

DISCUSSION PAPER SERIES

DP17908

APPLYING ECONOMICS – NOT GUT FEEL – TO ESG

Alex Edmans

**ASSET PRICING AND BANKING AND
CORPORATE FINANCE**

CEPR

APPLYING ECONOMICS – NOT GUT FEEL – TO ESG

Alex Edmans

Discussion Paper DP17908
Published 16 February 2023
Submitted 11 February 2023

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Asset Pricing
- Banking and Corporate Finance

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Alex Edmans

APPLYING ECONOMICS – NOT GUT FEEL – TO ESG

Abstract

Interest in ESG is at an all-time high. However, academic research on ESG is still relatively nascent, which often leads us to apply gut feel on the grounds that ESG is so urgent that we cannot wait for peer-reviewed research. This paper highlights how the insights of mainstream economics can be applied to ESG, once we realize that ESG is no different to other investments that create long-term financial and social value. A large literature on corporate finance studies how to value investments; asset pricing explores how the stock market prices risks; welfare economics investigates externalities; private benefits analyze manager and investor preferences beyond shareholder value; optimal contracting considers how to achieve multiple objectives; and agency theory examines how to ensure that managers pursue shareholder preferences, including non-financial preferences. I identify how conventional thinking on ten key ESG issues is overturned when applying the insights of mainstream economics.

JEL Classification: N/A

Keywords: N/A

Alex Edmans - aedmans@london.edu
London Business School, CEPR, and ECGI and CEPR

Applying Economics – Not Gut Feel – To ESG

Alex Edmans

London Business School, CEPR, and ECGI

Current draft: February 11, 2023

Abstract

Interest in ESG is at an all-time high. However, academic research on ESG is still relatively nascent, which often leads us to apply gut feel on the grounds that ESG is so urgent that we cannot wait for peer-reviewed research. This paper highlights how the insights of mainstream economics can be applied to ESG, once we realize that ESG is no different to other investments that create long-term financial and social value. A large literature on corporate finance studies how to value investments; asset pricing explores how the stock market prices risks; welfare economics investigates externalities; private benefits analyze manager and investor preferences beyond shareholder value; optimal contracting considers how to achieve multiple objectives; and agency theory examines how to ensure that managers pursue shareholder preferences, including non-financial preferences. I identify how conventional thinking on ten key ESG issues is overturned when applying the insights of mainstream economics.

Keywords: ESG, SRI, CSR, sustainable investing, responsible business

Email: aedmans@london.edu, London Business School, Regent's Park, London NW1 4SA. I thank Tom Gosling and Ken Pucker for helpful comments.

1. Introduction

Executives, investors, policymakers, the media, and business schools are taking environmental, social, and governance (“ESG”) issues more seriously than ever before. Given ESG’s potential to create long-term value for both shareholders and society, this attention is both much-needed and welcome.

However, the enthusiasm for ESG – and, in particular, the pressure to do something or say something about ESG to demonstrate your commitment – can often lead to actions or statements that shoot from the hip and apply gut feel rather than being based on careful analysis. The issues facing society are so urgent that it seems we must take action immediately rather than waiting several years for new academic research to be conducted, peer-reviewed and published. But the same urgency means that the dangers of taking wrong actions are also severe. Treating a serious illness is urgent, which is why doctors take the Hippocratic Oath to “first do no harm”.

Moreover, we don’t have to wait several years for new academic research. As explained in Edmans (2023), ESG is both “extremely important and nothing special” – it’s economically no different to other intangible assets that create long-term financial and social value. Decades of finance research have studied how to create long-term financial value; decades of economics research have studied how to enhance social welfare. For example, any model of investment can be applied to ESG, since ESG is an investment with short-term costs and long-term benefits. Any model of agency theory considers managerial private benefits; while a common application is to executive perks, such models can equally be applied to externalities that the manager cares about – a private benefit is any outcome that’s not captured in long-term firm value.¹

¹ There are also models in which shareholders enjoy private benefits, surveyed by Edmans and Holderness (2017) who stress that “private benefits need not be at the expense of other shareholders”. Thus, these models can apply to ESG that is value-enhancing as well as ESG that is not.

Instead, shooting from the hip has led to many claims, and current practices, that are contradicted by landmark papers. Gut feel has its value, but can only get you so far – gut feel would suggest not giving water to a baby with diarrhea because it will flow out the other end; research suggests otherwise. In some cases, the claims and practices that stem from shooting from the hip are inconsistent with not only academic research but also each other. They are an example of “ESG Doublethink”, named after the term used in George Orwell’s *1984* to describe holding two contradictory positions.

The purpose of this article is to apply the insights from finance and economics to some of the biggest challenges facing the world today. It is not to be an apologist for academia, to defend the status quo, or to argue that no new research is needed. In my book on ESG (Edmans, 2020), I stress how traditional net present value (“NPV”) analysis is incomplete and needs to be supplemented by new principles (multiplication, comparative advantage, and materiality) to properly evaluate ESG decisions; I recently joined a long-standing finance textbook (Brealey, Myers, Allen, and Edmans, 2022) and updated it significantly to incorporate ESG. As Managing Editor of the *Review of Finance*, I co-launched a Special Issue on Sustainable Finance, recognizing that many ESG topics are unique and cannot be answered by simply reapplying existing theories (Edmans and Kacperczyk, 2022).

Instead, the goal of the essay is to caution against throwing the baby out with the bathwater – even though some ESG issues need to be answered by new research, existing research can provide substantial insights on others. NPV analysis is incomplete, not useless; the most effective approach will combine both established tools and new ideas. Since many other articles tackle these new ideas, this paper will focus on how mainstream economics can help us better understand ESG.

For brevity, I’ll sometimes use “green” (“brown”) to refer to a company with good (bad) ESG performance, although the performance need not be limited to climate. Edmans (2023) argues that the term “ESG” should often be disaggregated – some discussions of ESG are actually only about ES,

since they consider the pursuit of stakeholder value regardless of its impact on shareholder value. I'll thus use ESG when discussing factors that aim to enhance shareholder value, ES for actions taken to improve stakeholder value, E for a company's environmental impact, and G for governance.

The list of existing research that can be applied to ESG is almost limitless. I will thus organize this essay by practical ESG topic, rather than academic sub-literature. I will discuss how our views on ten key ESG issues change when applying the insights of mainstream economics. These are as follow:

1. Shareholder Value is Short-Termist (*No, shareholder value is a long-term concept*).
2. Shareholder Primacy Leads to an Exclusive Focus on Shareholder Value (*No, shareholders have objectives other than shareholder value*).
3. Sustainability Risks Increase the Cost of Capital (*No, sustainability risks lower expected cash flows*).
4. Sustainable Stocks Earn Higher Returns (*No, sustainability may be priced in; tastes for sustainable stocks lead to lower returns*).
5. Climate Risk is Investment Risk (*No, climate risk is an unpriced externality*).
6. A Company's ESG Metrics Capture Its Impact on Society (*No, partial equilibrium differs from general equilibrium*).
7. More ESG Is Always Better (*No, ESG exhibits diminishing returns and trade-offs exist*).
8. More Investor Engagement Is Always Better (*No, investors may be uninformed or undermine managerial initiative*).
9. You Improve ESG Performance By Paying For ESG Performance (*No, paying for some ESG dimensions will cause firms to underweight others*).
10. Market Failures Justify Regulatory Intervention (*No, regulatory intervention is only justified when market failure exceeds regulatory failure*).

2. Applying Economics to ESG Issues

2.1 Shareholder Value is Short-Termist (*No, shareholder value is a long-term concept*)

Executives', investors', and business schools' focus on shareholder value is believed to be the cause of many of the ills of this world. Shareholder value, the argument goes, is focused entirely on the short-term and thus the enemy of long-term ESG issues; the European Commission's 2020 "Study on Directors' Duties and Sustainable Corporate Governance" highlights the perils of "short-term shareholder value". Such a claim has profound implications for practice. If shareholder value is indeed short-termist, we should change the objective of the firm. This in turn affects such fundamental issues as the duties of a director, who gets to vote in director elections, how executives should be evaluated and compensated, and whether we should protect firms from takeovers and shareholder activism. The entire literatures on the theory of the firm and corporate finance would need to be scrapped and rebuilt from scratch because they are founded on the wrong principles.

But Finance 101 teaches us that shareholder value is an inherently long-term concept. It is the present value of all future cash flows, from now until the end of time:

$$V_0 = E \left[\sum_{t=1}^{\infty} \frac{C_t}{(1 + r_t)^t} \right]$$

where V_0 represents shareholder value at time 0, E denotes expected values, C_t is the cash flow at time t , and r_t is the discount rate.² The key is the ∞ in the summation term \sum , which highlights that shareholder value takes into account all cash flows. "Maximizing shareholder value" would thus give a big green light to many ESG projects, such as clean energy, carbon capture, and workforce upskilling – even though near-term cash flows are negative, future cash flows are highly positive.

² For simplicity, I am assuming an all-equity firm, so cash flows to the company are cash flows to shareholders, and that these cash flows are paid to shareholders as dividends.

Indeed, NPV analysis – evaluating an investment by its effect on shareholder value – leads to far more long-term investments than the methodologies it replaced, such as payback period and accounting rate of return, which ignore distant cash flows.

There is nothing short-termist about shareholder value being the objective of the firm. Instead, the problem is too little focus on shareholder value, rather than too much; executives instead focus excessively on profits. Indeed, they freely admit to sacrificing shareholder value to meet short-term earnings targets (Graham, Harvey, and Rajgopal, 2005), and those with short horizons cut investment to do so (Edmans, Fang, and Lewellen, 2017).

Shareholder value can't be observed since it contains expectations of the future, so how do we know whether CEOs are maximizing it? If financial markets are efficient, then the time-0 stock price P_0 equals the market's assessment of shareholder value:

$$P_0 = E \left[\sum_{t=1}^{\infty} \frac{C_t}{(1 + r_t)^t} \middle| I_0 \right]$$

where I_0 is the information available at time 0. The market forms its expectations for future cash flows and the discount rate based on all available information. Note the word *all*. Critics argue that an efficient market only takes into account financial information and ignores ESG factors. This is not the case – it considers any information relevant for forecasting future cash flows and the discount rate, financial or non-financial. I_0 is simply information; it does not restrict it to be a particular type. Thus, a manager who only cares about today's stock price will still maximize long-term shareholder value.

The critical words are “in an efficient market”. In reality, markets may not be efficient, and many of my own papers demonstrate this – I'm far from a staunch defender of efficient markets theory. Inefficiencies may come from multiple sources. One frequently-accused culprit is shareholders having short horizons. Critics lament how holding periods have shortened over time – if a shareholder

(Andrea) only expects to hold her shares for one year, then she won't look beyond that year and will discourage a company from investing in clean energy because the fruits will arise after she's sold out.

But slide 2 in lecture 1 of my Masters in Financial Analysis course shows that these concerns are unfounded. If the investor intends to sell at the end of year 1, the value of the company to her is (dropping the expectations E and information I_0 for brevity):

$$\frac{C_1}{1 + r_1} + \frac{P_1}{1 + r_1}$$

At the end of the year, she receives the first year's cash flow as a dividend, and then immediately sells out and receives the year-end stock price. What will the year-end stock price be? Another investor, Bimal, buying the share at the end of year 1, knows that he will receive all future cash flows from year 2 onwards. He will pay:

$$P_1 = \sum_{t=2}^{\infty} \frac{C_t}{(1 + r_t)^{t-1}}$$

Substituting this into the prior equation yields the value to Andrea as:

$$\frac{C_1}{1 + r_1} + \frac{\sum_{t=2}^{\infty} \frac{C_t}{(1 + r_t)^{t-1}}}{1 + r_1} = \sum_{t=1}^{\infty} \frac{C_t}{(1 + r_t)^t}$$

Thus, Andrea cares about all future cash flows until the end of time. She explicitly receives the first-year cash flow as a dividend, and she implicitly receives all cash flows from year 2 onwards as they affect the stock price she receives at the end of year 1. Thus, her horizon is irrelevant.

Even though the above exposition is elementary, it is misunderstood by many practitioners, including regulators. For example, France's Loi Florange gives double the voting rights to shareholders who've held their stake for at least two years. The idea is that such shareholders are more long-term-oriented, and so we want them to have a greater say in the firm's actions. This argument is flawed for two reasons. First, how long an investor has held shares for in the past has no

clear relation to how long he'll hold shares in the future. Indeed, it may be inversely related – if the average investor has a three-year horizon, then a new shareholder cares about the next three years, but a “loyal” shareholder who's already been there for two won't look past the next year. Second, irrespective of how many years you intend to hold shares for in the future, you care about all future cash flows.

Worryingly, this elementary idea is also misunderstood by many academic papers, including those published in the most elite peer-reviewed journals. They use a shareholder's average holding period as a measure of short-termism, under the idea that shareholders with short holding periods will pressure companies to focus on cash flows that arise only during their tenure. However, all shareholders care about all future cash flows, regardless of their horizon.

But my above argument breaks down if there's a second source of inefficiency that's nothing to do with the holding period. What if P_t , the price that Bimal pays Andrea, is inefficient, and not the present value of all future cash flows from year 2 onwards? For example, if the market at $t=1$ underweights future cash flows, then Andrea won't receive the value of these cash flows when she sells to Bimal at $t=1$. As a result, at $t=0$, she will pressure the firm not to invest in projects that yield distant cash flows.

This is indeed a realistic possibility: in Edmans (2011, 2012) I show that the market underprices employee satisfaction, and cite several papers on the underpricing of other intangible assets. However, the market is just as likely to overvalue distant cash flows. The value anomaly, initially documented by Lakonishok, Shleifer, and Vishny (1994), indicates that the market overvalues growth stocks. Polk and Sapienza (2009) show that, when the market overvalues investment, this leads managers to overinvest. Bolton, Scheinkman, and Xiong (2006) demonstrate that, if investors only care about the short-term stock price and that price is inefficient due to speculation, they'll induce the CEO to undertake “castle-in-the-air” ventures. At the time they wrote their paper, those ventures involved

tech investments in the internet bubble, some of ended up paying off; nowadays, some of the most ambitious solutions to the climate crisis could be considered castles in the air and may be encouraged by investor short-termism. Indeed, the electric vehicle industry enjoyed a bubble in the early 2020s which facilitated capital raising. Thus, while stock markets are indeed inefficient, they're just as likely to overvalue future cash flows as undervalue them, and spur managers to overinvest in ESG as underinvest. Indeed, if stock markets systematically undervalued ESG performance, then ESG funds, which buy ESG outperformers, would systematically beat the market, but this is not the case (Renneboog, Ter Horst, and Zhang, 2008).

Finally, even if stock prices systematically underpriced (rather than overpriced) ESG factors, this would not be a reason to stop paying managers according to the stock price. Instead, it would be to lengthen the horizon of CEO equity, so that CEOs are concerned with the long-term rather than short-term stock price (Edmans, Gabaix, Sadzik, and Sannikov, 2012).

2.2 Shareholder Primacy Leads to an Exclusive Focus on Shareholder Value (*No, shareholders have objectives other than shareholder value*)

Almost as maligned as shareholder value is the so-called “doctrine” of shareholder primacy, which argues that shareholders should be considered superior to all other stakeholders. While I’m unaware of any papers that have proposed it (or any other economic concept) as a doctrine, this idea manifests in shareholders, and shareholders only, voting in director elections and on other corporate matters. Even if shareholder value were truly long-term, shareholder primacy still seems problematic. While many effects that companies have on wider society are eventually internalized and affect shareholder value, others don’t because they are *externalities* – societal effects that don’t feed through to profits, even in the long term. If true, directors’ fiduciary duties should be to more than just shareholders, and other stakeholders (such as workers) should elect directors. Indeed, the pursuit of

externalities is the distinguishing feature of ES; if ES factors always improved long-term profits, then ES would simply be shareholder value maximization (Edmans, 2023).

However, managers do not ignore other stakeholders, for two reasons. The first is that these stakeholders are protected. Some are protected by contract, such as workers, suppliers, and customers. If the firm underpays its workers or suppliers, or fails to provide its customers with the promised goods or services, it can be sued. Others are protected by laws, such as environmental regulations. Shareholders have no contractual rights. They invest their money into the firm and are guaranteed nothing in return; companies have no obligation to pay any dividend. Even if a dividend has been paid in the past, the company could cut it to zero, and shareholders can do nothing about it.³ It is due to the lack of contractual rights that shareholders have primacy.

But contracts are imperfect. Employment contracts can stipulate salaries, but not working conditions and corporate culture. Launching lawsuits is expensive; stakeholders' contracts may be breached but it's not worth it for them to sue. Regulation is imperfect; even though mainstream economics shows that a global carbon tax would be the most effective solution to climate change, one does not yet exist. Note that even Friedman (1970), in his claim that "the social responsibility of business is to increase its profits", recognized that companies should go beyond contracts and regulation – he stressed that they must conform "to the basic rules of society, both those embodied in law and those embodied in ethical custom." Indeed, there are social norms for issues such as prompt payment of suppliers and providing more than the legal minimum for severance pay, even though following these norms may reduce profits. One might argue that it's not enough to just rely on social norms; stakeholders deserve additional protection. But then the following questions crop up (Denis, 2016): who decides which stakeholders to protect? How do know how much shareholder value to

³ While shareholders could sell their shares, this won't reverse the dividend decision. Indeed, the stock price will have likely fallen due to the dividend cut, so shareholders also lose from a lower stock price.

sacrifice to protect them? How do we trade-off the interests of different stakeholders with each other? We could protect the environment by closing down coal-fired power stations, but this would hurt workers.

That's where the second consideration comes in. Shareholders are typically concerned with more than about shareholder value. A fundamental feature of nearly every economic model is that agents care about more than just financial factors; indeed, it's these trade-offs that make the problem non-trivial and require the model in the first place. Workers care about both wages and leisure. Consumers care about both prices and product quality. Society cares about both wealth and non-wealth utility; a key field in economics is welfare economics, which delivers axioms such as the Fundamental Theorems of Welfare Economics. Indeed, shareholders often care about wealth and externalities – they seek to maximize shareholder *welfare*, not shareholder value. For example, a citizen investing for retirement cares not only about his wealth in the future, but also the state of the planet as this will affect his welfare. This example highlights shareholders don't need to be altruistic to be concerned with more than shareholder value; an entirely self-interested shareholder would care about anything that affects his future standard of living.

One of the most influential finance models of all time – and one of the most criticized and misunderstood by ESG advocates (even though it's a model of G), is Jensen and Meckling (1976). Their Figure 1 studies the problem of an owner-manager, and shows that he will not maximize shareholder value, but a mix of shareholder value and private benefits.⁴ Jensen and Meckling (1976) list examples of such private benefits: “charitable contributions, personal relations (‘love’, ‘respect’, etc.) with employees” – two ES objectives. When the owner is no longer the manager, and we have the separation of ownership (by millions of investors) and control (by directors and executives), good

⁴ Jensen and Meckling (1976) use the term “non-pecuniary benefits”. We use “private benefits” as it is the more common term in the literature.

governance ensures that the firm pursues shareholders' interests – including any ES objectives. Thus, shareholder primacy does not mean that companies will focus exclusively on shareholder value and ignore stakeholders. ES advocates often criticize “agency theory”, a key strand of economic research in which Jensen and Meckling (1976) is a seminal paper, as focusing on shareholder value maximization (or, even more inaccurately, short-term profit maximization). But agency theory studies how to address the principal-agent problems that arise when agents (executives) have different objectives from the principal (shareholders). There is nothing in agency theory that restricts the principal's objective to just shareholder value.

I wrote that shareholders “typically” and “often” care about more than just shareholder value. However, it's possible that some shareholders don't realize that their future welfare will be affected by externalities; alternatively, the companies they own have such a small effect on aggregate externalities that investors want them to maximize shareholder value. Such shareholders exist, but in practice they have little power even in countries that give shareholders primacy, because such countries also give substantial latitude to directors. Investors' votes are limited to director elections, a narrow set of corporate actions (e.g. executive pay and mergers), plus advisory votes on shareholder proposals. In the US, the Business Judgment Rule protects directors from liability, even if their decisions end up eroding shareholder value *ex post*, as long as they exercised sound business judgment *ex ante*. This gives companies substantial freedom – they can give millions of dollars to charity, pay furloughed workers, and donate free products both inside and outside of the pandemic, even though these actions do not clearly increase shareholder value. Since it's impossible to prove that these actions clearly destroy shareholder value (charitable donations may boost a company's image and attract employees and customers), courts defer to directors' business judgement. In the UK, the Companies Act 2006 states that directors' duties are to “promote the success of the company”, not to increase shareholder value. While the next words are “for the benefit of its members” (referring to

shareholders), legal scholars and practitioners point out that this “benefit” includes non-financial benefits – such as a cooler planet as discussed earlier.

What shareholder primacy does restrict is executives’ ability to pursue goals that are clearly destructive to shareholders. A CEO cannot engage in gross negligence, transactions in which she has a conflict of interest, or ES investments that have very little probability of increasing shareholder welfare (such as donating to a charity that only she cares about). Shareholder primacy does not limit a CEO’s ability to make ES investments that increase shareholder value or produce externalities that shareholders care about. Indeed, shareholders are often pressuring companies to make more ES investments than they would do otherwise.

2.3 Sustainability Risks Increase the Cost of Capital (*No, sustainability risks lower expected cash flows*)

The list of ESG risks that a company faces is almost endless – stranded assets, executive malfeasance, employee unrest, customer boycotts, regulatory fines, media shaming, and so on. Since these are risks, it seems a no-brainer that investors should incorporate them into valuations by increasing the discount rate – after all, Finance 101 tells you that the discount rate depends on risk. Indeed, the PRI’s (2016) comprehensive survey of valuation approaches used in practice concludes: “Some investors adjust the beta or discount rate used in company valuation models to reflect ESG factors: corporate governance, operational management, general quality of management, its strategic decision making etc.” (It is not clear why operational management, general quality of management, and strategic decision making are “ESG” factors; these are simply attributes of a great company (Edmans, 2023). ESG investing loses credibility if we try to label anything that might be good about a company as “ESG”). It contains numerous case studies of investors that do so. One investor explains how they lower the discount rate for companies with above-average gender diversity, arguing that

“Diversity also helps to reduce company-specific risk in the long term, leading to a lower cost of capital.” Similarly, question 12 of the specimen exam for the CFA UK Level 4 Certificate in ESG Investing is “What impact will a high ESG rating have on a company’s cost of capital?”. The answer key gives the correct response as “A: A lower cost of capital”.

But the main effect of a risk is to change the expected cash flows in the numerator of a valuation. If a cash flow should be \$1 million, and there’s a 10% chance of a major ESG catastrophe that reduces it to \$200,000, a 15% risk of a moderate disaster that lowers it to \$300,000, and a 20% probability of a mild calamity that decreases it to \$400,000, the expected cash flow becomes \$695,000.

These risks need not increase the discount rate. Finance 101 tells us that the discount rate is affected *only* by market risk, and *not* by company-specific risk, in contrast to the statement by the investor in the PRI survey. The discount rate only increases if the risk of the disaster is correlated with market conditions – i.e. a disaster is more likely in bad times. In contrast, many ESG scandals, such as Volkswagen cheating emissions tests, Wells Fargo opening fake bank accounts, or Rio Tinto blowing up Juukan Gorge, are company-specific, not market-wide – they’re no more likely in a down market than an up market. In fact, one could argue that some ESG scandals are more likely in an up market – when times are good, companies get sloppy and don’t impose as tight controls. Now it’s true that some ESG risks affect the broader market – for example, a carbon tax will impact many firms. However, what matters is not only whether a risk affects the broader market, but whether it’s more likely to manifest in up or down markets. As Giglio, Kelly, and Stroebel (2021) show, the government could be more likely to implement a carbon tax when the market is booming, since economic activity leads to high emissions and thus the need for a tax; then, brown assets actually bear lower market risk, since they do worse in a booming economy. Similarly, the model of Baker, Hollifield, and Osambela (2022) demonstrates that they are a hedge against the reduction in welfare

that arises if the government fails to take action on climate change. Thus, while ESG risks definitely reduce expected cash flows, the impact on the discount rate is far from clear.

However, changing expected cash flows seems a hassle – you need to estimate the probability of different scenarios and what will happen to the firm in each scenario. Since such estimates will only be approximate, you’ll then need to do sensitivity analyses around your base case. Because this is cumbersome, some investors instead increase the discount rate because it’s easier – you just need to change one number – but the discount rate shouldn’t change if the risk is company-specific. Now it’s true that you can “fudge” the valuation by increasing the discount rate to reflect the effect of lower cash flows. But there’s no guidance as to how much this increase should be – there’s not even a logical midpoint about which to do a sensitivity analysis. It would be like incorporating the effect of higher energy prices by hiking the discount rate rather than raising the costs in the cash flow statement.

The textbook of Brealey et al. (2022) states the correct approach clearly: “Remember that a project’s cost of capital depends only on market risk. Diversifiable events can affect project cash flows but they do not increase the cost of capital... Don’t give in to the temptation to add fudge factors to the discount rate to offset things that could go wrong with the proposed investment... Adjust cash-flow forecasts instead. Fudge factors in discount rates are dangerous because they displace clear thinking about future cash flows.”

Note that there *are* reasons why brown companies might suffer a higher cost of capital, but they are more nuanced than “there’s a risk that cash flows fall”. The first is if the ESG risks are more likely to manifest in bad times – for example, if firms are more likely to engage in fraudulent behavior in recessions because they’re desperate, or if a climate disaster both reduces shareholder welfare and spurs government action. The second has nothing to do with risk. If investors have a taste for green stocks and dislike holding brown stocks, they’ll demand a higher cost of capital to own the latter (see

Pastor, Stambaugh, and Taylor (2021) for a model of both types of risk). However, unless either condition is satisfied, sustainability only has an effect on cash flows, and not the cost of capital.

2.4 Sustainable Stocks Earn Higher Returns (*No, sustainability may be priced in; tastes for sustainable stocks lead to lower returns*)

Many investors claim that ESG enhances shareholder value, and thus green companies will enjoy superior returns – hence justifying the practice of ESG investing. However, Finance 101 teaches us that a shareholder’s return is relative to the price he pays: the 1-year return on a stock is:

$$\frac{C_1 + P_1}{P_0} - 1.$$

If ESG enhances shareholder value, then next year’s dividend C_1 and stock price P_1 will be high. However, if the high ESG was already known at $t=0$, the market would already price it in, leading to a high P_0 . Thus, the stock won’t generate superior returns.

This simple idea is not fully understood even by leading investors. For example, Larry Fink’s 2022 letter to shareholders states: “In today’s globally interconnected world, a company must create value for and be valued by its full range of stakeholders in order to deliver long-term value for its shareholders.” This sounds undisputable, but is far from unambiguous. Such a firm delivers *shareholder value*, but it may not deliver *value to its shareholders*. A company only delivers value, i.e. returns, to its shareholders if it creates shareholder value in ways that the market did not expect – if so, this value creation would not have been initially not priced in and so investors could have bought the stock on the cheap. If the firm were “valued by its full range of stakeholders”, including its shareholders, then the initial price P_0 would have already reflected its superior ESG. Thus, even if this ESG did subsequently “deliver long-term value” in the form of a high C_1 and P_1 , shareholders would not enjoy abnormal returns. Similarly, the UK’s leading broker states that “Study after study

has shown that businesses with ESG (environmental, social and governance) characteristics have outperformed their peers. So, what’s good for the environment and society could be good for your wealth too.” The evidence is far more mixed than claimed, and the broker cited no evidence so it is not clear on what dimensions ESG companies “outperformed” according to these studies. But if the outperformance is on any accounting measures, such as sales, operating profit, or net income, then it will already be priced in and thus not be “good for your wealth”.

ESG factors may indeed lead to superior shareholder returns if they are unanticipated. For example, Edmans (2011, 2012) shows that employee satisfaction was not fully priced in over 1984-2011, leading to abnormal returns of 2.3-3.8%/year, and Boustanifar and Kang (2022) document that the mispricing continued through 2020. In addition, ES factors may yield higher returns if there is a shift in investor tastes towards green stocks (Pastor, Stambaugh, and Taylor, 2022). However, it is equally possible that ESG factors may be overpriced, if investors overestimate the value of ESG, as was the case for electric vehicle companies in 2021. Similarly, tastes may shift away from ES stocks given the current ES backlash, leading to inferior returns.

In equilibrium, there is no mispricing because investors fully recognize the value of ESG, and tastes are stable. If so, green stocks may lead to lower returns for the two reasons given at the end of the prior section: they hedge against government action on climate change and action is prompted by a welfare-reducing climate disaster, or investors have a taste for green stocks (Pastor, Stambaugh, and Taylor, 2021). In both cases, investors are fully aware of green stocks’ lower returns but are willing to hold them anyway for their hedging properties or a warm glow feeling. Indeed, in equilibrium, *investor returns equal the cost of capital*. If sustainability does indeed lower the cost of capital, it must also lead to lower investor returns. Some ESG advocates claim that ESG is good for companies because it lowers the cost of capital, and is also good for investors because it increases shareholder returns. This is an example of ESG Doublethink that can never be simultaneously true in equilibrium.

2.5 Climate Risk is Investment Risk (*No, climate risk is an unpriced externality*)

This is another seemingly unambiguous phrase in a Larry Fink CEO letter, but is much more complex than it appears. Indeed, it is another example of ESG Doublethink. As discussed in Section 2.2, one of the main – and entirely valid – arguments for ES investing is that there are unpriced externalities that investors care about, and so investors should reduce externalities even at the expense of shareholder returns. But if there are unpriced externalities, then climate risk is not investment risk – investors don’t bear the consequences of climate change because polluting companies don’t have to pay for the damage they cause (for example, through a carbon tax). Indeed, this was the key point in Stuart Kirk’s famous 2022 speech that led to him being pressured to resign as HSBC’s head of responsible investing – climate change is a serious risk to society, but not a risk to investors in the absence of government regulation. Indeed, Bolton and Kacperczyk (2021) find that companies with higher carbon emissions earn higher stock returns, consistent with such companies being able to get away with the negative externalities they cause.⁵

But what if the government finally takes action and imposes a carbon tax, so that investors do bear climate risk? Indeed, an alternative interpretation of the Bolton and Kacperczyk (2021) results are that investors are worried that a carbon tax might be passed, and thus will only hold emitting stocks if they earn a higher return to offset this risk (assuming that this risk manifests in bad times). If so, the phrase “climate risk is investment risk” is, strictly speaking, correct. However, investors often use this phrase to justify not investing in emitting stocks. For example, the first paragraph of the section of Larry Fink’s letter entitled “Climate Risk is Investment Risk” states that “Our investment conviction is that sustainability- and climate-integrated portfolios can provide better risk-

⁵ This will lead to higher stock returns if the market was not fully aware that companies don’t bear the consequences of such externalities.

adjusted returns to investors.” But if investors earn a return from bearing risk, there is no loss in risk-adjusted returns from owning brown firms. Finance 101 tells you that, in an efficient market, risk is compensated for by a higher return, and so a risky stock is just as good as investment as a safe one. As Edmans (2023) points out, “Holding stakes in young firms, tech companies, and emerging markets bears investment risk, but that risk is compensated for by a return. If an asset manager wanted to avoid investment risk, it would ironically eschew clean energy and carbon capture.”

A separate problem is that “climate risk” is an ambiguous term. There are two types of climate risk. The first is physical risk, the risk that a company experiences from a warming planet. A real estate company with waterfront properties bears significant physical risk, as does an agricultural firm confronting the threat of drought. The second is transition risk, the risk that a company faces from a move to a low-carbon economy, such as a carbon tax or customers boycotting emitting companies.

Carbon emissions is the most common measure of “climate risk”. However, it is only a valid measure of transition risk; it does not capture physical risk. The latter is how much your company suffers when the climate changes, not how much your company affects the climate – just like how, in Finance 101, market risk measures how much your company suffers when the market declines, not how much your company affects the market. Indeed, it’s reasonable to argue that physical risk is even more important to investors than transition risk, given limited government action on climate change and the little attention paid to climate change mitigation. In this case, carbon emissions are a particularly poor measure of “climate risk.”

2.6 A Company’s ESG Metrics Capture Its Impact on Society (*No, partial equilibrium differs from general equilibrium*)

Investors, stakeholders, and the media are demanding that companies report an ever-increasing set of ESG metrics so that they can better value them, and (for ES metrics) so that they can hold them

accountable for their societal impact. Implicit in this trend is the assumption that a company's ES metrics capture its externalities. However, a company can improve its ES metrics at the expense of other firms, leading to a zero or negative effect on aggregate externalities. For example, it can increase gender diversity by hiring female employees away from its competitors; doing so may not increase gender representation across the industry (and may decrease it since the females lose firm-specific human capital and incur switching costs). Alternatively, it can reduce its environmental footprint by selling polluting plants to its peers, who run them less efficiently, thus increasing aggregate pollution.

Such concerns are not merely hypothetical possibilities. Duchin, Gao, and Xu (2022) study how industrial companies respond to the greater scrutiny that follows environmental risk incidents. They improve their E metrics by divesting polluting plants; however, under new ownership, per-employee pollution increases. However, ESG rating agencies don't realize this, perhaps taking emissions numbers at face value, and reward the divesting companies with superior ESG ratings.⁶

This concern is linked to the economic distinction between partial equilibrium and general equilibrium analysis. The former considers the effect of a policy on a single market; the latter includes the knock-on effects on other markets – for example, a tobacco tax might encourage consumers to switch to alcohol. I'm abusing terminology slightly since a partial *equilibrium* analysis refers to an entire market, rather than just one company. However, the problem is similar – partial thinking, such as studying one company in isolation, ignores knock-on effects on society in general.

Moreover, this point applies beyond companies to investors. Many asset managers are concerned with decarbonizing their portfolios or aligning them with net zero, but decarbonizing your portfolio

⁶ While ES investors should be concerned about aggregate externalities, rather than only those produced by the companies they own, it's possible that some investors' tastes are such that they only care about the social impact of their holdings. For example, some shareholders exclude tobacco companies, and don't engage with the government to increase tobacco taxes or restrict tobacco advertising as they don't feel it's their responsibility if they don't hold tobacco. Even if we had a literal definition of "its" in "A company's ESG metrics capture its impact on society", i.e. we don't include the pollution produced by plants outside a company's control, this measure still does not fall. The sellers don't lose access to the divested plants, since they are typically bought by firms with supply chain relationships or joint ventures with the sellers.

does not decarbonize society. The only way that you can sell your shares in an energy company is if another investor buys; it's analogous to clearing up the trash from your garden by throwing it over the fence into your neighbor's.⁷ The carbon that a company emits depends on the assets that it owns, not its liabilities (how it finances those assets). It's true that reducing a company's liabilities will also reduce its assets, but selling shares affects neither the assets nor the liabilities side of the balance sheet – the total shares outstanding stays the same; the shares are merely owned by different people. It's also true that those who own your liabilities may influence which assets you hold; shareholders enjoy voting rights. However, investors who choose to buy your energy stocks may have little concern for climate change and not engage with the company to cut its emissions. Addressing the partial issue of making your numbers look good actually worsens the general problem.

The same issue applies to citizens concerned with their personal carbon footprint. One way to lower it is to replace your petrol car with an electric car. However, if you sell your petrol car, the new owner will drive it. In the UK, every four new cars purchased results in three additional cars on the road, so buying an electric car increases the number of cars in circulation by 0.75 (Gosling, 2019). Thus, while it's unambiguous that an electric car reduces aggregate emissions if your petrol car has come to the end of its life, it's far less obvious that you should replace the petrol car prematurely.

2.7 More ESG Is Always Better (*No, ESG exhibits diminishing returns and trade-offs exist*)

Even ignoring the general equilibrium considerations detailed in Section 2.6, companies and investors often treat ESG metrics as if higher numbers are always better.⁸ However, ESG is an

⁷ Even though another investor buys your shares, the act of selling lowers the stock price, which makes it harder for the company to raise capital in the future and harms the CEO's wealth and reputation. Thus, a CEO might be concerned with investor selling even though it does not directly deprive a company of capital. But even if so, Edmans, Levit, and Schneemeier (2022) show that the optimal divestment strategy may be tilting (leaning away from a brown industry but being willing to hold a best-in-class company) rather than blanket exclusion. Exclusion gives the firm no incentives to reform as it will be excluded anyway; tilting incentives the company to reform and become best-in-class.

⁸ More precisely, higher numbers are always better for "good" measures such as diversity, and lower numbers are always better for "bad" measures such as carbon emissions.

investment, and any investment features diminishing returns and/or increasing costs. From a financial standpoint, a company should stop investing when the benefits no longer exceed the cost, meaning that the optimal level is finite. Moreover, this optimum might vary between firms and over time. For example, Edmans, Pu, Zhang, and Li (2023) find that the payoff to being a Best Company to Work For is decreasing in the labor market rigidity of the country in which the firm is located. When labor market regulations already guarantee a minimum level of worker welfare, companies that have a high level of employee satisfaction compared to their peers may be overinvesting in it.

While increasing employee satisfaction is often directly costly (e.g. through paying higher wages or providing superior working conditions), diminishing returns may exist even for ESG metrics where there's no direct cost. The evidence on the benefits of demographic diversity is mixed to negative (Fried, 2021) but there is evidence on the payoffs to cognitive diversity. Even if cognitive diversity were measurable, and even if it were possible to increase cognitive diversity without any cost (e.g. there's no trade-off with other employee characteristics), Aggarwal et al. (2019) find a hump-shaped relationship between cognitive diversity and team learning. Too much diversity backfires because team members “speak different languages” and have a hard time understanding each other. These diminishing returns apply not only to ES but also to G – in Section 2.8 we will discuss how excessive governance erodes managerial initiative.

Moreover, even if we ignore financial performance and only consider externalities, higher ES numbers aren't always better due to trade-offs with other ES factors not captured in standard metrics. Lower carbon emissions could result from a company producing less output, thus generating less consumer surplus, providing fewer revenues to suppliers and fewer jobs to employees. Even when considering carbon intensities (emissions divided by sales), an energy company could reduce this number by rapidly closing down its most polluting plants without for finding new jobs for the employees working in those plants.

2.8 More Investor Engagement Is Always Better (*No, investors may be uninformed or undermine managerial initiative*)

When companies erode financial or social value, investors are often blamed for not preventing it. For example, the 2009 Walker Report partly attributed the UK financial crisis to inadequate stewardship by investors; shareholders are similarly chastised for failing to rein in excessive executive pay. On the flipside, society praises investors for engaging on ES issues, such as pushing companies to reduce their carbon footprint, pay higher wages, and increase workforce diversity. The underlying assumption is that more investor engagement is always better. This is an example of ESG Doublethink, since ES advocates often oppose shareholder primacy on the grounds that it will force companies to focus narrowly on shareholder value (Section 2.2).

Perhaps we should refine this assumption – more engagement on ES is always better, but more engagement on shareholder value is always worse? But that’s also incorrect, because Section 2.7 highlighted how more ES is sometimes suboptimal. It’s also contradicted by classic economic models which highlight that investor engagement may undermine managerial initiative.

Two seminal papers by Aghion and Tirole (1997) and Burkart, Gromb, and Panunzi (1997) show that it’s optimal for shareholders to give managers latitude to pursue different objectives from theirs. A rough intuition is as follows. Investors only care about shareholder value (we’ll relax this assumption shortly). The manager cares about both shareholder value and private benefits, such as ES. It takes effort for the manager to search for value-creating projects. Once she’s found a range of projects, investors might step in and force her to take the one that maximizes shareholder value, even if it doesn’t maximize ES. Since the CEO cares about ES, if she expects such interference, she’ll put in less effort to come up with the projects to begin with, and shareholders are worse off. Shareholders’ desire to split the pie in their favor reduces the CEO’s incentive to grow the pie in the first place.

It's thus in shareholders' interest to commit not to interfere. In Aghion and Tirole (1997), they do this by giving "real authority" to the CEO, i.e. delegating decisions to her even though they possess "formal authority". This is consistent with the wide discretion given to directors in practice, and the limited votes given to shareholders. In Burkart, Gromb, and Panunzi (1997), they take a small stake to begin with, reducing their incentives to interfere. Moreover, while both models feature investors caring only about shareholder value and the manager caring also about private benefits, the principle applies to any difference in goals. It may be that the manager's objective is purely shareholder value and investors' objective includes ES, or both objectives include ES but different dimensions. Regardless of the source of differences, the same result applies: more shareholder engagement *ex post* deters managerial initiative *ex ante*.

Even ignoring any effect on *ex ante* initiatives, more shareholder engagement is still not always better. First, shareholders may have objectives other than firm value, such as ES goals that are at the expense of profits. A long literature, starting with Shleifer and Vishny (1986), studies the incentives of blockholders (large shareholders) to take value-creating actions. While these actions are often interpreted as engagement that enhances firm value, the literature review of Edmans and Holderness (2017) stresses that they may also refer to "not taking private benefits". The key finding of this research is that the incentive to create value is increasing in the blockholder's stake. Thus, shareholders with small stakes are likely to pursue private benefits when they engage, i.e. their own ES goals. Empirically, Gantchev and Giannetti (2021) find that shareholder proposals by "gadfly" investors, who own small stakes in hundreds of firms which they use to advance proposals at each one, typically destroy long-term firm value. Gadflies are often single-issue investors, who define success by how many firms they can persuade to adopt their one issue, irrespective of whether it's material to their long-term performance or a significant externality for that particular company.

One counterargument is that this section is entitled “More Investor Engagement Is Always Better”, but better for who? The discussion has focused on shareholder value, but perhaps engagement improves societal welfare – any reduction in shareholder value is more than outweighed by the positive effect on externalities. If other shareholders also care about externalities (Section 2.2), then shareholder welfare improves overall. However, this may not be the case. There are several models of blockholder private benefits, starting with Zwiebel (1995), where different shareholders pursue different private benefits, which in our setting correspond to contrasting ES preferences. A shareholder may engage to persuade firms to pursue his favored ES goal, at the expense of both firm value and other ES objectives. For example, Agarwal (2012) finds that labor union pension funds vote for labor-friendly directors, even if they may oppose the closure of polluting plants.

In addition to differing ultimate objectives (which, if any, ES goals to pursue), a second source of contrasting incentives is that shareholders may also be concerned with attracting flows from clients: see Dasgupta, Fos, and Sautner (2021) for a survey of the literature. Indeed, the press releases and significant media outreach following such proposals suggests that shareholders may have reputational reasons for launching them.⁹ Thus, they may initiate proposals even if they erode firm value, and even if they’re not the best way to achieve ES goals because investors are less informed than management, thus leading to micromanagement. Dasgupta and Prat (2006) show that mutual funds may trade, even if they are uninformed, in order to convince their clients that they have good trading ideas due to being skilled. If we apply these models to engagement rather than trading, they deliver the prediction that investors may launch shareholder proposals even if they are uninformed. The one-size-fits-all nature of some ES proposals is consistent with this prediction.

None of the above means that shareholder engagement is bad; indeed, there is ample evidence that it can improve both long-term financial and social value, even if the engagement is primarily

⁹ Note that there may be innocuous reasons for such outreach, e.g. to garner support from other investors and the public.

motivated by financial considerations (e.g. Brav et al., 2007, 2015, 2018) or social considerations (Dimson, Karakas, and Li, 2015). Instead, just like many other ESG issues, economic principles suggest that engagement is not black-and-white. Shareholder engagement is likely to create financial value when the shareholder has a large stake in the company, social value if the shareholder has similar objectives to other shareholders, and both types of value if the engagement is initially private and thus less likely to be flow-motivated.

2.9 You Improve ESG Performance By Paying For ESG Performance (*No, paying for some ESG dimensions will cause firms to underweight others*)

The linkage of executive pay to ESG targets is one of the fastest-growing ESG trends. Based on a large-scale survey, PwC (2022) found that 92% of US companies and 72% of UK firms now use ESG metrics in incentive plans; based on interview evidence, a third of companies implemented ESG-linked pay for the first time in 2022. (In practice, these schemes are almost exclusively linked to ES performance, but we will use the term ESG as this is how they're often described to the public). Some investors have publicly stated that all firms should link their pay to ESG performance, and regulators are contemplating making such an association mandatory. It seems a no-brainer – the best way to improve performance is to pay for performance; you get what you pay for.

However, as essayist H.L. Mencken is often paraphrased, “Every complex problem has a solution which is simple, direct, plausible – and wrong,” and this may be the case with ESG-linked pay. A seminal paper by Holmstrom and Milgrom (1991) studies the case of multi-tasking, where a CEO has multiple responsibilities, each of which is associated with a different output. Paying for output A indeed increases the CEO's effort on task A – but reduces her effort on other outputs. You might think that a solution is to also pay her for outputs B, C, and D, but if these outputs are harder to measure, then doing so won't increase effort by much. If you can't pay for those other outputs, it

might be better not to pay for output A either, even though output A is measurable, because of the distortions that it will induce on the other tasks.

For example, let task A be increasing demographic diversity and task B be improving inclusion. The most common ESG metric in executive pay is demographic diversity. However, it's almost impossible to measure the inclusiveness of a company's corporate culture, so pay is very rarely linked to inclusion. Thus, diversity-linked pay may lead to box-ticking diversity – companies hiring minorities but putting little effort into creating an inclusive culture. In Holmstrom and Milgrom (1991), tasks A and B indirectly conflict because the CEO has limited time and so allocating it to one task is necessarily at the expense of other tasks. Moreover, in many ES applications, they directly conflict. For example, one way to increase diversity in senior management is to hire minorities from the outside, but doing so may reduce inclusion since the firm is now led by outsiders who don't fully understand employees' key concerns.

Moreover, there may be direct conflicts with other ES tasks. If task C is increasing cognitive diversity, this may conflict with demographic diversity. A white male may provide cognitive diversity due to being from a different socioeconomic background or age cohort, but reduce demographic diversity. If task D is improving workforce quality, then paying only for diversity may encourage companies to hire diverse applicants irrespective of quality.

This is an example of the Theory of the Second Best, a general theorem in welfare economics. The first-best solution is to achieve optimal outcomes in all areas. However, if optimality is impossible in one area, the second-best solution may be to allow for suboptimalities in other areas to prevent distortions. For example, if it's impossible to incentivize task B because it is difficult to measure, it may be optimal not to incentivize task A – even though you can – as doing so will further disincentivize B and lead to lower welfare overall, even when taking into account the superior

performance on A. It is also linked to the prior concern in Section 2.6 of “partial equilibrium” thinking, where fixing one problem can exacerbate others.

What’s the solution? One might argue that ESG-linked pay is problematic, but it’s better than any other alternative, to paraphrase a quote typically applied to capitalism.¹⁰ If you paid for nothing at all, you’d get no performance on any of the tasks, so it’s better to pay for task A. Another argument is that CEO pay is already linked to a multitude of financial targets, which cause them to skew their efforts to the financial at the expense of the non-financial. ESG-linked pay is the only way to redress the balance. But there is a solution – to scrap all bonuses, on both financial and nonfinancial targets, and instead pay CEOs with shares that they are required to hold for many years (e.g. five to seven) and must retain beyond their departure. Since material ESG factors ultimately improve the long-term stock price, this holds them accountable for material ESG issues, regardless of how measurable they are. Indeed, Flammer and Bansal (2017) show that long-term pay improves not only financial performance, but ES performance as well, and conduct a regression discontinuity to demonstrate a causal relationship.

2.10 Market Failures Justify Regulatory Intervention (*No, regulatory intervention is only justified when market failure exceeds regulatory failure*)

Externalities are a market failure. The free market fails to bring about an efficient allocation of resources since, left to their own devices, companies ignore the externalities they exert on society. Some ES advocates argue that market failures automatically mean that regulatory intervention is called for. However, mainstream economics has uncovered multiple sources of regulatory failure, and

¹⁰ The quote is “capitalism is the worst economic system, except for all the others”. This is based on Winston Churchill’s phrase that “democracy is the worst form of government, except for all those others that have been tried.”

so government intervention is only optimal when the market failure is so large that it exceeds even the regulatory failure.

One regulatory failure arises from crowding out. For example, the government directly investing in clean energy may crowd out similar investments from the private sector. Moreover, more subtle forms of crowding out also exist. Some E advocates call for “green capital requirements” that require banks to hold more capital against loans that bear climate risk (although whether this is physical or transition risk is often unclear). This seems to kill two birds with one stone – not only does it increase financial stability, by requiring banks to hold more capital against risky loans, but it also makes it less attractive for banks to lend to brown clients, thus encouraging them to shift their lending to green. However, Oehmke and Opp (2022) show that such requirements can actually crowd out green lending – since the bank has to hold more capital against the brown loans, it has less capital to finance green loans.

A second type of crowding out arises because governments can only regulate measurable factors. This may crowd out performance on difficult-to-measure factors, similar to the problem of paying for ES performance discussed in Section 2.9. For example, the cities of Portland and San Francisco tax companies whose pay ratios exceed a certain level, due to the concern that high ratios worsen social inequality. However, such a ratio captures only pay, and not other elements of the employee value proposition such as flexible working hours, hybrid working, on-the-job training, and corporate culture. It can lead to companies increasing pay and underinvesting in these other dimensions.

Regulatory failures become even more severe when we also consider the impact on firm performance.¹¹ One cost is the mere compliance cost of regulation. The Securities and Exchange

¹¹ Laws are often motivated by the desire to address externalities, rather than improve firm performance, as companies have sufficient incentives to take value-increasing actions by themselves. However, policymakers should also consider the impact on companies, so that any correction of externalities is not at too great a financial cost.

Commission estimated a first-year implementation cost of \$1.3 billion and ongoing annual costs exceeding \$520 million for the pay ratio disclosure that it mandated from 2018.

A second cost is the distortion in decisions arising because regulations can only target measurable factors. Earlier we discussed how this limitation affects social performance – it crowds out social performance on non-measurable dimensions. Moreover, this limitation can also worsen financial performance, as a company is forced to prioritize the regulated dimension at the expense of other performance-relevant ones. For example, when considering the appointment of new directors, a board may consider gender diversity, other forms of diversity, and other characteristics (e.g. experience, contacts) beyond diversity. Mandating diversity quotas may force firms to underweight the second and third factors to the detriment of business performance: Ahern and Dittmar (2013) find negative performance consequences of Norway’s gender quota, and Greene, Intintoli, and Kahle (2020) document similar results for California.

The existence of regulatory failure does not mean that governments should refrain from regulating. Instead, it merely means that regulation has costs and benefits, and that there are no one-size-fits-all solutions, e.g. always to regulate or never to regulate. Even if there is regulatory failure, regulation is optimal if the market failure is even larger, such as with climate change. Moreover, regulatory failure is reduced when the regulation is directly targeted at the underlying market failure. For climate change, the market failure is that carbon emissions are themselves an externality, not that carbon emissions increase bank risk and thus the likelihood of a bank collapse which creates externalities. Indeed, Oehmke and Opp (2022) show that a carbon tax is more effective than a green capital requirement.

3. Conclusion

This article has highlighted how the application of finance and economics principles significantly enhances our understanding of major ESG issues. This application only enhances our understanding; it does not complete our understanding. Thus, the new ideas being proposed and studied by academics and practitioners have great value. However, the excitement and urgency of ESG issues does not mean that we should shoot from the hip and go with gut feel or untested ideas to appear radical. As we learned from the fable of the tortoise and the hare, we need more haste, but less speed.

References

- Aggarwal, Ishani, Anita Williams Woolley, Christopher F. Chabris and Thomas W. Malone (2019): “The impact of cognitive style diversity on implicit learning in teams.” *Frontiers in Psychology* 10, 112.
- Aghion, Philippe and Jean Tirole (1997): “Formal and real authority in organizations.” *Journal of Political Economy* 105, 1–29.
- Agrawal, Ashwini K. (2012): “Corporate governance objectives of labor union shareholders: Evidence from proxy voting.” *Review of Financial Studies* 25, 187–226.
- Ahern, Kenneth R. and Amy K. Dittmar (2012): “The changing of the boards: The impact on firm valuation of mandated female board representation.” *Quarterly Journal of Economics* 127, 137–197.
- Baker, Steven D., Burton Hollifield, and Emilio Osambela (2022): “Asset prices and portfolios with externalities.” *Review of Finance* 26, 1433–1468.
- Bolton, Patrick and Marcin Kacperczyk (2021): “Do investors care about carbon risk?” *Journal of Financial Economics* 142, 517–549.
- Bolton, Patrick, José Scheinkman and Wei Xiong (2006): “Executive compensation and short-termist behaviour in speculative markets.” *Review of Economic Studies* 73, 577–601.
- Boustanifar, Hamid and Young Dae Kang (2022): “Employee satisfaction and long-run stock returns, 1984-2020.” *Financial Analysts Journal* 78, 129–151.
- Brav, Alon, Wei Jiang, Frank Partnoy, and Randall Thomas (2008): “Hedge fund activism, corporate governance, and firm performance.” *Journal of Finance* 63, 1729–1775.
- Brav, Alon, Wei Jiang and Hyunseob Kim (2015): “The real effects of hedge fund activism: Productivity, asset allocation, and labor outcomes.” *Review of Financial Studies* 28, 2723–2769.
- Brav, Alon, Wei Jiang, Song Ma and Xuan Tian (2018): “How does hedge fund activism reshape corporate innovation?” *Journal of Financial Economics* 130, 237–264.
- Brealey, Richard A., Stewart C. Myers, Franklin Allen and Alex Edmans (2022): *Principles of corporate finance* (14th edition). McGraw-Hill.
- Burkart, Mike, Denis Gromb, and Fausto Panunzi (1997): “Large shareholders, monitoring, and the value of the firm.” *Quarterly Journal of Economics* 112, 693–728.
- Dasgupta, Amil, Vyacheslav Fos, and Zacharias Sautner (2021): “Institutional investors and corporate governance.” *Foundations and Trends in Finance* 12, 276–394.

- Dasgupta, Amil and Andrea Prat (2006): “Financial equilibrium with career concerns.” *Theoretical Economics* 67–93.
- Denis, Diane (2016): “Corporate governance and the goal of the firm: In defense of shareholder wealth maximization.” *The Financial Review* 51, 467–480.
- Dimson, Elroy, Oğuzhan Karakaş, and Xi Li (2015): “Active ownership.” *Review of Financial Studies* 28, 3225–3268.
- Duchin, Ran, Janet Gao, and Qiping Xu (2022): “Sustainability or greenwashing: Evidence from the asset market for industrial pollution.” Working Paper, Boston College.
- Edmans, Alex (2011): “Does the stock market fully value intangibles? Employee satisfaction and equity prices.” *Journal of Financial Economics* 101, 621–640.
- Edmans, Alex (2012): “The link between job satisfaction and firm value, with implications for corporate social responsibility.” *Academy of Management Perspectives* 26, 1–19.
- Edmans, Alex (2020): *Grow the pie: How great companies deliver both purpose and profit*. Cambridge University Press.
- Edmans, Alex (2023): “The end of ESG.” *Financial Management*, forthcoming.
- Edmans, Alex, Vivian W. Fang, and Katharina A. Lewellen (2017): “Equity vesting and investment.” *Review of Financial Studies* 30, 2229–62271.
- Edmans, Alex, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov (2012): “Dynamic CEO compensation.” *Journal of Finance* 67, 1603–1647.
- Edmans, Alex and Clifford G. Holderness (2017): “Blockholders: A survey of theory and evidence.” Chapter 8 in Benjamin E. Hermalin and Michael S. Weisbach (eds.), *Handbook of the Economics of Corporate Governance*, Elsevier: Amsterdam, 541–636.
- Edmans, Alex and Marcin Kacperczyk (2022): “Sustainable finance.” *Review of Finance* 26, 1309–1313.
- Edmans, Alex, Doron Levit, and Jan Schneemeier (2022): “Socially responsible divestment.” Working paper, London Business School.
- Edmans, Alex, Darcy Pu, Chendi Zhang, and Lucius Li (2022): “Employee satisfaction, labor market flexibility, and stock returns around the world.” Working paper, London Business School.
- Flammer, Caroline and Pratima Bansal (2017): “Does a long-term orientation create value? Evidence from a regression discontinuity.” *Strategic Management Journal* 38, 1827–1847.
- Fried, Jesse M. (2021): “Will Nasdaq’s diversity rule harm investors?” Working paper, Harvard University.

- Friedman, Milton (1970): “The social responsibility of business is to increase its profits.” *New York Times*, September 13.
- Gantchev, Nickolay and Mariassunta Giannetti (2021): “The costs and benefits of shareholder democracy: Gadflies and low-cost activism.” *Review of Financial Studies* 34, 5629–5675.
- Giglio, Stefano, Bryan Kelly, and Johannes Stroebe (2021): “Climate finance.” *Annual Review of Financial Economics* 13, 15–36.
- Gosling, Tom (2019): “Air miles, road miles.”
- Graham, John R., Campbell R. Harvey, and Shivaram Rajgopal (2005): “The economic implications of corporate financial reporting.” *Journal of Accounting and Economics* 40, 3–73.
- Greene, Daniel, Vincent J. Intintoli, and Kathleen M. Kahle (2020): “Do board gender quotas affect firm value? Evidence from California Senate Bill No. 826.” *Journal of Corporate Finance* 60, 101526.
- Holmstrom, Bengt and Paul Milgrom (1991): “Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design.” *Journal of Law, Economics & Organization* 7, 24–52.
- Jensen, Michael C. and William H. Meckling (1976): “Theory of the firm: Managerial behavior, agency costs, and ownership structure.” *Journal of Financial Economics* 3, 305–360.
- Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny (1994): “Contrarian investment, extrapolation, and risk.” *Journal of Finance* 49, 1541–1578.
- Oehmke, Martin and Marcus M. Opp (2022): “Green capital requirements.” Working Paper, London School of Economics.
- Pastor, Lubos, Robert F. Stambaugh and Lucian A. Taylor (2021): “Sustainable investing in equilibrium.” *Journal of Financial Economics* 142, 550–571.
- Pastor, Lubos, Robert F. Stambaugh and Lucian A. Taylor (2022): “Dissecting green returns.” *Journal of Financial Economics* 146, 403–424.
- Polk, Christopher and Paola Sapienza (2023): “The stock market and corporate investment: A test of catering theory.” *Review of Financial Studies* 22, 187–217.
- PRI (2016): “A practical guide to ESG integration for equity investing.”
- PwC (2022): “Paying for good for all.”
- Renneboog, Luc, Jenke ter Horst, and Chendi Zhang (2008): “The price of ethics and stakeholder governance: The performance of socially responsible mutual funds.” *Journal of Corporate Finance* 14, 302–322.

Shleifer, Andrei and Robert W. Vishny (1986): "Large shareholders and corporate control." *Journal of Political Economy* 94, 461–488.

Zwiebel, Jeffrey (1995): "Block investment and partial benefits of corporate control." *Review of Economic Studies* 62, 161–185.