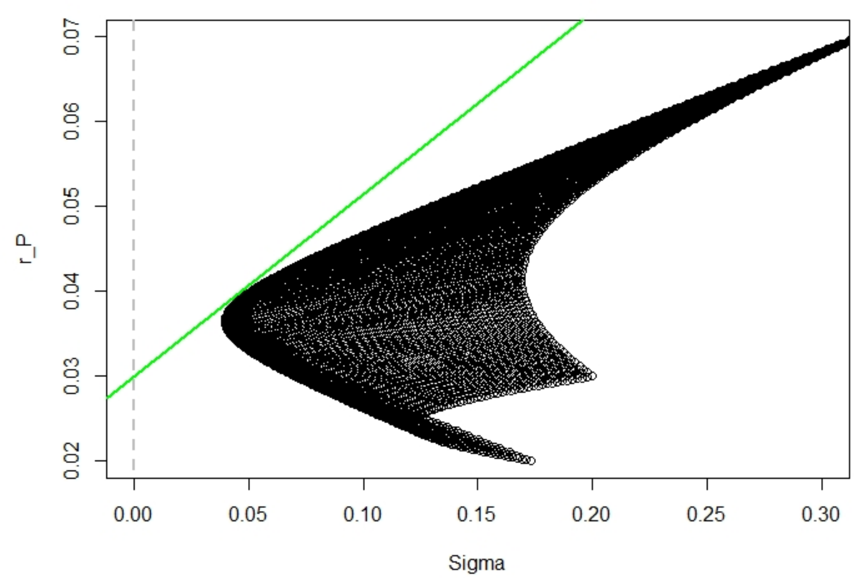

For this *very special* portfolio, **also known as the market portfolio**, every single portfolio I could construct by either borrowing or lending using the risk free asset is *above* the efficient frontier. This portfolio is the one on the efficient frontier that forms a *tangent* when a line is drawn from it to the risk-free rate. Thus, *every portfolio on this line is more efficient than any portfolio I could make by using only the stocks available*. This is a **huge** result. By using the risk-free asset, I can access *a much better risk-reward profile*. I'm also able to access *any* level of risk that I like. This is a **cool result**.

3.5.1 Do investors do this in practice?

Yes!! On the lending side, pretty much every institutional investor will hold some risk-free assets to manage their portfolio. On the borrowing side, some investors perform what is known as *buying stocks on margin*: they borrow money from their brokerage in order to invest more in their portfolio. When done effectively, this can access higher expected returns for an increase in risk than could be achieved by simply selecting other stocks. Of course, buying on margin can also be a catastrophe if the *true* returns turn out not to be so great... remember, these returns are only *expected*: they could be negative!

3.5.2 Summary and caveats on borrowing and lending

This concept is a little bit tough to grasp, but once you've got it, I hope you'll see how cool this is. We started off by saying: which portfolios are *bad* choices. We ended up outlining a *set* of portfolios that are *good* choices. We then introduced the idea of *borrowing and lending* at the risk-free rate. With this tool, we can blend our portfolio between any efficient portfolio and our risk-free asset. We then showed that, *if this is an option*, then there is only *one* portfolio on the efficient frontier that we should choose! The one that forms a tangent to the efficient frontier when we draw a straight line to the risk-free rate.

Whilst this is very nice in theory, there is a major caveat that is absolutely worth emphasising: **If you are not the US government, you cannot borrow at the risk-free rate**. Whilst it's nice to imagine a theoretical model in which I can borrow and lend at the risk-free rate, this is not true in practice. I can for sure lend at the risk-free rate: I just buy a US T-bill. But I cannot issue risk-free assets (i.e. issue a bond/ ask for a loan) as an independent investor, *because I am risky!* You will not be able to convince anyone that lending to you is as free of risk as lending to the US government. As such, the idea that we can have a *straight line* is bogus in practice.

4 The Market Portfolio

In the discussion above, we identified one *and only one* portfolio that, in the presence of borrowing and lending, should be selected. We called this **the market portfolio**. This portfolio displays the highest possible return *for every risk level*—we *cannot* do better than this portfolio, at least in theory.

4.1 The Assumptions

What assumptions prop up this theory? Critically, this theory holds *only if everybody has access to the same information, and interprets that information in the same way.* If we all agree on the expected returns *and* the expected variances, *and* the expected co-variances, of every stock in the market place, *then we have to agree on what is the best portfolio.* In practice, this assumption is unlikely to be true.

4.2 Why the 'market' portfolio?

Note that for this portfolio to have the property that it achieves the highest return for a given level of risk, *the converse must also be try*: the portfolio delivers the *lowest* risk for a given return. For this to be true, it must be that the portfolio is *well diversified*, so that *all* idiosyncratic risk has been eliminated, or as close as is possible. *It can be shown that this portfolio will contain non-zero amounts of every available stock*, but showing this is a little beyond the scope of this class. However, the intuition is still there, right? If we can lower idiosyncratic risk by adding more stocks, we should add as many as we can.

So, this portfolio *contains every stock in the market.* As such, 'market portfolio' is an appropriate term for this portfolio.

4.3 No idiosyncratic risk... but the risk is still non-zero?

That's right! This is the *market* risk we talked about above. There are some factors that *cannot be diversified away.* These factors, as we saw above, *tend to affect most stocks in the same way* —the β values of most stocks are greater than 0.

So, what is the β of the market portfolio? It should be one! It would be very strange if it weren't. To confirm this, I took the same data I used to calculate the individual stock βs, and just pooled it all together. What I'm now doing is asking, how much do *all* the stocks returns covary with the *market* return we'd achieve if we built a portfolio containing all stocks? Here's what I find:

Dependent Variable:	$r_{i,t}$
Model:	(1)
Variables	
(Intercept)	0.0070*
	(0.0042)
$r_{m,t}$	1.002***
	(0.0042)
Fit statistics	
Observations	544,258
R^2	0.09335

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 1: Caption

I love it when a plan comes together!

5 The Capital Asset Pricing Model (CAPM)

In the last section, we brought together the notion of portfolio risk and portfolio return and showed how they relate to one another. In particular, we showed that there are *inefficient* portfolio, and *efficient* portfolios. An inefficient portfolio is one where we can achieve a *strictly higher* expected return by rebalancing our portfolio, *with not change in the expected risk*.

We also showed that there was *one* portfolio, the *market portfolio*, that is the unique best *if* we can access risk free borrowing and lending. Recall our discussion of β? β was a measure of how much a given stock moves with the *market*. Our market portfolio should, therefore, have a β that equals **one**. What about the risk free asset? Well, as that asset has *no* variance, it doesn't move *with anything*. So its $\beta = 0$. What about other stocks/portfolios?

First, recall the notion of a risk premium: for stock i, the *risk premium* is the additional return that the stock delivers to *compensate us for risk*. So, if we let r_{rf} be the return of the risk free asset, then:

$$Risk\ Premium_i := r_i - r_{rf}$$

What is the risk premium of holding the risk free asset?

$$Risk\ Premium_{rf} = r_{rf} - r_{rf} = 0$$

Okay, sensible! What about the market portfolio? Let the return of the market portfolio be given by r_m; then:

$$Risk\ Premium_m = r_m - r_{rf} \neq 0$$

Note that the risk premium *of the market portfolio* reflects *only* market risk. By definition, this portfolio has effectively eliminated idiosyncratic risk.

So, a natural question we might ask is, if I have a portfolio *other* than the market portfolio, and hence $\beta \neq 1$, then what does the expected risk premium look like in that case? Note that this portfolio may contain idiosyncratic risk if it is not well-diversified! Does the idiosyncratic risk, or the market risk of the portfolio matter? It is this exact idea that motivates the development of the CAPM.

5.1 Overview

The CAPM was developed by William Sharpe, John Lintner, and Jack Treynor in the 60s. Their goal was precisely to answer the question of what the risk premium of portfolios *other* than the market portfolio ought to be. It turns out that, under certain assumptions, the *theoretical* relationship between the risk premium of a portfolio and its β is just:

$$r_i - r_{rf} = \beta_i(r_m - r_{rf})$$

That is, the expected returns of portfolio i is just the β of portfolio i, *multiplied by the market premium.* In other words, *the only risk that is rewarded by the market is market risk.* You should not expect to be compensated for taking on idiosyncratic risk.

I'm not going to go through the derivation, as it's a little too complicated and isn't necessary to build intuition. For that, consider the following figure:

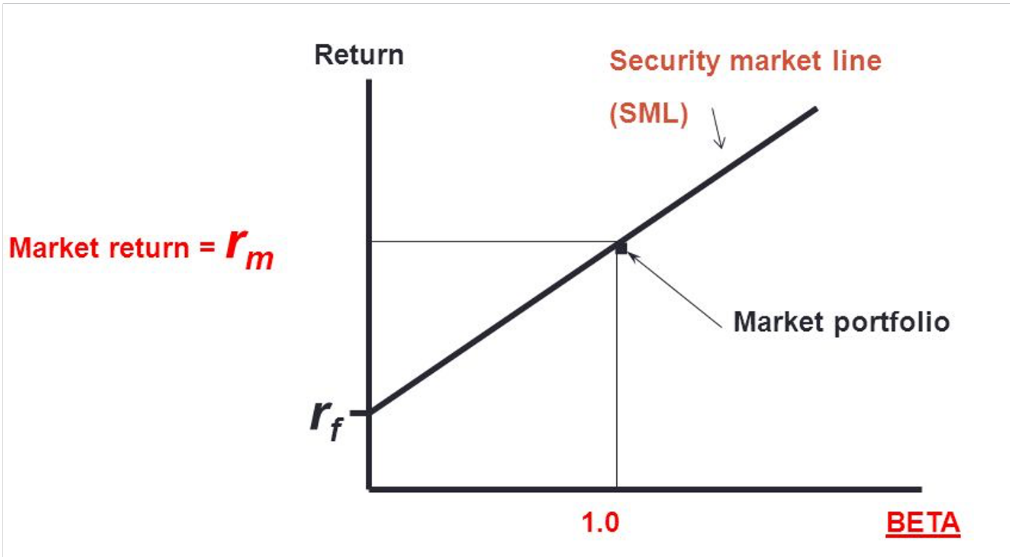

Figure 3: Source: Brealey, Myers, and Allen—Principles of Corporate Finance, 10th Edition

Here we have plotted the hypothetical model of Sharpe, Lintner, and Treynor. What this model says is that the risk *premium* is a linear function of the β of the stock. For the market portfolio, this $\beta = 1$, and the y-axis correponds to the market return. For the risk-free asset, the $\beta = 0$, and the y-axis corresponds to the risk free return. The relationship is captured by the *security market line*, or SML—the theory tells us every stock should lie on this line.

Now suppose we consider some alternative portfolio, i. Suppose this portfolio has $\beta = 0.5$, but the expected return lies *below* the security market line. Would you buy it? Well, if I can lend at the risk free rate, then I can mix the market portfolio with the risk free asset, guarantee a $\beta = 0.5$ for my portfolio, *and ensure a higher reward*. So no one would want to buy this portfolio, and so its price would have to go down. As its price goes down, the returns go up, until we eventually hit the SML.

Now suppose another alternative portfolio, j. Suppose this portfolio also has $\beta = 0.5$, but the expected return lies *above* the security market line. Would you buy it? Absolutely! This is better than I can get by mixing between the market portfolio and the risk free rate! So, everyone rushes to buy this asset... meaning the price goes up, and the return goes down.

Hopefully you can see that the intuition here is *very similar* to the notion that risk and reward are related. However, the CAPM says something even stronger than this. The CAPM tells us that *only* market risk matters in the relationship between risk and reward. The return profile of an asset is *only* a function of market risk, because there is *no* idiosyncratic risk in the market portfolio. Cool huh?

5.2 Some additional intuition on this result

In many ways, it would be a little strange if the market rewarded us for taking on idiosyncratic risk. Because we can always diversify away from idiosyncratic risk, its presence in our portfolio is sort of *unnecessary.* We don't usually reward *unnecessary risks.* For example, you are no more likely to win the lottery if you buy 10 tickets and destroy 9 at random. You have introduced an unnecessary risk to your chance of winning the lottery.

5.3 Putting CAPM to the sword!

The prediction of CAPM is a strong one. Its also one we can test! Suppose we go back to all those β_i's that we estimated before, and we make 10 portfolios, where each portfolio contains 500 stocks, and they're selected according to the decile of their β, i.e. portfolio 1 is the 1st decile, portfolio 2 is the 2nd, and so on. I label each portfolio with the notation, $\{\beta^1, \beta^2, ..., \beta^{10}\}$. Then we say, how will did these portfolios do, *over time?*

To do this, I first estimate all 5,000 β_i's from 2000 to 2016. I then form my portfolios, and see how well they do: I calculate their total *cumulative* return from 2017-2021. I also calculate the *marker* return over the same period. I'll then plot the actual return and the $beta^i$ of each portfolio. What do I find?

So here is a plot of the median estimated β of the 10 portfolios, plus the market portfolio, on the *cumulative returns* of each portfolio if you were to buy it at the start of 2004, and hold it until the end of 2007. Overlaid is the *prediction* of the CAPM: every portfolio should lie on that straight line. To find this line, I need to find the risk free rate. I do this by finding the annual T-bill return for the years 2017-2021 (17.6%), and then draw the line between this rate and the market portfolio where $\beta = 1$.

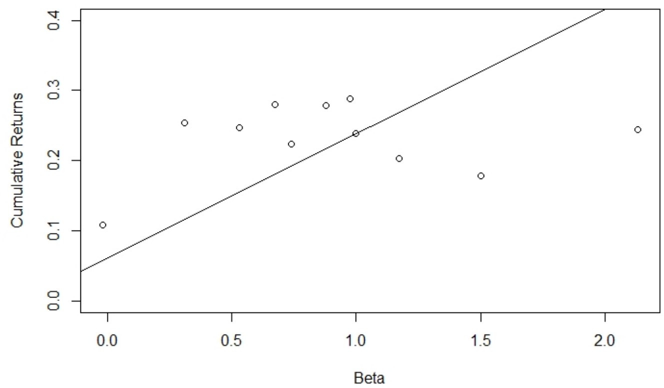

So what's the verdict? Hmm, not great... It seems as though the CAPM model *underestimates* the returns of portfolios that have $\beta < 1$, and *overestimates* the returns of portfolios that have $\beta > 1$. If we go further, and run a regression to see if β_i is a good predictor of future returns, the results look even worse. Suppose I run the following regression:

$$r_{i,t} - r_{rf,t} = \alpha_0 + \alpha_1 \beta_i$$

If the CAPM is true, what should α_1 be? Well it should be roughly equal to the *market risk premium*[2]. Looking at our figure above, we can see that number is around $0.24 - 0.05 = 0.19$. Similarly, α_0 should equal 0. What do we actually find?

Dependent Variable:	$r_{i,t} - r_{rf,t}$
Model:	(1)
Variables	
(Intercept)	0.8289***
	(0.0007)
β	0.0011*
	(0.0006)
Fit statistics	
Observations	91,829
R^2	3.51×10^{-5}
Adjusted R^2	2.42×10^{-5}

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

[2]This comes directly from the model

We are way off! We can *just about* say that the relationship is not zero, but barely. It gets worse: in the past we've suggested that *past* returns might be a good predictor of *future* returns. Suppose we include *both* β *and* past returns, r_i^{past} in our regression?

$$r_{i,t} - r_{rf,t} = \alpha_0 + \alpha_1 \beta_i + \alpha_2 r_i^{past}$$

Dependent Variable:	$r_{i,t} - r_{rf,t}$
Model:	(1)
Variables	
β	0.0009
	(0.0021)
r_i^{past}	0.1390*
	(0.0705)
Fixed-effects	
date	Yes
Fit statistics	
Observations	91,829
R^2	0.05506
Within R^2	0.00034

Clustered (date) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2: Caption

Oh dear! What do these results say? They say that once we include past returns as a potential predictor, *we can't even claim that the relationship between β and future returns is different from zero.* Even *worse*, it seems like past returns *actually may* predict future returns!

So what gives? Why is this such a famous theory, so much so it may well be the *only* theory any of you had heard of before you started this course?

5.4 A history diversion

Using the data from 2000-2016 to construct the β_i's for each stock, and then seeing how well these β_i's correlate with future returns from 2017-2021, we seem to reject the CAPM. It doesn't look like a stock or portfolio's β is a good predictor of its returns.

However, if we were to perform the same exercise using data from *much further in the past*, say, from 1930-1970, we might find something altogether different:

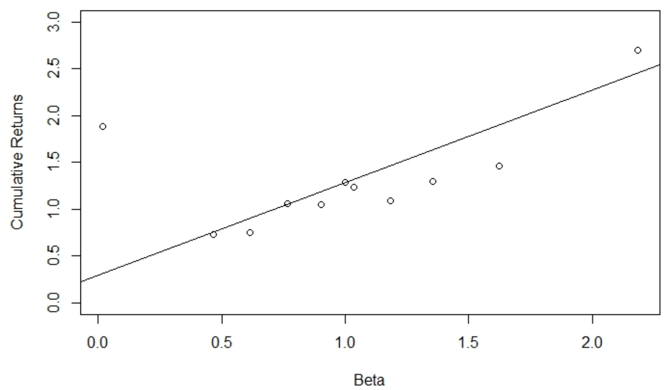

That is a *much* better fit. Apart from the portfolios with very low β, the fit is *remarkable*. Comparing the performance of β to past returns by running the same regression as above, we find more striking evidence:

Dependent Variable:	$r_{i,t} - r_{rf,t}$	
Model:	(1)	(2)
Variables		
(Intercept)	0.0174***	0.0076
	(0.0006)	(0.0184)
β	0.0040***	0.0040***
	(0.0005)	(0.0005)
r_i^{past}		0.0097
		(0.0181)
Fit statistics		
Observations	92,811	92,811
R^2	0.00059	0.00059
Adjusted R^2	0.00058	0.00057

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

So using this old data, the CAPM looks very very good. Okay, the actual *coefficient* is a little off (the actual slope should be much bigger), but the *basic* intuition holds very well indeed.

5.5 Temporal Validity

Note that at the time of its construction, the dataset that Sharpe, Lintner, and Treynor were using was essentially the older data we just used. So, at the time, it seemed like they'd really nailed it. This is why the model was so famous and has been so influential.

One remarkable feature of finance and economics, is that models that had been 'proven' correct by the data at some point in time, turn out to be rejected by the data at other times. This could be for two reasons:

- The original analysis was poorly done, at least relative to modern standards of rigor.

- The world has changed.

In the case of the former, we can simply conclude that we didn't *actually* prove that the model was correct, we were just consulting tea leaves. In the case of the latter though, we have a real problem. The concern that our models may only have validity *at a specific point in time* is called a concern of *temporal validity*[3].

Without more investigation, it's hard to say which reason dominates in our example. Suffice to say, the profession recognised the breakdown of the result and have sought to find improved models that can perform better where CAPM fails. Nonetheless, it is an important note to emphasise that, as with all things, *intellectual humility* rarely goes unrewarded. We should always be conscious that, when it comes to finance, *it is extremely difficult*, perhaps impossible, to build a model that is true both now *and* in the future. Perhaps we are constantly engaged in a Sissyphean task of modelling and remodelling our understanding of how financial markets work, *simply because they are always changing.*

One natural mechanism by which this change occurs is that, once the CAPM was developed and popularised, *people started trading with it!* This is a *huge* change in *how* people buy stocks. With this in mind, its perhaps not surprising at all that the principal effect of the development and dissemination of the CAPM was simply to sow the seeds of its own destruction.

[3] Cool phrase right?

6 An Alternative: The Three Factor Model

Okay, that's enough history! How have people improved on the CAPM, now that we have seen how poorly it seems to describe reality?

The great successor to CAPM is the 'Three Factor Model', henceforth referred to as **TFM**. This model was developed by two economists, Eugene Fama and Kenneth French. Eugene Fama is an absolute *legend* in the field of finance (though a controversial one! He is very much a believer in the efficiency of markets!), and he won the Nobel prize in 2013 in part because of this model.

The TFM can basically be thought of as CAPM+. The only change to CAPM is to include two *additional* variables:

- **SMB**: Stands for 'small minus big'
  - This is a measure of the difference in the returns of small market cap stocks versus the returns of big market cap stocks.
- **HML**: Stands for 'high minus low'
  - This is a measure of the difference in the returns of high book-to-market stocks versus the returns of low book-to-market stocks.

So, in the TFM, we posit that the risk premium on stock or portfolio i is given by the following expression:

$$(r_{i,t} - r_{rf,t}) = \beta_{market}(r_{m,t} - r_{rf,t}) + \beta_{SMB}(r_{s,t} - r_{b,t}) + \beta_{HML}(r_{h,t} - r_{l,t})$$