Topic 11: The Efficiency of Markets

Eddie Shore

Corporate Finance – ECONS4280

August 8, 2022

In this set of notes we discuss the concept of market *efficiency*. We begin with a reminder of what the price of a stock '*should*' be, with a focus on the role of *expectations*. We then offer a definition of the *efficient market hypothesis*, noting that there are different *degrees* to which this hypothesis can hold. We discuss why market efficiency is important, particularly in finance. We then show evidence that both supports and rejects the EMH.

1 Why do we care about market efficiency?

At two major points in the course, we've justified our arguments using the claim that *market forces* squeeze profitable opportunities out of availability. The first example was in our discussion of risk and return:

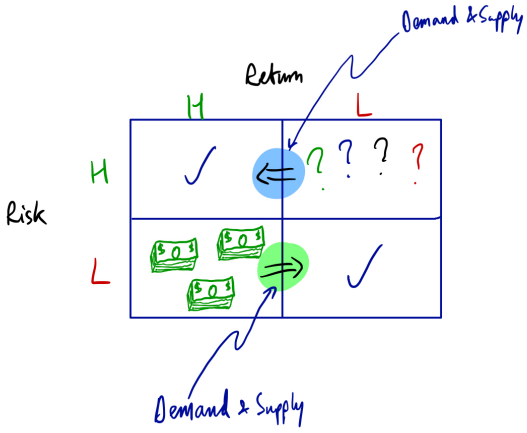

We argued that risk and return are *two sides of the same coin*, because any projects in the off-diagonals (upper right and lower left) will either have too little demand, or too much demand. As such, the *market forces* will push projects in these quadrants into the diagonals (upper left and lower right).

We used a similar argument to justify CAPM: we said that, because we can invest in the market portfolio and then blend that portfolio with the risk free rate, we can access the **security market line**. Any portfolio or stock above that line will be so desirable that everyone will buy it, so the return would have to come down. For every stock below that line, no one will buy it, so the returns will have to go up!

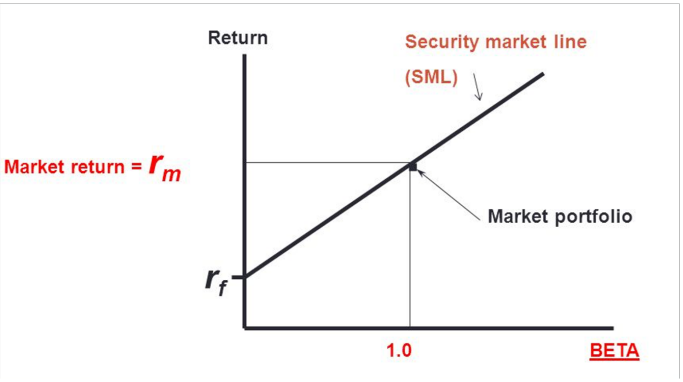

In both cases, we invoked the power of the market to quickly squirrel away these opportunities or anomalies. In fact, we did more than that: we invoked the **efficient market hypothesis**. Before we get to that though, let's remind ourselves of where the price of a stock comes from.

2 What is the price of a stock?

Recall that we found that the price of a stock was just equal to the *expected*, discounted flows of future dividends?

$$P_t = \mathbb{E}_t \left[\sum_{j=1}^{\infty} \left(\frac{1}{1 + r_{t+j}} \right)^j DIV_{t+j} \right]$$

We arrived at this result from a very simple and *mechanical* reworking of an

accounting *identity*. Recall that this result follows directly from:

$$r_{t+1} = \frac{DIV_{t+1} + P_{t+1} - P_t}{P_t}$$

So, in a sense, there is nothing 'assumed' about this framework for price, it is just true *by definition*. Why does that matter? Well it means that the expression above *must* be the true price, a true reflection of the value of a stock. But note that it *relies heavily on how we form expectations*.

If everybody forms expectations in the same way, then everyone should have the same beliefs over what the price of a stock is. Recall that expectations are just a function of two things:

- Our understanding of the system that governs random processes.
- The information that we use to form our beliefs.

So beliefs could differ if and only if there was some disagreement on the system that governs random processes, *or* some people have access to different sets of information, *or* both. In essence, what the efficient market hypothesis argues is that *none of these things can be true in equilibrium*.

3 Defining the Efficient Market Hypothesis

The efficient market hypothesis, or EMH, is defined in the following way:

Definition (The Efficient Market Hypothesis). In an efficient market, prices reflect **all** available information.

The *degree* of efficiency in a market depends on two dimensions:

- The *type* of information that is incorporated into price, i.e. *which information is available?*
- The *speed* with which new information is incorporated into price, i.e. *how fast is information reflected?*

Suppose that we assume financial markets are *not* efficient. That means that they either do not incorporate all information, or the information is processed slowly. What would that mean for traders? It would mean there were opportunities to profit using information relevant for pricing but not yet reflected in prices!

3.1 An example from British history

The Rothschild family is possibly the most influential in the history of modern Europe. As the family that bankrolled the British during the Napoleonic wars, the Rothschilds tilted the course of European history towards Britain, and have played a sizeable role in British history ever since. The course of this family's rise to extraordinary wealth and influence is a fascinating narrative, and one I strongly recommend that you learn about[1]. Their story also provides a fascinating narrative on the *efficiency of markets* and the value of information.

Nathan Rothschild, son of the family patriarch Mayer Rothschild, was the founder of the London branch of the Rothschild firm[2]. By all accounts, Nathan was a brilliant businessman, and had a keen understanding of the value of information in the context of finance. He made his fortune initially by supplying the Duke of Wellington's army in Spain and France with gold and silver to pay the troops.

Rapid and reliable communications were crucial for his complex and risky payments and arbitrage operations. He set up a private courier system with shipping agents in Dover, Calais, and Ostend, with fast light vessels ready to sail at any time. There were relays of horses to speed messages from the Channel to London, and even a farm on the coast of Hythe for courier pigeons.

As night descended at Waterloo on Sunday, June 18th, 1815, a Rothschild agent dashed to Dunkirk. Conveyed by Rothschild ship and Rothschild steeds, Nathan received news of the victory on the night of Monday 19th, just 24 hours later. Wellington's official messenger did not arrive until Wednesday evening! Nathan proceeded to the Prime Minister's residence, but was refused entry by a butler as the PM was sleeping.

His duty done, Nathan proceeded to the Stock Exchange where he alone knew that the French had been defeated. It is not known exactly how much Nathan made from Waterloo, but their fortune went from £500,000 in the spring of 1815 to £1,000,000 by July 1816. In modern units, that's a shift from £31,000,000 to £68,000,000. Not a bad return for a day's work!

The Iron Duke observed that Waterloo was: 'A damn close-run thing —the nearest run thing you ever saw in your life'. Perhaps this was true on the battlefield, but not on the Stock Exchange with Rothschild's information system.

[1] Niall Ferguson's history of the Rothschilds, considered his greatest piece of *academic* work, is a great place to start

[2] This is the same Rothschild that still operates today!

3.2 Applying the notion to prices

So if the EMH holds, what does this mean for our price expression above? If a price contains *all* information, then could there be differences in beliefs over future dividends such that profitable activities of the kind associated with *asymmetric* information could flourish?

Yes! If some people understand how the system operates better than others, *they can take advantage of that information.* In other words, they could make *positive NPV* on buying stock prices, earning abnormal returns. But can this last in equilibrium? If you keep getting screwed because you have a bad system for thinking about how the random processes work, *how long will you stick around for?* Will you not learn that you're getting things wrong? Thus, the *pressure* of the market will squeeze participants into behaving appropriately, *lest they suffer continuous losses.* If a buyer is making a positive NPV by buying a stock, *it must be true that the seller is making negative NPV.* This is not sustainable *if everyone has the same information*, because you will *consistently* get things wrong.

Similarly, suppose you uncover a system that does better than everyone else, and you keep it a secret. In the short run, maybe you can eek out some positive NPV. But remember, if *everyone* has the same info, then *at some point* someone will work out what system you're using.

Think of it this way: suppose you're a clockmaker, and a rival manufacturer has just developed the most accurate watch anyone has yet made. You are fuming, as that used to be your title! You want to take it back. In a market with *perfect* information, then you would be able to see every piece of machinery that went into that clock! You could buy one, reverse engineer it, and make a copy of the design. It might be hard, but eventually, you would work out how they did it.

By contrast, if information is not uniform and all encompassing, we could imagine that fluctuations in the availability of information could maintain disagreements over prices that provide opportunities for positive and negative NPVs, *but we should think these should be as random as the random access to information.* In other words, it seems unlikely that someone is consistently accessing better information than everyone else, if it is indeed the case that not all information is publicly available. Nathan Rothschild did not make millions on *every* trade, he just got access to a couple of *crucial* bits of info.

Thus, the critical feature of the EMH relates to the notion of the *availability* of

information. If *everyone* has access to the same information, then *in the medium to long run*, no one should be able to consistently deliver positive NPV.

4 Categories of market efficiency

Note that within the definition of the EMH, we noted that the *degree* of efficiency depends on both the **type** and the **speed** at which new information is incorporated into price. There are typically three categories of market efficiency corresponding to different levels of these criteria:

- Weak-form efficiency
- Semi-strong efficiency
- Strong efficiency

Let's take a look at each of them.

4.1 Weak-form efficiency

The weakest form of market efficiency is that *prices reflect all information contained in market trading data*. That is to say:

- Past prices
- Past volume
- Past dividends
- Past interest rates
- etc.

What would this mean? If this level of efficiency holds, would you be able to consistently make profits by using past prices to identify mispriced securities?

No, you would not. If prices already *contain* the information from the past, then using the past to find *positive* NPV will not work! Everyone knows the stuff that you're using to trade on already.

To illustrate with an analogy, it would be like Nelson's messenger, arriving on Wednesday two days after the Rothschild's messenger, going to the stock market to

trade on the news of Waterloo. Everyone already knows! The prices have already gone up to reflect the news. So you can't hope to use information that is already available to everyone to *beat* everyone.

4.1.1 Evidence for Weak Efficiency

The evidence here is very strong: **financial markets display weak efficiency**. People have tried again and again and again to predict stock returns using past data, and it doesn't work. You can't just plug publicly available financial numbers into a computer, and expect to make money doing that.

The evidence that supports this claim often consists of tests of '*serial correlation*'. Serial correlation occurs in a time series when a variable and a lagged version of itself (for instance a variable at times t and at $t-1$) are observed to be correlated with one another over periods of time. If we believe stock returns have serial correlation, then the hypothesis is that returns at time $t+1$ are some function of past returns, perhaps in the following way:

$$r_{t+1} = \rho_0 r_t + \rho_1 r_{t-1} + ... + \rho_j r_{t-j} + \epsilon_{t+1}$$

So tomorrow's returns are some function of past returns, plus some shock, ϵ_{t+1}. When we test this framework what we find is:

$$\rho_0 = \rho_1 = ... = \rho_j = 0 \implies r_{t+1} = \epsilon_{t+1}$$

In other words, *there is no serial correlation of past returns.* Below is a table outlining the results of tests of serial correlation with just one lag:

USA	0.03	UK	0.08
France	-0.01	Italy	-0.02
Germany	0.08	Holland	0.03
Belgium	-0.02	Switzerland	0.01
Sweden	0.06		

Figure 1: Source: B. Solnik, "A Note on the Validity of the Random Walk for European Stock Prices." Journal of Finance (December 1973).

4.2 Semi-strong efficiency

Semi-strong efficiency suggests that prices reflect all *publicly* available information. That means:

- CEO Scandals
- The price of copper in Belarus
- Traffic data around corporate HQs
- National Championship Football game results

Again, what would this level of efficiency mean? It would mean that an investor *cannot use publicly available information about firms to pick stocks.*

4.2.1 Evidence for semi-strong efficiency

Most of the evidence around semi-strong efficiency relates to the speed of adjustment of prices to news. Here are a few examples:

Meta

In January of this year, Meta announced their annual earnings for 2021. Amongst many details of their release, they noted a drop in total Facebook users and sign ups. What does this suggest? That Facebook has *finished its growth trajectory.* How did markets react?

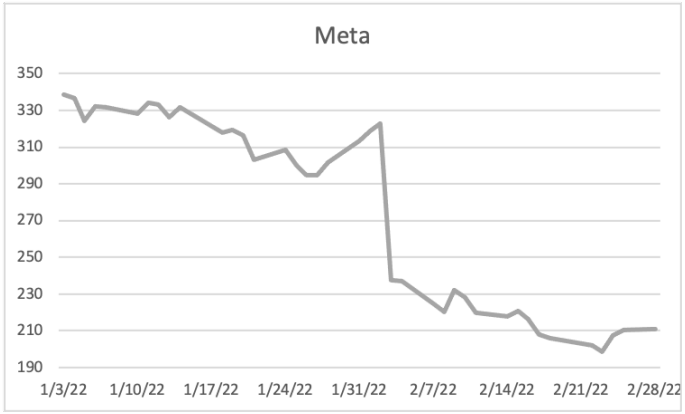

Figure 2: Meta stock price around earnings announcement

What we observe is a *very* rapid drop in the share price. In the space of a few minutes, Meta dropped from roughly $320 a share to around $235, a **26% loss** in value.

Recall why people own stocks? Either because they believe the value of the company will *grow*, or because the stock pays *income*. The news that Facebook was no longer attracting new users, and was even losing its existing ones, made it hard to argue that Facebook is a 'growth' stock. For Facebook to maintain value, it would need to start *paying dividends*. Historically, Facebook has paid a dividend of only 1.2%, considerably less than the risk-free rate associated with Treasury bills. Thus the market takes the view that, without a significant change in dividend policy, or some further sign of growth opportunities, *the value of owning a share in Facebook is now considerably less than it was before.*

Volkswagen

In 2015, Volkswagen became embroiled in one of the biggest corporate scandals to hit German business of all time. It was revealed by regulators (specifically, the US Environmental Protection Agency) that, since 2008, VW had been manipulating the emissions data on their diesel engines.

Beginning in 2008 the car maker fraudulently installed engine control unit (ECU) software determined to be a "defeat device", in violation of the Clean Air Act, to circumvent environmental regulations of emissions by diesel engine 2009–2015 model year Volkswagen and Audi cars. The software detects when the cars were being subject to emissions testing, and then fully enables the ECU emission controls to successfully pass. However, during normal driving conditions, emission control software was shut off in order to attain greater fuel economy and additional power, resulting in as much as 40 times more pollution than allowed by law. What happened to the stock price?

Within a matter of days, presumably as more information about the degree and extent of Volkswagen's manipulation became apparent, the share price fell from around $17 to around $11, a 33% loss in value.

Why did the market react in this way? Scandals of this kind *damage brand value.* By damaging *brand value*, VW will find it *harder to sell cars.* If they sell fewer cars, they make less earnings, which means less money to invest (*growth*), and less money to pay out in dividends (*income*).

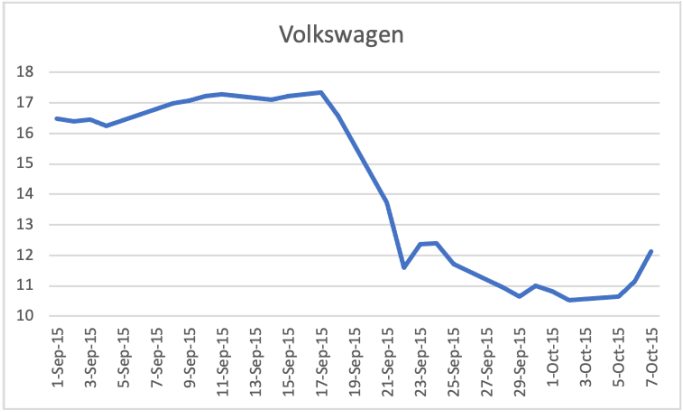

Figure 3: VW stock price around emissions scandal

4.2.2 Hang on...

One thing we can notice in both cases, but particularly Meta, is that there appears to be a clear pre *and* post trend in these stock price data. Is that consistent with semi-strong efficiency? The answer is **no**. To illustrate, consider the following figure which shows potential reactions to a *positive* announcement:

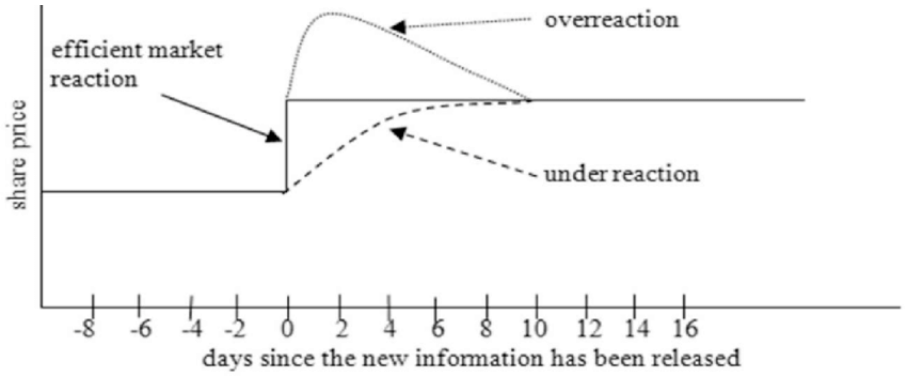

Upon hearing of a new announcement, the stock market could either react appropriately, *or under/over react*. In these latter two cases, the figure above notes what we'd expect to see in terms of stock price movements: in the case of over reaction, the stock price jumps above the 'correct' market reaction, and in the case of under reaction the stock price jump is *insufficient*. Over time, as the market learns of its over/under reaction, the price stabilises to the 'correct' value.

For Meta, what we appear to have is *under-reaction*. The stock price clearly drifts after the announcement, indicating that the *speed* at which this new information is converted into market activity is *slower than in the efficient case*.

This phenomenon is known as *post earnings announcement drift*. It is a remarkably stable phenomenon within the stock market: unlike CAPM, knowledge of this phenomenon *does not appear to have eliminated it!* Why might this be? There are many potential explanations, but I tend to think this is due to the marketplace being made up of a very heterogeneous spectrum of investors, some of whom react quickly, whereas others react slowly.

This slow reaction means that the stock price *does not* contain all publicly available information, meaning that *you can* make money by trading on public information. You just have to be quick!

4.3 Strong efficiency

Now for the MacDaddy: strong efficiency. Under strong market efficiency, prices contain *all information*. That means, all public *and* all private information. **All of it**. To clarify what this means, we need to provide more structure on what constitutes 'private' information. Private information is composed of two distinct types:

- **Inside information:** Info known only to company management but not yet made public.
  - Knowledge of takeover bid
  - Knowledge that earnings will be lower than expected
- **Private assessment based on public information**.
  - An analyst's report based on public accounting statements.

4.3.1 Evidence for strong efficiency

Most evidence around this claim relates to the ability to *outperform* the market. If market's have strong efficiency, *then it is not possible to consistently outperform the market*. Instead, we are back in our CAPM world. So, do people consistently outperform the market?

The answer appears to be *no*. Several papers have documented that mutual funds *underperform* relative to the market. Similarly, the evidence for hedge fund performance is far from clear cut. Is this sufficient to justify strong efficiency?

Also no! There are many stock market phenomena that are very difficult to reconcile with the view that markets are efficient. One economist who has spent much of their life identifying and analysing these phenomena is Nobel Laureate, Robert Shiller. Amongst his most famous work is the observation that the *volatility* of stock prices is *too large* to be reconciled with a world of rational expectations, and hence a perfectly efficient market. Take a look at the figure below:

Here, Shiller has mapped the *actual* price of the S&P 500 (solid black line), alongside the *ex-post* 'correct' price (dashed line) using the expression for the price noted above:

$$P_t = \mathbb{E}_t \left[\sum_{j=1}^{\infty} \left(\frac{1}{1 + r_{t+j}} \right)^j DIV_{t+j} \right]$$

Here we say *ex-post* because Shiller is using the *actual* returns to compute the 'optimal price. Note that the actual price is *considerably* more volatile. Why? How could this be, if people are using all available information, and have good models of how random processes work? The Nobel committee agreed this was a good question.

4.4 Summary

Of the three forms of market efficiency, we have good evidence for weak form, reasonable evidence for semi-strong, and some evidence for strong form. I think the takeaway from this section should be something like the following:

> Markets are not efficient, *but they are very close to being.*

In this respect, the same intellectual humility ought to dominate our thoughts when placing investment opportunities into 'risk-return' quadrants. Maybe you have found something that no one else has... but probably not!

5 Does the market encourage of discourage bad behavior?

One interesting point that could be made of the examples above, is that VW's decision to cheat on their emissions was *ultimately very costly*. The drop in the share price increased VW's cost of raising equity capital *considerably*, making a bad situation *even worse*. However, we've also seen examples of villainous managers like Jeff Skilling behaving very irresponsibly with respect to the health of his business, whilst enjoying sizeable rewards from the market. A natural question then emerges: do market places *discipline* for bad behavior, or encourage it?

In this section we assess the role that market forces play in guiding managerial behavior, especially as it relates to *environment*, *social*, and *governance* concerns. In the 'lingo', these concerns are labelled as **ESG** concerns. ESG is a very hot topic in academia at present, much as it is in the real world. In the interests of time, I'm going to focus exclusively on the environment, but the findings are roughly the same across all three.

What I hope to show you is that, slowly but surely, a consensus is emerging, and the details of that consensus may surprise you! Throughout the focus will be that *long-run risks matter, and they certainly matter to investors.*

5.1 Markets and the environment

A commonly held belief is that companies, in particular oil and gas companies, are at best indifferent, and at worst actively and enthusiastically committed, to carbon emissions. The 'capitalist system', with its ruthless, maniacal obsession with profits, leads to an aggressive and indiscriminate evisceration of our planet's natural resources, all to service the fat cats who stuff $100 bills into their cigar lighters while laughing hysterically.

Is this narrative accurate? It will depend on your perspective, but I personally would say a hard no. Investors, and businesses, *typically care about the medium and long-run risks that their companies face.* The risks related to climate change *are* relevant for businesses! Are these risks perfectly incorporated into investors' preferences and resultant stock prices: no. But do investors care about carbon emissions?? Absolutely!

5.1.1 Investors in their own words

One easy way to assess whether investors care about the environment, is to ask them. Now of course, talk is cheap, but it's a good place to start. A recent paper by Krueger, Sautner, and Starks (2020) did exactly this. They asked a large number of institutional investors (i.e. those who work at banks, mutual funds, hedge funds, etc.) a series of questions designed to establish:

- Whether they believed in climate change.
- How serious they thought the problem was.
- How their beliefs affect their investment strategies.
- How their beliefs affect their engagement with shareholders.

Here are a few quotes that help to summarise what they found:

- "Across all respondents, only 3% do not expect any temperature increase, 16% expect an increase by up to one degree, and 30% by up to 2∘C. Moreover, four in ten respondents expect a temperature rise that exceeds the Paris 2∘C target, with 12% expecting an increase of more than 3∘C."
- "More than half of the respondents that incorporate climate risks started to do so within the past 5 years.", "only a small percentage (7%) having chosen no approach to manage their climate risks during the 5 years preceding the survey"
- "divestment was the least used course of action when investors were dissatisfied with firm responses to their engagement (only 17% exited under such circumstances)."
- "Close to 30% of the investors submitted shareholder proposals on climate risk issues, and a similar fraction voted against management on proposals because of climate risk concerns."
- Only 16% reported no engagement over climate issues
- "...agreement is strongest for two motives: the protection of the investor's reputation (30% strongly agree), which can arise from both financial and non-financial motives, and moral/ethical reasons to consider climate risks (27.5%),

which would be a purely nonpecuniary motive. Institutions also tend to agree with the motive of incorporating climate risks due to a legal obligation/fiduciary duty (27%). Purely financial motives also score relatively high, especially the idea that incorporating climate risks is beneficial to returns (25% strongly agree) and reducing portfolio risk (24%) or tail risk (21%)"

So there appears to be *strong* evidence that investors are (i) aware of the problem, (ii) concerned by the problem, (iii) have taken action to mediate their concerns, and (iv) do so for a mixture of financial and ethical reasons... at least according to their own words!

5.2 Putting their money where their mouth is

Is this just investor nonsense? They are under no obligation to say anything other than the thing that makes them look good! So whilst it is *interesting* that this is how investors report their attitudes and behavior, *it is not enough*. To assess whether these claims bear out in action, we have to look at the data.

In a recent paper, Bolton and Kacperczyk (2021) show that as the emissions goes up, *institutional ownership goes down*. This is true in the US, *and globally*. They are not the only ones to show this result. So do the following four papers: Fernando et al. (2017), Starks et al. (2017), Dyck et al. (2019), Nofsinger et al. (2019).

If investors shun carbon stocks, what should we expect to be true of returns on those stocks? Recall the demand *and* supply nature of the stock market. If you're a carbon emitting firm, and you want to raise capital through equity, *then you need to offer some return*. If people shun your stock because of moral or long-term financial reasons, *raising equity will be very expensive*. One way around this problem is to *issue lots of dividends*. This would mean your stock offers *high* returns, even though the demand is *low*. So, do we see this?

Bolton and Kacperczyk (2021) give us the answer: yes. Carbon emitting firms consistently deliver greater returns relative to non, or low carbon emitting firms. This is true not just in the US, *but across the globe*. Investors are demanding *additional* compensation in order to bear the *long-run risk* of carbon emissions. So what does that mean for carbon-emitting firms? The lack of demand for their stocks from institutional investors lowers the stock price of carbon emitting firms, *and* the market exerts additional pressure to increase dividends to attract investors, placing *additional* pressure on firms to innovate away from carbon.

5.2.1 Engagement not divestment

Many commentators call for 'divestment' from fossil fuels, but investors tend to prefer engagement. Is this because they actually 'don't care'?

An excellent paper that explores the effectiveness of investor engagement is Dyck et al. (2019). They look at the BP oil spill and argue that this event *increased the salience of environmental concerns for the health of businesses*. What they want to explore is whether the composition of *investors* affected the trajectory of companies' environmental performance *after* the BP oil spill.

Here's what they do. They first establish for a large set of firms the percentage of their shares owned by institutional investors. They then look at the environmental performance of these firms *pre* and *post* the BP oil spill. Their hypothesis is that companies with more institutional investors will improve their environmental scores by *more* than those with lower levels of institutional ownership. This would be true *if institutional investors actually do put pressure on managers to improve environmental performance*. What do they find? Well, exactly that! Firms with higher institutional ownership improved their environmental performance by more than those with low institutional ownership.

Another very famous paper that explores investor engagement is Flammer (2015). In this paper, Flammer looks at the impact on companies of investor-led proposals that are voted on at board meetings that relate specifically to ESG. Flammer looks at firms that just passed these measures (slim margin of victory) versus those firms that just failed (slim margin of defeat), and is then able to claim that the only difference in these two sets of firms is the *fact* that the proposal passed.

What is the impact of proposals that 'just pass'? They are two-fold: firstly, and fortunately, these proposals do indeed appear to improve ESG performance. Secondly, and more importantly, *these firms do better in the stock market*. They post higher returns, have inflows of institutional investment, and perform better than rivals who *failed* to pass these measures.

5.2.2 Has this worked?

It is hard to say, but the numbers associated are sizeable. As of 2022, 5% of oil and gas capital expenditure goes on renewable energy. That may not seem like a lot, but that is a *five-fold increase* since 2019. It also amounts to roughly $7.2bn of investment in the US alone!

As the world has come to understand the likely trajectory of the global environment, and these issues have gained more prominence, investors, just like the rest of us, have responded by *divesting* and *engaging* with companies to try and *reduce* harmful environmental behavior.

Have they done this perfectly? No! Have they/will they solve the problem? Probably not! But the story is more nuanced than the newspapers, or your friends, would have you believe.

Even more crucially, note that if markets worked perfectly, the long-run risks of climate change *would by definition* be incorporated into the stock price, thus punishing carbon-reckless firms by increasing their cost of raising capital! The imperfectly efficient world of markets means that this market feedback is only a weakened version of the impetus that would otherwise *entirely* internalise the risk associated with climate change.