COUNTERPRODUCTIVE IMPACT INVESTING: THE IMPACT ELASTICITY OF BROWN AND GREEN FIRMS*

SAMUEL M. HARTZMARK

KELLY SHUE

February 2023

First Draft November 2022

Abstract

We develop a new measure of impact elasticity, defined as a firm's change in environmental impact due to a change in its cost of financing. We show empirically that a reduction in financing costs for firms that are already green leads to small improvements in impact at best. In contrast, increasing financing costs for brown firms leads to large negative changes in firm impact. Thus, impact investing that directs capital away from brown firms and toward green firms may have the counterproductive effect of worsening aggregate firm impact on society. We further show that the impact investing movement primarily rewards green firms for economically trivial improvements in their environmental impact, thereby providing very weak incentives for brown firms to become more green. Impact investor flows and engagement that targets the incentives of green firms would be more effective if targeted at brown firms.

*Respective contact information and affiliations are: samuel.hartzmark@bc.edu, Boston College and NBER; kelly.shue@yale.edu, Yale School of Management and NBER. We thank Yiyuan Wang for excellent research assistance. This research was funded in part by the International Center for Finance at Yale SOM.

A wise man gets more use from his enemies than a fool from his friends. -Baltasar Gracian

I. Introduction

There has been dramatic growth in financial investments made with sustainable impact objectives. While a variety of tactics have been employed, the dominant strategy involves making investments in firms that are perceived to be "green" (positive environmental impact) and divesting away from firms that are perceived to be "brown" (negative environmental impact). To the extent that impact investing can meaningfully change firm financing conditions, it rewards green firms by lowering their cost of financing and punishes brown firms by raising their cost of financing. A consistent message from impact-driven investors and fund managers is the hope that impact investing will motivate all firms to become more green, thereby improving the aggregate impact of firms on society.

The success of the impact investing movement depends critically on how firms alter their behavior in response to changing financing conditions. If green firms provide the most *additional* green in response to a reduction in their cost of financing, impact investors should reward green firms by investing in them. Likewise, if brown firms provide the largest reduction of brown in response to an increase in their cost of financing, impact investors should punish brown firms by divesting away from them. In this paper, we develop a new measure of "impact elasticity," defined as a firm's change in impact due to a change in its cost of financing. Without knowledge of the relative impact elasticities of green and brown firms, it is unclear which and how firms should be targeted by impact investors.

This paper shows that, if the dominant impact investing strategy of directing capital away from brown firms toward green firms succeeds in changing firms' financing costs, such a strategy would likely have the counterproductive effect of worsening aggregate firm impact on society. We show empirically that firms that are considered green based on their green house gas emissions have little scope for further improvement in their impact. Green firms exhibit low variability in impact and have close-to-zero impact elasticity with respect to changes in their cost of financing. In contrast, brown firms have substantially greater scope for change (about 180 times as much as similarly-sized green firms) and exhibit large negative impact elasticities with respect to their cost of financing. Brown firms become meaningfully less brown in response to easier access to capital and substantially more brown if pushed toward financial distress. We further show that the dominant impact investing strategy provides very weak incentives for brown firms to become less brown. Instead, the impact investing movement primarily rewards green firms for economically trivial reductions in their already low levels of emissions.

Whether the massive inflows of money into impact investing over the past several decades have had a major impact on firm cost of financing is open to debate. For example, Chava (2014) presents empirical evidence of green firms having lower cost of financing, whereas Berk and van Binsbergen (2021) argue the impact-driven investment flows are not large enough to impact market conditions given that non-green funds can offset these flows. Our paper suggests that the best case scenario for aggregate firm impact is the one where the dominant impact investing strategy has not yet shifted firms' cost of financing. To the extent that such a strategy succeeds in changing firms' cost of financing, it would have the counterproductive effect of lowering the overall green impact of firms on society.

A simple case study may help illustrate our paper's intuition. Travelers is an insurance firm in the S&P 500 that looks spectacular on environmental, social, and governance (ESG) metrics. Travelers widely advertises its low green house gas emissions. In 2021, it emitted 33,477 metric tons of carbon, which is about 1 ton per million dollars of revenue.[1] At the opposite extreme lies Martin Marietta Materials, another S&P 500 firm that supplies heavy building materials. Amongst ESG ratings providers, Martin Marietta is uniformly considered poor. In 2021, it emitted about 5.1 trillion tons of carbon, corresponding to about 1,000 tons per million dollars of revenue.[2] Relative to Travelers, Martin Marietta has 1000 times as much emissions intensity, measured as emissions scaled by revenue.

The most common impact investing strategy dictates that investors with sustainable objectives should invest in Travelers and avoid Martin Marietta. With that said, if money flows toward Travelers allowing further investments in green projects at subsidized rates, where would it go? If Travelers were able to cut emissions by 100%, it would be equivalent to Martin Marietta cutting its emissions by 0.1%. As an insurance firm, Travelers would also struggle to become more brown. If it ceased to care at all about carbon emissions, perhaps its emissions might double, but this would still mean that Travelers' overall emissions intensity would be less than 1% of those of Martin Marietta.

On the other hand, Martin Marietta has the capability of becoming much more green or brown. While the company emits a large amount of carbon, it does so after having made significant green investments to improve its environment impact. After significant investment, it cut its emission per ton of cement from 0.836 to 0.77. While this might seem minor, given they produced about 3.5 million tons of cement in 2021, this decreased emissions by roughly 7 times Travelers's annual total. Martin

[1] https://sustainability.travelers.com/iw-documents/sustainability/Travelers_ESGAnalystData2021.pdf

[2] https://mcdn.martinmarietta.com/assets/sustainability/flip/sustainability2021-f/index.html

Marietta has also considered a number of green technologies for future adoption, although the firm currently deems these green investments to be unprofitable absent additional financial incentives.[3] Thus, if investors were willing to make such green investments profitable, Martin Marietta's carbon emissions could decrease further (recall that a 1% decrease would be greater than the entirety of Travelers's emissions). Similarly, if the market forced the firm to worry about its short-term financial survival, the company could invest more into its existing brown projects which deliver more front-loaded cash flows compared to green investment projects. Simply reversing its efficiencies per dollar on carbon emissions since 2016 would result in an increase in emissions of approximately one million tons, equal to 30 times Travelers' annual level of emissions. Overall, Martin Marietta has the potential to significantly decrease emissions, and if pushed to cut corners, could meaningfully increase emissions as well.

In our empirical analysis, we show the intuition of this example is reflective of the broader data. When asked, impact investors report environmental concerns as their main priority (Hartzmark and Sussman, 2019). We measure firm environmental impact using green house gas (e.g., carbon) emissions, because it is a standardized measure that clearly matters to impact investors, and because carbon emissions directly contribute to the risk of global climate change. The importance of carbon emissions to impact investors is reflected in recent SEC communications concerning mandatory disclosure of emissions in the portfolio holdings of all funds that consider environmental factors.[4] Similarly, a recent cover story in *The Economist* argued that emissions were so important that they should be the sole focus of environmental focused investors.[5]

We begin by showing that brown firms have much greater year-to-year variability in their impact than green firms. Measures of variability provide a useful upper bound on the absolute value of a firm's impact elasticity; variability that is close to zero implies that the impact elasticity must also be close to zero. We divide firms into quintiles by their emissions intensity (defined as scope 1 and scope 2 emissions scaled by revenue, hereafter referred to as "emissions" for brevity) in the previous year,

[3]For example, the firm cites the friendly regulatory environment in Europe for green investment and argues that it would be unprofitable to implement similar changes in the US without such regulation.

[4]"ESG-focused funds that consider environmental factors... would be required to disclose the carbon footprint and the weighted average carbon intensity of their portfolio. The requirements are designed to meet demand from investors seeking environmentally focused fund investments for consistent and comparable quantitative information regarding the GHG emissions associated with their portfolios and to allow investors to make decisions in line with their own ESG goals and expectations." https://www.sec.gov/files/ia-6034-fact-sheet.pdf

[5]The article states "The environment is an all-encompassing term, including biodiversity, water scarcity and so on. By far the most significant danger is from emissions, particularly those generated by carbon-belching industries. Put simply, the e should stand not for environmental factors, but for emissions alone." https://www.economist.com/leaders/2022/07/21/esg-should-be-boiled-down-to-one-simple-measure-emissions

with the top and bottom quintiles representing green and brown firms, respectively. We show that green firms have very low emissions and very little year-to-year variation in that level. In contrast, the average brown firm has 260 times as much emissions as the average green firm, and experiences approximately 150 times larger absolute changes in emission levels from year to year. The average absolute annual *change* in emissions for a brown firm is equal to the average *level* of emissions from 35 green firms combined.

Some impact investment funds adjust their allocations so that their industry or sector weights resemble those of the overall market, and they invest in green firms and divest away from brown firms within each industry. We show that brown firms within an industry also exhibit much greater variability in impact than green firms within the same industry. Thus, "industry-adjusted" impact investment is nevertheless tilted toward green firms with little scope to change their impact.

Next, we examine the impact elasticity of green and brown firms to changes in their cost of financing. To obtain sufficient statistical power, we examine a broad class of shocks to firms' cost of financing without restricting our attention to changes in financing costs that are due to impact investment. The underlying assumption is that, if the dominant impact investing strategy succeeds in altering firm's financing costs, firms would react in a similar fashion as they react to other changes in their financing costs. We explore potential violations of this assumption later in this paper.

Across a variety of tests, we find that brown firms have greater negative impact elasticity than green firms. First, brown firms exhibit greater reductions in their emissions following improvements in their real or financial performance. Stronger firm performance likely eases the firm's financial constraints and lowers the firms' cost of financing. To establish causality, we examine the relation between firm emissions and the firm's industry average performance, calculated excluding the focal firm. The intuition is that industry performance shocks strongly affect firm-level performance and access to finance, but individual firm choices of emissions should not affect the industry average performance calculated excluding the focal firm. We find that brown firms are much more elastic to industry shocks to firm performance than green firms.

Second, we examine how firm emissions respond to financial distress shocks, as proxied by interest coverage, Altman Z-scores, or industry performance in the lowest decile within our sample. We find that brown firms react to financial distress by increasing their emissions. In contrast, green firms exhibit a much smaller and less significant response, sometimes in the opposite direction.

Third, we measure impact elasticity with respect to the firm's implied cost of capital (ICC). The

ICC is the internal rate of return that equates the firm's market value to the present value of its expected future cash flows. A higher ICC corresponds to higher cost of financing. Because ICCs have been shown to be noisy with conflicting estimates across methodologies (Lee et al., 2021), we present these tests as a complement to our earlier analyses. Across a variety of ICC measures, we find that brown firms exhibit a large negative impact elasticity with respect to their ICC while green firms exhibit elasticities close to zero. To establish causality, we examine the relation between firm emissions and the firm's industry value-weighted average ICC, calculated excluding the focal firm. We again find the brown firms substantially increase emissions following an increase in industry ICC, while green firms exhibit an economically small and insignificant response.

These empirical patterns are consistent with basic corporate finance theory. Brown firms likely face a choice between dark-brown investment projects (e.g., continuing or expanding existing high-pollution operations, or cutting corners on pollution abatement) and light-brown investments (e.g., shifting toward cleaner production or green energy). Because the light-brown investment project entails a departure from existing production methods, it likely requires investment in new capital which costs more up front and delivers back-loaded cash flows compared to the dark brown project. Financial distress or an increase in the cost of financing will make short-term cash flows more attractive relative to long-run cash flows. Thus, an increase in the cost of financing causes the dark-brown project to look relatively more attractive, leading brown firms to have a negative impact elasticity.[6] In contrast, a green firm likely operates in a line of business (e.g., insurance in the case of Travelers) where the firm cannot generate substantial environmental regardless which investment projects it chooses to pursue, leading green firms to have impact elasticities close to zero.

Finally, we investigate a potential concern with our conclusions. Our measure of impact elasticity captures how firms react to general shocks to their cost of financing. However, changes in the cost of financing due to impact investing may differ from other shocks because impact investing could (in theory) incentivize firms to become more green. Specifically, brown firms may choose to become more green if impact investors reward their change by lowering their cost of financing in the future.

We believe that an impact investing strategy targeted at changing firms' incentives could be very

[6] For instance, Thomas et al. (2022) show that firms cut back on pollution abatement costs to meet earnings targets. The simple intuition here is also similar in spirit to a more sophisticated and specialized model presented in Eisfeldt and Rampini (2007), which shows that financial distress or an increase in the cost of capital induces firms to favor investment in "used capital" (e.g., older technologies or existing projects) that is cheaper up front and may require greater maintenance costs in the future. The ability of brown firms to choose between projects with more or less front-loaded cash flows is underscored by recent evidence from Gilje et al. (2020) showing that financial distress causes firms in the oil and gas industry to pull forward drilling in existing oil wells at the expense of long-run project returns.

promising. However, we show empirically that the dominant impact investing strategy in practice provides no such incentives. Using data on the holdings of impact investment funds over the past three decades, we find that impact investment funds indeed overweight firms that have improved their impact over the past several years, consistent with these funds rewarding firms who transition to becoming more green. However, changes in impact are measured in the wrong units. Impact investors appear to suffer from a proportional thinking bias (Tversky and Kahneman (1981) and Shue and Townsend (2021)) in which they reward firms with large *percentage* reductions in emissions rather than large *level* reductions in emissions. Popular ESG environmental ratings similarly reward percentage reductions in emissions rather than level reductions in emissions.

The distinction between percentage and level changes in emissions is important because, holding constant firm size, the average brown firm emits 260 times as much pollution as the average green firm. Comparing a brown and green firm of equal size, an increase in emissions by a brown firm of 1% has the same actual environmental impact as an increase in emissions by a green firm of 260%. It is also much easier for a green firm to purchase a small quantity of carbon offsets to completely offset its initially low level of emissions and become carbon neutral. However, this 100% reduction in emissions is far less economically meaningful than a brown firm reducing its emissions by a mere 1%.

Perhaps most surprisingly, we find that impact investment funds and ESG ratings reward green firms much more than brown firms for the same percentage reduction in emissions (the logic above implies that it should be the other way around). This additional mistake is consistent with a behavioral affect heuristic (e.g., Slovic et al., 2007), in which impact investors choose to disassociate from or punish brown firms that they dislike, despite the fact that brown firms have the greatest scope to change their environmental impact. The view that impact investment in practice largely ignores the potential for brown firms to improve is underscored by recent evidence from Cohen et al. (2020) showing that energy sector firms are considered brown and excluded from impact investing portfolios, despite the fact that these firms produce the most highly-cited green patents. These green innovations have the potential to dramatically improve firm impact in the long run. While Cohen et al. (2020) does not measure the impact elasticity of brown firms, their results complement ours by suggesting that brown firms have the greatest scope to change their long term impact.

Overall, we show that the dominant impact investment strategy targets green firms with the least scope to change their environmental impact. To the extent that impact investing increases the cost of financing for brown firms and/or pushes them toward financial distress, brown firms react by becoming more brown. Because it measures changes in impact in the wrong units, the dominant impact investment strategy also provides very weak incentives for brown firms to become less brown. We caveat these conclusions by noting that our analysis is meant only as a critique of the dominant impact investing strategy and not all impact strategies. Our analysis suggests that impact investing flows and engagement that targets the incentives of green firms would be more effective if targeted at brown firms.[7]

Our paper complements recent theoretical work by Edmans et al. (2022), which argues that an optimal impact investing strategy should tilt away from brown industries but be willing to hold brown firms that take actions to improve their impact. Their model describes a managerial agency problem in which a self-interested manager of a brown firm benefits from a short-term increase in the stock price. Our paper is similar in spirit because we show that divestment away from brown firms can be suboptimal. However, our conclusions does not require a managerial agency conflict. Rather, brown and green can firms fundamentally differ in their impact elasticities due to differences in their investment opportunity set.

Our evidence that brown firms have a negative impact elasticity is consistent with earlier evidence in Hong et al. (2012) showing that firms on average do more good when they are doing well, i.e., when their financial constraints are relaxed. We differ in focus from their paper by showing that the magnitude of the relation between firm goodness and financial constraints strongly varies by whether the firm is initially brown or green, which has important implications for the effectiveness of the dominant impact investing strategy.

Another related literature has highlighted problems in the current system of evaluating firm ESG and sustainability (e.g., Berg et al., 2019). Our analysis shows that ESG ratings are flawed because they evaluate changes in emissions in percentage units, thus favoring green firms with little scope for real improvement. Heath et al. (2021) show that socially responsible investment funds buy firms with green characteristics, but these characteristics do not meaningfully improve after they are purchased. Our paper offers a complementary explanation for why green firms do not improve—green firms have little scope to improve, even when incentivized to do so.

Our paper also contributes to the broader academic literature exploring the optimal impact investment strategy. Much of the debate concerns the choice between divestment strategies (holding green firms and avoiding brown firms) versus engagement strategies (direct corporate governance interac-

[7]While engagement is not as common in the real world, there are some notable examples such as Engine No. 1, which has successfully engaged with Exxon Mobil to change its impact.

tions such as voting your shares or obtaining board seats). In practice, the divestment camp represents the bulk of impact investing (see e.g. Pástor et al. (2021)), though the engagement strategy has been widely discussed in the academic literature (e.g. Broccardo et al. (2020)). Our paper complements the existing literature in three ways. First, the theory literature generally models green investment with costs and/or benefits that are not a function of the firm's current level of green. Our paper suggests that adding an elasticity term that is a function of the firm's current level of green would both more closely reflect reality and possibly lead to different conclusions. Second, our findings provide empirical support for the idea that engagement efforts directed at brown firms will be more effective than engagement efforts directed at green firms. Because brown firms have the greatest scope to become more brown or green, impact investing aimed at changing project-specific incentives for brown firms may be particularly effective. Third, while other research has argued that a divestment strategy may ineffective (e.g., Berk and van Binsbergen (2021) and Broccardo et al. (2020)), we show conditions where such a strategy can be actively counterproductive.

II. Framework: Impact Elasticity

We define impact elasticity as the firm's change in impact in response to a change in its cost of financing:

$$\text{Impact Elasticity} \equiv \frac{\partial \text{ Impact}}{\partial \text{ Cost of Financing}}.$$

Our primary contribution is to document heterogeneity in the impact elasticity as a function of the firm's initial level of green. In this paper, we focus on firm impact as measured by green house gas emissions. Greater emissions implies a more negative firm environmental impact on society. Therefore, an increase in emissions following a positive shock to a firm's cost of financing translates to a negative impact elasticity.

To obtain sufficient statistical power, we examine a broad class of shocks to firms' cost of financing without restricting our attention to changes in financing costs that are due specifically to impact investing. The underlying assumption is that, if the dominant impact investing strategy succeeds in altering firm's financing costs, firms would react in a similar fashion as they react to other changes in their financing costs.

Two important considerations apply to our measure of the impact elasticity. First, changes in the cost of financing due to impact investing may differ from other shocks to the cost of financing because impact investing could incentivize firms to become more green. In other words, firms could be motivated to become green by the inverse of the impact elasticity: the future change in a firm's cost of financing in response to a change in the firm's environmental impact. For example, a brown firm may choose to pursue green investment projects because it anticipates that impact investors will reward its positive change in impact by lowering its cost of financing in the future. While we believe this incentive channel could be promising in theory, we will present empirical evidence that the dominant impact investing strategy provides no such incentives for brown firms. Instead, impact investors and ESG ratings primarily reward firms that are already green for economically trivial, but large percentage, reduction in their emissions.

Second, impact elasticity is a measure of firm-level changes in impact in response to firm-level changes in the cost of financing. In theory, firms should assess potential investment projects using a project-specific cost of capital that reflects project-specific risk rather than a firm-wide cost of capital. The impact elasticity measure can accommodate a project-specific valuation method. We assume that firm-level shocks to the cost of financing shifts the firm's project-specific discount rates equally across all projects. We leave the study of the effectiveness of alternative impact investing strategies, such as those that subsidize specific green investment projects, to future research. For now, we note that the dominant impact investing strategy of investing in green firms and divesting away from brown firms is an example of a firm-level shock to the cost of financing, and its effect would depend on the firm-level impact elasticity.

III. Data

Our data sample covers the years 2002 to 2020. Data on green house gas (GHG) emissions comes from S&P Global Trucost. GHG emissions are gas emissions that trap heat in the atmosphere and contribute to the risk of global climate change. The primary green house gases emitted in the U.S. in 2020 are carbon dioxide (79%), methane (11%), nitrous oxide (7%), and fluorinated gases (3%) such as hydrofluorocarbons and perfluorocarbons.[8] We use data on scope 1 and 2 emissions. Scope 1 emissions are the most directly tied to the firm as they represent emissions from equipment that the firm owns. Scope 2 emissions are the indirect emissions associated with the purchase of electricity, steam and heating, so they occur at a location not controlled by the firm, but are directly tied to firm actions. We present our main results for total scope 1 and scope 2 emissions and present our main results separately for each type of emissions in the Appendix.

Accounting data concerning firm financial and real performance, leverage, earnings, and revenue

[8] See https://www.epa.gov/ghgemissions/overview-greenhouse-gases.

are obtained from the Compustat database. Price and return information is taken from the Center for Research in Securities Prices CRSP database.

A natural reason for firms to vary in their emissions is differences in size. It is not obvious that a larger firm should be considered less green purely because it emits more green house gases due to its larger scale. Therefore, we follow a convention commonly used by ESG ratings companies and prior studies, and focus on emissions intensity, defined as scope 1 and scope 2 emissions scaled by revenue. Hereafter, we refer to emissions intensity as just "emissions" for brevity, unless otherwise noted. We divide firms into quintiles by their emissions in the previous year, with the top and bottom quintiles representing green and brown firms, respectively. We classify firms in the middle three quintiles as neutral.

Analysis on the holdings of green investment funds are based on data generously shared by the authors of Cohen et al. (2020). We categorize an investment fund as green if it is defined as green based on any of the measures described in the Cohen et al. (2020) paper. Specifically, we classify funds as green if the fund name contains "ESG" or "Green" or if the fund is classified as an impact investment by either the Forum for Sustainable and Responsible Investment (USSIF) or Charles Schwab. We merge data on impact funds with data on monthly holdings by mutual funds from CRSP. For each stock-month, we measure the extent to which it is overweighted by green investment funds relative to the value-weighted market index. For example, if a stock represents 3% of the value of the combined portfolio of all green funds and 2% of the value of the total market portfolio, then we would estimate that green funds overweight the stock by 50%.

Data covering annual firm implied cost of capital (ICC) are generously shared by the authors of Lee et al. (2021). Following the advice on best practices in Lee et al. (2021), we use the Gebhardt et al. (2001) (GLS) mechanical ICC as our preferred measure of the ICC. This measure is also similar to those used in prior papers in the ESG literature (e.g. Chava (2014)). As shown in Lee et al. (2021), estimation of firm-level ICC is difficult and suffers from substantial noise due to the necessity of making assumptions about expected future cash flows and non-unique numerical solutions. To mitigate the problems of noise, we also show that our results are robust to using a simple average of four published ICC measures.

Table 1 presents summary statistics of the main variables used in our analysis. The distribution of the variables across the 10th, 50th, and 90th percentile indicate that emissions is extremely right skewed. Total raw emissions (unscaled) in the 90th percentile is equal to nearly 2000 times total raw emissions in the 10th percentile. After scaling by revenue to account for differences in firm size (our preferred measure of emissions), emissions in the 90th percentile is still 155 times as large as in the 10th percentile. The absolute value of level changes in emissions is similarly extremely right skewed, with the 90th percentile equal to 442 times the 10th percentile. In contrast, the absolute percentage change in emissions in far less skewed. However, as we will show, percentage changes in emissions are a poor measure of the true change in firm impact because green firms with extremely low levels of emissions tend to be associated with large percentage changes in their emissions. These summary statistics offer an early indication of our main results: brown firms have the greatest environmental impact and the greatest scope to change their impact.

IV. Results

A. Variability in firm impact

We begin our analysis by showing that brown firms have dramatically greater levels of and year-to-year variability in emissions compared to green firms. Variability provides a useful bound for our ultimate measure of interest, the impact elasticity, because variability that is close to zero implies that the impact elasticity must also be close to zero. Variability is also useful because one reasonable estimate for how much a firm can change its impact is how much it has changed in the past.

We divide firms into quintiles based on their level of emissions in each year, with quintile 5 representing firms with the lowest emissions. In subsequent analysis, we refer to firms in quintile 5 as "green," firms in quintile 1 as "brown," and firms in quintiles 2 through 4 as "neutral." In Figure 1 Panel A, which examines the raw level of carbon emissions (unscaled), the average brown firm releases more than 1,700 times as much carbon as the average green firm. Of course, these differences in carbon emissions across firms could be due to differences in firm size; it would be natural for larger firms to emit more carbon. Therefore, we use emissions scaled by same-year firm revenues as our baseline measure of emissions. In Panel B, we show that, even after scaling by revenues, brown firms release 261 times as much carbon per unit of sales as green firms in quintile 5. Using both the raw and scaled measures of emissions, neutral firms in quintiles 2 through 4 are associated with emissions levels much closer to that of green firms than of brown firms.

The dramatic differences in carbon emissions across the five quintiles shown in Figure 1 indicates that any analysis focusing on firm's annual *percentage* change in emissions is unlikely to be informative. Consider a green firm. Even if it doubled or halved its emissions in a single year, its change in behavior would have minimal environmental impact because its baseline level of emissions is several orders of magnitude smaller than the emissions of brown firms of similar size. In contrast, if the average brown firm doubled its emissions, the real carbon impact would be equivalent to the average green firm increasing its emission by 26,400%.

Thus, instead of focusing on percentage changes, we focus on the annual absolute level change in emissions (always scaled by revenues). In Table 2, we regress the annual absolute value of the change in emissions on indicators for the firm's emissions quintile, calculated in the previous year. Brown firms in quintile one represent the omitted reference category. The graphical analogue is presented in Figure 2.

We find a robust pattern in which brown firms exhibit substantially greater variability in their emissions. In the first column of Table 2, we include a fiscal year fixed effect to assess raw differences between green and brown firms after removing a general time effect. We find that the annual change in absolute emissions by brown firms exceeds that of green firms by approximately 180 tons per dollar of sales. Graphically in the left panel of Figure 2, this implies that the average annual variability of emissions by brown firms is 164 times the variability of emissions by green firms. Recall that green firms have a level of emissions intensity of 5, which means the average *variability* in brown firms is about 35 times the average emissions *level* of green firms. Variability in emissions declines monotonically from quintile 1 to 5, but the largest gap lies between firms in quintile 1 (brown) and quintile 2; the change in absolute emissions by brown firms exceeds that of firms in quintile 2 by 163 tons.

A potential concern with the results in Column 1 of Table 2 is that the extreme differences in emission is driven by small firms that have high emissions per unit of revenue, but low overall emissions due to there small size. To account for such a possibility, in Column 2, we weight observations by firm market value as a fraction CRSP market value in each year. We find similar patterns which shows that the large gap in variability of emissions between green and brown firms is not driven by small outlier firms.

In Column 3, we test whether the variability gap between brown and green firms holds within industry. We sort firms into quintiles according to their previous-year emissions rank within their SIC2 industry and control for SIC2 industry fixed effects. We continue to find a similar pattern in which brown firms within an industry year exhibit significantly greater variability in emissions than green firms. These within-industry results have the important implication that impact investment that is "adjusted for industry" will nevertheless target green firms within an industry that have substantially lower variability in emissions than brown firms within the same industry.

In supplementary results shown in the Appendix, we examine annual absolute changes in emissions separately for scope 1 (direct emissions from the company's owned or controlled sources) and scope 2 (indirect emissions from purchased energy) emissions. We find significant variability gaps between green and brown firms for both types of emissions. The gap for scope 1 emissions is much larger consistent with the fact that the level of scope 1 emissions is much larger than the level scope 2 emissions. We also find similar patterns using the absolute level of total emissions, without scaling by firm sales.

In Figure 3, we show that differences in variability in emissions across brown and green firms disappear if we measure changes in emissions using percentage changes instead of level changes. Across three possible specifications, green firms have close to or greater percentage variability in their emissions compared to brown firms. However, we caution that a large percentage change in the emissions of green firms is not economically meaningful, because green firms are associated with levels of emissions several orders of magnitude smaller than the level of emissions for similarly-sized brown firms.

B. Impact Elasticity

In this section, we estimate the impact elasticity of green and brown firms by examining how emissions by each type of firm changes following changes or shocks to the cost of financing. The dominant impact investing strategy seeks to lower the cost of financing for green firms by directing capital toward them and to increase the cost of financing for brown firms by divesting away from them. If green firms react to a lower cost of financing by becoming more green (i.e., green firms have a negative impact elasticity), and brown firms react to a higher cost of financing by becoming more green (i.e., brown firms have a positive impact elasticity), then we expect that the dominant impact elasticity strategy will improve aggregate firm impact on society if it succeeds in altering firms' cost of financing. However, as we will show, the actual impact elasticity of green firms is close to zero and the impact elasticity of brown firms is large and negative. Together, these measures imply that the dominant impact investing strategy may have the counterproductive effect of worsening aggregate firm impact on society.

We begin by examining the relation between changes in emissions and changes in firm performance. A positive performance shock is likely to ease the firm's access to financing. Likewise, a negative performance shock should increase the firm's shadow cost of capital.

One limitation of simply looking at the relation between firm emissions and firm performance is that any measured correlation could be driven by reverse causality (if becoming more green causes the firm to perform differently), or by omitted variables bias (e.g., if the arrival a talented new green-oriented CEO causes both a shift in green production and firm performance). To better estimate the causal impact of firm performance on firm emissions, we examine the relation between firm emissions and the firm's industry value-weighted average performance, calculated excluding the focal firm. The intuition is that industry performance shocks strongly impact firm-level performance, but individual firm choices should not have a strong impact on the industry average performance calculated excluding the focal firm.

In Table 3, we find that brown firms are significantly more elastic to shocks to firm performance than green firms across a variety of specifications. Column 1 shows a large and highly significant negative coefficient on the firm annual return for brown firms indicating that as firm financial performance improves, brown firms reduce their emissions. Symmetrically, the negative coefficient implies that negative performance by brown firms is associated with an increase in emissions. In terms of magnitude, the coefficient of -49.95 implies that a 10% financial return for brown firms is associated with an approximate 5 tons per million dollar of revenue reduction in emissions. This *change* in emissions by brown firms due to a modest change in financial returns is equal to the average *level* of emissions for green firms. In contrast, neutral and green firms have close-to-zero and insignificant coefficients on the firm annual return. These results are consistent with brown firms having a large negative impact elasticity and green firms having a close-to-zero impact elasticity.

In Column 2, we find similar patterns using industry returns (calculated excluding the focal firm) instead of firm returns. These results using industry returns imply that our estimates are unlikely to be due to reverse causality or omitted variables. Rather, exogenous shocks to the firm performance, as proxied by the industry return, are associated with large declines in emissions by brown firms and small insignificant changes in emissions by green firms.

In Columns 3 and 4 of Table 3, we repeat the analysis using changes in real firm or industry performance, with real performance measured as return on assets. Column 3 shows a large negative and significant coefficient on ROA for brown firms of -133.1, implying that a 10% increase in firm ROA is associated with a decline in emissions by brown firms that is equal to 2.6 times the average level of emission by green firms. In contrast, neutral and brown firms display smaller positive changes in emissions following an increase in ROA. In Column 4, we find qualitatively similar patterns using changes in industry ROA instead of firm ROA, consistent with a causal link between real performance shocks and differential changes in emissions across brown and green firms.

In all specifications in Table 3, brown firms pollute less following positive performance shocks and pollute more following negative performance shocks. In contrast, green firms have smaller and inconsistently signed changes in emissions. We can also confidently reject the hypothesis that brown and green firms have equal elasticities (p-values for a test of equality in coefficients are below 0.01). In some specifications, green firms even exhibit the opposite relation; they increase emissions intensity following positive performance shocks. These empirical patterns are consistent with the notion that impact investment that rewards green firms by lowering their cost of capital and punishes brown firms by raising their cost of capital would have the counterproductive effect of increasing global emissions intensity.

Perhaps the most common refrain from impact investors on why they avoid brown firms is that they wish to starve those firms of capital, thereby forcing them into financial distress. If this results in brown firms becoming more green then the strategy would have the desired effect. Before getting to the empirical results, we note that there are reasons ex-ante to be skeptical of this line of argument. An increase in the firm's cost of capital should cause the firm to prefer investment projects that deliver front-loaded cash flows over those with back-loaded cash flows. In particular, a firm that is in a liquidity crisis or has a high risk of bankruptcy faces a high discount rate, such that the firm will favor investments offering short term gains. Since transitioning to greener production by brown firms usually entails adoption of new equipment and technologies that differ from their existing brown investment projects, these new green investments are unlikely to pay off in the very short run and should be less attractive to firms in financial distress.

In Table 4 we measure financial distress in four ways. First, we examine an indicator for for whether the firm is likely to face challenges in making interest payments on its existing debt. The indicator is equal to one if firms have positive interest payments and negative earnings, or the firm has an earnings to interest ratio that is in the bottom decile within our sample. Second, we measure each firm's Altman Z-score (a commonly-used predictor of the firm's probability of bankruptcy, see Altman (1983)), and set the low Z-score indicator equal to one if the firm has a Z-score in the bottom decile within our sample. Lastly, we use indicators for whether the firm's financial or real performance is in the bottom decile within our sample. To establish a causal channel, we also examine firm reactions to industry shocks, using indicators for whether industry performance (calculated excluding the focal firm) is in the bottom decile of our sample.

Across all six specifications in Table 4, we find that brown firms react to distress by increasing their emissions. In contrast, neutral and green firms exhibit smaller, inconsistently signed, and less significant responses to the proxies for distress. P-value tests of equality shows that we can reject the null that brown and green firms have equal changes in emissions after experiencing distress. The coefficient magnitudes imply that brown firms increase their emissions by approximately 20 to 50 tons per million in revenue after experiencing a distress shock associated with being in the lowest decile of some measure (interest coverage, Z-score, financial returns, or ROA) within our sample period. Given that the average level of emissions by green firms is only 5 tons per million in revenue, these results imply that brown firms react to distress by increase their emissions by at least four times the level of emissions of the average green firm.

So far, we have explored the relation between a firm's environmental impact and various proxies for the firm's cost of financing, such as bankruptcy risk and firm and industry performance. We can also attempt to directly measure the implied cost of capital (ICC) for each firm-year. Before proceeding, we stress that these tests should be viewed as a complement rather than as a replacement for our previous results because direct estimates of a firm's cost of capital have been shown to be noisy. Lack of unique numerical solutions along with small differences in the timing of measurement and assumptions regarding the path of future cash flows can lead to large differences in estimates of the ICC (for more details, see Lee et al. (2021)).

The ICC is the internal rate of return that equates the firm's market value to the present value of its expected future cash flows, based on a valuation model that can vary depending on the researcher's chosen set of assumptions. Thus, the ICC represents the expected return to investors of the firm and the firm's cost of raising capital from the same investors. We use estimates of firm ICCs generously shared by Lee et al. (2021). As our baseline, we follow the recommendations of Lee et al. (2021) and use ICCs estimated following the Gebhardt et al. (2001) (GLS) method where the inputs for future cash flows consist of mechanical forecasts from the cross-sectional forecast model of Hou et al. (2012). To ensure robustness, we also present results using a composite ICC that is the equal-weighted average for four ICC variants.

In Table 5, we regress the firm's change in emissions intensity on the firm's change in implied cost of capital over the previous year, interacted with indicators for whether the firm is brown, neutral, or green. We also control for the direct effects of the firm type indicators (brown, green, or neutral), fiscal year and SIC2 fixed effects, and in some specifications, we also control for SIC2 industry fixed effects.

We find that brown firms significantly increase their emissions following an increase in their cost of capital. Once again, neutral and brown firms experience smaller and less significant changes in their emissions following changes to their cost of capital. This is true both using the GLS as well as the composite ICC estimates. For example, the coefficient in Column 1 implies that brown firms increase emission by 68 tons per million follow a 10 percentage point increase in their ICC. Because an increase in emissions translates to a negative change in environmental impacts, these results again imply that brown firms have large negative impact elasticities with respect to their cost of financing, while neutral and green firms have smaller impact elasticities closer to zero.

Similar to our earlier analysis of firm performance, a limitation of looking at the relation between firm emissions and firm ICC is that any measured correlation could be driven by reverse causality (if becoming more green causes the firm to have a lower cost of capital), or by omitted variables bias (e.g., if the arrival a talented new green-oriented CEO causes both a shift in green strategy and cost of capital). To better estimate the causal impact of firm cost of capital on firm emissions, we examine the relation between firm emissions and the firm's industry value-weighted average ICC, calculated excluding the focal firm. The intuition is that industry cost of capital shocks strongly impact firm-level cost of capital, but individual firm choices should not have a strong impact on the industry average cost of capital calculated excluding the focal firm. We find a similar large elasticity of brown firm emissions intensity with respect to industry ICC, and an smaller and insignificant relation for neutral and green firms. In supplementary results, we find that these patterns are robust to using the equal-weighted average for four ICC variants.

C. Additional Incentive Effects of Impact Investing

So far, we have shown that brown firms have large negative impact elasticities and green firms have impact elasticities that are close to zero. Together, these elasticities imply, if dominant impact investing strategy succeeds in altering firms' cost of financing, it would have the *direct effect* of worsening aggregate firm environmental impact on society. In this section, we explore the possibility that the dominant impact investing strategy could have an additional indirect incentive effect on firm behavior. Specifically, if it is known that impact investors reward firms that improve their impact, then brown firms may be incentivized to become more green to access a lower cost of financing in the future.

We believe that developing an impact investing strategy that motivates brown firms to become more green is a promising agenda. However, we show empirically that the dominant impact investing strategy in practice has not yet provided any such incentives.

To study these indirect incentive effects, we examine the extent to which the dominant impact investing strategy over the past two decades has rewarded green and brown firms who have improved their environmental impact. Using data on the holdings of impact investment funds, we test whether impact investors increase their holdings of firms that have lowered their emissions, holding the current level constant. Using data on ESG ratings released by a leading impact investment advisory firm, we also test whether firms are rewarded for a decrease in emissions with improvements in their environmental ESG ratings.

Our analysis yields nuanced results which ultimately imply very weak incentives for brown firms to become more green. We find that impact investment funds indeed overweight firms that have improved their impact over the past several years, consistent with these funds rewarding firms who transition to becoming more green. However, changes in impact are measured in the wrong units. Impact investors appear to suffer from a proportional thinking bias (see e.g. Tversky and Kahneman (1981) and Shue and Townsend (2021)) in which they reward firms with large *percentage* reductions in emissions rather than large *level* reductions in emissions. Popular ESG environmental ratings similarly reward percentage reductions in emissions rather than level reductions in emissions. Because it is much more costly for a brown firm with high levels of pollution to have a large percentage reduction in emissions, the dominant impact investing strategy primarily rewards firms that are already green that have large percentage, but economically trivial, reductions in emissions.

Using data generously shared by Cohen et al. (2020), we classify funds as green if the fund name contains "ESG" or "Green" or if the fund is classified as an impact investment by either the Forum for Sustainable and Responsible Investment (USSIF) or Charles Schwab. To assess whether a firm is favored by green funds, we compare the holdings of two portfolios: the aggregated holdings of all green funds within a year and the holdings of a hypothetical market portfolio that holds all firms in CRSP in proportion to their market value as of the beginning of the year. We measure the extent to which a firm is rewarded by green funds using it's “overweight,” defined as the difference between the stock's portfolio weight in the aggregate green fund portfolio and the market portfolio, scaled by its weight in the market portfolio.

In Table 6, we regress the firm's overweight in the aggregate green fund portfolio on the firm's current level emissions as well as the firm's change in emission in the past one or two years. In Columns (1) and (2), we measure the firm's change in emissions in levels. We argue that this is the correct measure of the change in real firm environmental impact. Note that because we measure emissions as emissions intensity per dollar of revenue, measuring the change in emissions in levels is already adjusted for differences in firm size. In Columns (3) and (4), we measure change in emissions as the percentage change. We argue that this is the incorrect measure of the change in real firm environmental impact. As shown in Figure 3, green firms are associated with very large absolute percentage changes. These large percentage changes in emissions by green funds are economically trivial because their level of emissions is several orders of magnitude smaller than the level of emissions of similarly-sized brown firms.

The estimates in Table 6 show that the current level of emissions and percentage changes in emissions are very strong predictors of green fund holdings. Green investment funds reward firms with lower current levels of emissions as well as larger percentage reductions in their emissions. However, the estimated coefficients on the level changes in emissions in Columns (1) and (2) are close to zero and statistically insignificant. In other words, green investment funds, as a whole, fail to reward firms for reducing emissions in the units that actually matter for real environmental impact.

In Table 7, we find very similar results using the firm's ESG environmental rating as the dependent variable. We find that ESG ratings reward firms with lower current levels of emissions as well as larger percentage reductions in their emissions. However, the estimated coefficients on the level changes in emissions in Columns (1) and (2) are again close to zero and statistically insignificant.

The distinction between percentage and level changes in emissions is very important because, holding constant firm size, the average brown firm emits 260 times as much pollution as the average green firm. Comparing a brown and green firm of equal size, an increase in emissions by a brown firm of 1% has the same actual environmental impact as an increase in emissions by a green firm of 260%. It is also much easier for a green firm to purchase a small quantity of carbon offsets to completely offset its initially low level of emissions and become carbon neutral. However, this 100% reduction in emissions is far less economically meaningful than a brown firm reducing its emissions by a mere 1%.

Perhaps most surprisingly, we find in Table 8 that impact investment funds and ESG ratings reward green firms more than brown firms for the same percentage reduction in emissions. They should do the reverse and reward brown firms more for the same percentage reduction in emissions. This additional mistake is consistent with a behavioral affect heuristic (e.g., Slovic et al., 2007), in which impact investors choose to disassociate from or punish brown firms that they dislike, despite the fact that brown firms have the greatest scope to change their environmental impact.

V. Conclusion

This paper shows that the dominant impact investing strategy of directing capital toward green firms and away from brown firms can be counterproductive. We develop a new measure of impact elasticity, defined as a firm's change in environmental impact due to a change in its cost of financing. We show empirically that a reduction in financing costs for firms that are already green leads to small improvements in impact at best. Increasing financing costs for brown firms leads to large negative changes in firm impact. We further show that the dominant impact investing strategy provides very weak incentives for brown firms to become less brown. Instead, the impact investing movement primarily rewards green firms for economically trivial reductions in their already low levels of emissions.

Altogether, our results imply that the best case scenario for aggregate firm impact is the one where the dominant impact investing strategy has not yet shifted firms' cost of financing. To the extent that such a strategy succeeds in changing firms' cost of financing, it would have the counterproductive effect of lowering the overall green impact of firms on society.

Our findings and conclusions are not meant as a negative assessment of all possible impact investment strategies. Rather, they highlight potential problems with the most popular impact investment strategy to date, which divests from brown firms and invests in green firms. Our analysis suggests that impact investing flows and engagement that targets the incentives of green firms would be more effective if targeted at brown firms.

References

Altman, Edward I, "A complete guide to predicting, avoiding, and dealing with bankruptcy," *Corporate Financial Distress, New York*, 1983.

Berg, Florian, Julian F Koelbel, and Roberto Rigobon, "Aggregate confusion: The divergence of ESG ratings," *Forthcoming Review of Finance*, 2019.

Berk, Jonathan and Jules H van Binsbergen, "The impact of impact investing," *Available at SSRN 3909166*, 2021.

Broccardo, Eleonora, Oliver Hart, and Luigi Zingales, "Exit vs. Voice," 2020.

Chava, Sudheer, "Environmental externalities and cost of capital," *Management science*, 2014, *60* (9), 2223–2247.

Cohen, Lauren, Umit G Gurun, and Quoc H Nguyen, "The ESG-innovation disconnect: Evidence from green patenting," Technical Report, National Bureau of Economic Research 2020.

Edmans, Alex, Doron Levit, and Jan Schneemeier, "Socially responsible divestment," *European Corporate Governance Institute–Finance Working Paper*, 2022, (823).

Eisfeldt, Andrea L and Adriano A Rampini, "New or used? Investment with credit constraints," *Journal of Monetary Economics*, 2007, *54* (8), 2656–2681.

Gebhardt, William R, Charles MC Lee, and Bhaskaran Swaminathan, "Toward an implied cost of capital," *Journal of accounting research*, 2001, *39* (1), 135–176.

Gilje, Erik P, Elena Loutskina, and Daniel Murphy, "Drilling and debt," *The Journal of Finance*, 2020, *75* (3), 1287–1325.

Hartzmark, Samuel M and Abigail B Sussman, "Do investors value sustainability? A natural experiment examining ranking and fund flows," *The Journal of Finance*, 2019, *74* (6), 2789–2837.

Heath, Davidson, Daniele Macciocchi, Roni Michaely, and Matthew C Ringgenberg, "Does Socially Responsible Investing Change Firm Behavior?," *European Corporate Governance Institute–Finance Working Paper*, 2021, (762).

Hong, Harrison, Jeffrey D Kubik, and Jose A Scheinkman, "Financial constraints on corporate goodness," Technical Report, National Bureau of Economic Research 2012.

Hou, Kewei, Mathijs A Van Dijk, and Yinglei Zhang, "The implied cost of capital: A new approach," *Journal of Accounting and Economics*, 2012, *53* (3), 504–526.

Lee, Charles MC, Eric C So, and Charles CY Wang, "Evaluating firm-level expected-return proxies: implications for estimating treatment effects," *The Review of Financial Studies*, 2021, *34* (4), 1907–1951.

Pástor, L'uboš, Robert F Stambaugh, and Lucian A Taylor, "Sustainable investing in equilibrium," *Journal of Financial Economics*, 2021, *142* (2), 550–571.

Shue, Kelly and Richard R Townsend, "Can the Market Multiply and Divide? Non-Proportional Thinking in Financial Markets," *The Journal of Finance*, 2021, *76* (5), 2307–2357.

Slovic, Paul, Melissa L Finucane, Ellen Peters, and Donald G MacGregor, "The affect heuristic," *European journal of operational research*, 2007, *177* (3), 1333–1352.

Thomas, Jake, Wentao Yao, Frank Zhang, and Wei Zhu, "Meet, beat, and pollute," *Review of Accounting Studies*, 2022, *27* (3), 1038–1078.

Tversky, Amos and Daniel Kahneman, "The framing of decisions and the psychology of choice," *science*, 1981, *211* (4481), 453–458.

Figure 1: Average carbon emission in each quintile

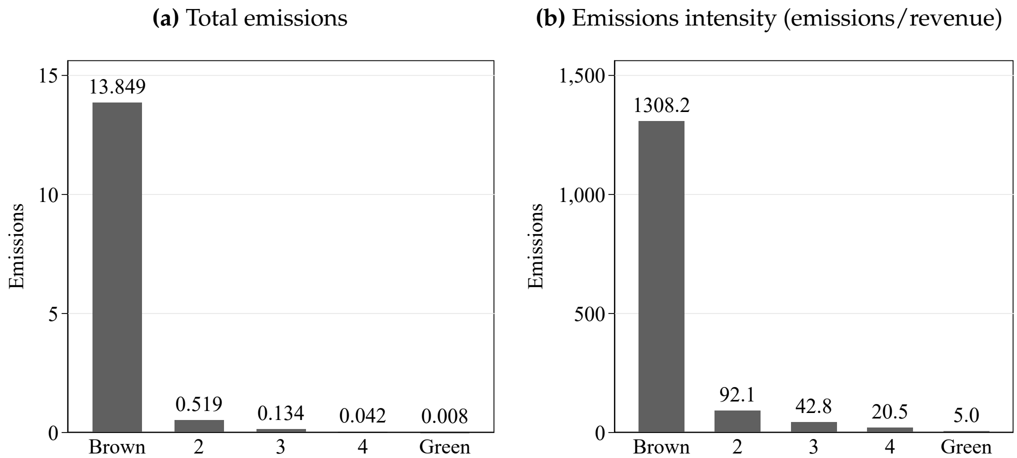

This figure plots the average emissions of scope 1 and scope 2 green house gases by firms, sorted into quintiles within each year, with quintile 1 representing brown firms and quintile 5 representing green firms. In Panel (a), emissions are measured as million tons of CO_2 equivalents. In Panel (b), emissions are measured as tons of CO_2 equivalents emitted per million dollars of revenue.

Figure 2: Absolute variation of emissions intensity by quintile

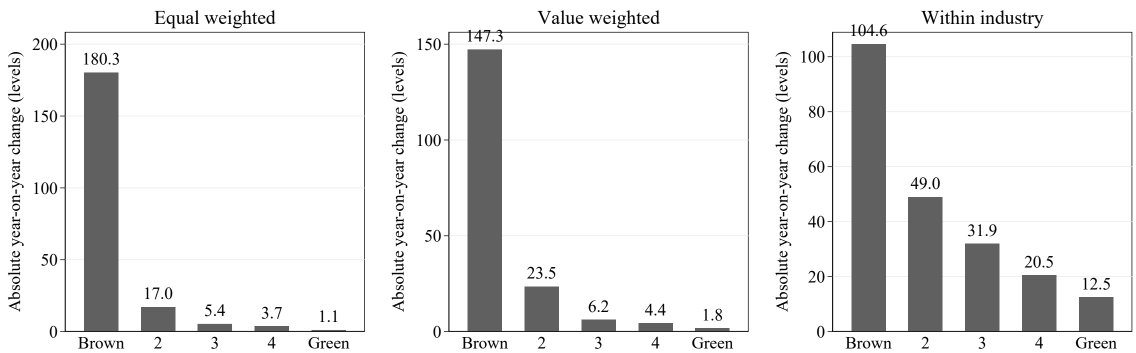

This figure plots year-on-year variation in emissions across quintiles for the level of emissions. Variation in emissions is $|e_{t+1} - e_t|$, the absolute change in scope 1 and scope 2 greenhouse gases emissions intensity, where emissions intensity is measured in tons of CO_2 equivalents emitted per million dollars of revenue. In the left and middle panels, quintiles are computed within each fiscal year. In the right panel, quintiles are computed within each year×SIC2 industry. Observations in the left and right panels are equal weighted. In the middle panel, observations are weighted by the firm's market capitalization as a fraction of total CRSP market size in the same year.

Figure 3: Absolute percentage variation of emissions intensity by quintile

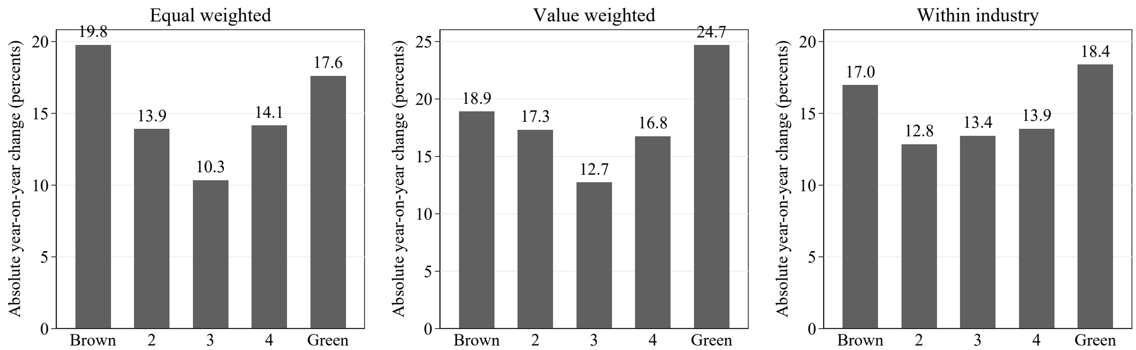

This figure plots year-on-year percentage variation in emissions across quintiles for the level of emissions. Variation in emissions is $|\frac{e_{t+1}-e_t}{e_t} \times 100|$, the absolute percentage change in scope 1 and scope 2 greenhouse gases emissions intensity, where emissions intensity is measured in tons of CO_2 equivalents emitted per million dollars of revenue. In the left and middle panels, quintiles are computed within each fiscal year. In the right panel, quintiles are computed within each year$\times$SIC2 industry. Observations in the left and right panels are equal weighted. In the middle panel, observations are weighted by the firm's market capitalization as a fraction of total CRSP market size in the same year.

Table 1: Summary statistics

	Mean	SD	p10	p50	p90
Total emissions	2.9063	14.2767	0.0019	0.0938	3.7277
Emissions intensity (emissions/revenue)	257.9102	733.6114	3.4575	40.6942	534.5153
Absolute changes in emissions	41.5508	138.0794	0.1899	2.2738	83.8643
Absolute percentage changes in emissions	0.1960	1.0431	0.0130	0.0634	0.3352
Changes in emissions	-5.3007	108.4962	-25.1790	-0.4828	12.9887
Annual return	0.1343	0.5298	-0.3835	0.0676	0.6227
Industry annual return	0.1452	0.2213	-0.1106	0.1370	0.4000
ΔROA	-0.0022	0.0896	-0.0747	-0.0000	0.0649
ΔIndustry ROA	-0.0008	0.0238	-0.0230	0.0010	0.0225
ΔICC	0.0008	0.0265	-0.0265	-0.0000	0.0300
ΔIndustry ICC	-0.0002	0.0132	-0.0127	-0.0007	0.0151
ΔICC composite	0.0085	0.0631	-0.0527	0.0022	0.0795
ΔIndustry ICC composite	0.0030	0.0346	-0.0285	-0.0003	0.0407

This table presents summary statistics for our main analysis sample, consisting of observations at the firm-year level. Total emissions is measured as million tons of CO_2 equivalents. Emissions intensity is total emissions per million dollars of revenue. Hereafter, we refer to emissions intensity as emissions for brevity. Absolute change in emissions is the absolute value of the annual change in the level of emissions. Absolute percentage change in emissions is the absolute value of the annual fractional change in emissions. Annual return is the annual return of the firm. Industry annual return is the annual value-weighted return within each SIC2 industry, calculated excluding the focal firm. ΔROA is the annual change in firm ROA. ΔIndustry ROA is the annual value-weighted change in industry ROA, calculated excluding the focal firm. ΔICC is the annual change in the firm implied cost of capital estimated using the mechanical GLS method, as described in Lee et al. (2021). ΔIndustry ICC is the annual value-weighted change in industry ICC, calculated excluding the focal firm. Δ ICC composite is the annual change in the firm implied cost of capital estimated using the composite method, as described in Lee et al. (2021). ΔIndustry ICC composite is the annual value-weighted change in industry ICC composite, calculated excluding the focal firm.

Table 2: Absolute change in emissions intensity by quintile

	Absolute changes in emissions		
	(1)	(2)	(3)
Quintile 2	-163.3***	-124.1***	-55.39***
	(8.448)	(13.11)	(5.934)
Quintile 3	-175.1***	-140.8***	-72.06***
	(8.409)	(12.01)	(5.988)
Quintile 4	-176.6***	-142.8***	-86.44***
	(8.393)	(12.04)	(6.258)
Quintile 5	-179.2***	-146.0***	-92.64***
	(8.390)	(12.09)	(6.401)
Year FE	Yes	Yes	No
SIC2 industry FE	No	No	Yes
Value-weighted	No	Yes	No
Within SIC2 industry	No	No	Yes
N	24345	24330	24280
R^2	0.262	0.259	0.372

This table shows year-on-year variation in emissions by quintiles representing the level of emissions. The dependent variable is $|e_{t+1} - e_t|$, the absolute change in scope 1 and scope 2 greenhouse gases emissions, where emissions is measured in tons of CO_2 equivalents emitted per million dollars of revenue. We regress the dependent variable on quintile dummies of emissions intensity in year t. Estimated coefficients represent the difference in absolute year-on-year variation in emissions relative to the omitted category, quintile 1 representing brown firms. Standard errors are clustered at the firm-level. In columns (1) and (2), quintiles are computed within each fiscal year, while in column (3), quintiles are computed within each fiscal year$\times$SIC2 industry. Columns (1) and (2) include year fixed effects, and column (3) includes year and SIC2 industry fixed effects. In column (2) regressions are value-weighted where each observation is weighted by the firm's market capitalization as a fraction of total CRSP market size in the same year. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Table 3: Change in emissions and changes in performance

	Changes in emissions			
	(1)	(2)	(3)	(4)
Brown × Annual return	-49.95***			
	(8.150)			
Neutral × Annual return	1.260			
	(0.944)			
Green × Annual return	1.245			
	(1.101)			
Brown × Industry annual return		-72.84***		
		(16.96)		
Neutral × Industry annual return		-0.315		
		(5.014)		
Green × Industry annual return		1.445		
		(5.487)		
Brown × ΔROA			-133.1***	
			(37.82)	
Neutral × ΔROA			9.273***	
			(3.112)	
Green × ΔROA			7.486	
			(5.461)	
Brown × ΔIndustry ROA				-222.2**
				(104.7)
Neutral × ΔIndustry ROA				43.49*
				(24.51)
Green × ΔIndustry ROA				140.8***
				(40.65)
p-value: Brown × X = Green × X	0.000	0.000	0.000	0.001
Type FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
SIC2 industry FE	Yes	Yes	Yes	Yes
N	23818	24271	21863	24262
R^2	0.0493	0.0432	0.0421	0.0392

This table shows changes in firms' emissions following changes in firm or industry performance. The dependent variable is $e_{t+1} - e_t$, the change in scope 1 and scope 2 greenhouse gases emissions, where emissions is measured in tons of CO_2 equivalents emitted per million dollars of revenue. We regress the dependent variable on the interactions between firm- or industry-level financial or real performance in the previous year and indicators for whether the firm is brown, neutral, or green. All other variables are as define in Table 1. All columns include year fixed effects, SIC2 industry fixed effects, and indicators for whether the firm is brown, neutral, or green. Below each regression, we report the p-value for the test of whether coefficients on the interaction with brown firms and the interaction with green firms are equal. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Table 4: Change in emissions and financial distress

	Changes in emissions					
	(1)	(2)	(3)	(4)	(5)	(6)
Brown × Low interest coverage	26.15*					
	(14.80)					
Neutral × Low interest coverage	-6.135***					
	(1.189)					
Green × Low interest coverage	-4.218***					
	(1.289)					
Brown × Low Z-score		34.90***				
		(12.96)				
Neutral × Low Z-score		-5.333***				
		(1.273)				
Green × Low Z-score		0.508				
		(2.117)				
Brown × Low annual return			53.74***			
			(10.88)			
Neutral × Low annual return			-4.854***			
			(1.365)			
Green × Low annual return			-4.225**			
			(1.733)			
Brown × Low industry annual return				51.19***		
				(11.17)		
Neutral × Low industry annual return				-1.077		
				(3.303)		
Green × Low industry annual return				0.202		
				(3.001)		
Brown × Low ΔROA					29.80**	
					(11.80)	
Neutral × Low ΔROA					-4.523***	
					(1.085)	
Green × Low ΔROA					-2.003	
					(1.390)	
Brown × Low ΔIndustry ROA						20.27**
						(8.736)
Neutral × Low ΔIndustry ROA						-2.474
						(1.618)
Green × Low ΔIndustry ROA						-9.815***
						(1.732)
p-value: Brown × X = Green × X	0.041	0.009	0.000	0.000	0.007	0.001
Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
SIC2 industry FE	Yes	Yes	Yes	Yes	Yes	Yes
N	19747	19069	23818	24271	21863	24262
R^2	0.0404	0.0425	0.0440	0.0433	0.0413	0.0392

This table shows changes in firms' emissions following financial distress. The dependent variable is $e_{t+1} - e_t$, the change in scope 1 and scope 2 greenhouse gases emissions, where emissions is measured in tons of CO_2 equivalents emitted per million dollars of revenue. We regress the dependent variable on the interactions between indicators for financial distress in the previous year and indicators for whether the firm is brown (quintile 1 in terms of emissions in year t), neutral (quintiles 2-4), or green (quintile 5). The low interest coverage indicator is equal to one if the firm has positive interest payments and negative earnings, or the firm has an earnings to interest ratio that is in the bottom decile within our sample. The low Z-score indicator is equal to one if the firm has an Altman Z-score in the bottom decile within our sample. All other variables are as defined in Table 1. All columns include year fixed effects and indicators for whether the firm is brown, neutral, or green, and columns (2) and (4) further include SIC2 industry fixed effects. Below each regression, we report the p-value for the test of whether coefficients on the interaction with brown firms and the interaction with green firms are equal. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Table 5: Change in emissions and changes in ICC

	Changes in emissions			
	(1)	(2)	(3)	(4)
Brown × ΔICC	680.6***			
	(140.8)			
Neutral × ΔICC	-14.36			
	(19.15)			
Green × ΔICC	-18.40			
	(16.90)			
Brown × ΔIndustry ICC		452.8*		
		(240.9)		
Neutral × ΔIndustry ICC		-26.56		
		(62.93)		
Green × ΔIndustry ICC		-44.19		
		(57.17)		
Brown × ΔICC composite			418.3***	
			(107.3)	
Neutral × ΔICC composite			-10.35	
			(18.38)	
Green × ΔICC composite			-22.82	
			(16.78)	
Brown × ΔIndustry ICC composite				402.1***
				(122.8)
Neutral × ΔIndustry ICC composite				19.59
				(23.34)
Green × ΔIndustry ICC composite				38.02*
				(22.98)
p-value: Brown × X = Green × X	0.000	0.042	0.000	0.003
Type FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
SIC2 industry FE	Yes	Yes	Yes	Yes
N	15978	24162	6801	22809
R^2	0.0521	0.0389	0.0639	0.0390

This table shows changes in firms' emissions following changes in firm or industry implied cost of capital (ICC). Measures of the ICC are as defined in Table 1. The dependent variable is $e_{t+1} - e_t$, the change in scope 1 and scope 2 greenhouse gases emissions intensity, where emissions intensity is measured in tons of CO_2 equivalents emitted per million dollars of revenue. All columns include year fixed effects, SIC2 industry fixed effects, and indicators for whether the firm is brown, neutral, or green. Below each regression, we report the p-value for the test of whether coefficients on the interaction with brown firms and the interaction with green firms are equal. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Table 6: Portfolio holdings of green funds and changes in emissions

	Overweight in green funds			
	(1)	(2)	(3)	(4)
Emissions	-0.00726***	-0.00808***	-0.00697***	-0.00764***
	(0.00237)	(0.00253)	(0.00233)	(0.00248)
$\Delta_{t,t-1}$Emissions (changes in levels)	-0.00196			
	(0.00554)			
$\Delta_{t,t-2}$Emissions (changes in levels)		0.000651		
		(0.00455)		
$\Delta_{t,t-1}$Emissions (changes in percents)			-0.106***	
			(0.0385)	
$\Delta_{t,t-2}$Emissions (changes in percents)				-0.0584**
				(0.0279)
Year FE	Yes	Yes	Yes	Yes
N	24345	21118	24345	21118
R^2	0.0106	0.0113	0.0108	0.0114

This table shows how the relation between the holdings of green investment funds and emissions varies depending on how changes in emissions are measured. The dependent variable measures the extent to which a firm is overweighted by green funds relative to the stock's weight in a value-weighted market portfolio (overweight is calculated as total dollar value of holdings in the relevant stock as a fraction of the total dollar value of all green funds, minus the value of the stock as a fraction of total market value). All columns control for the level of emissions. Columns (1) and (2) control for the one- and two-year change in the level of emissions, respectively. Columns (3) and (4) control for the one- and two-year percentage change in emissions, respectively. All columns include year fixed effects. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Table 7: Environmental score and changes in emissions

	Environmental score			
	(1)	(2)	(3)	(4)
Emissions	-0.0190***	-0.0197***	-0.0188***	-0.0196***
	(0.00346)	(0.00361)	(0.00348)	(0.00364)
$\Delta_{t,t-1}$Emissions (changes in levels)	-0.00514			
	(0.00748)			
$\Delta_{t,t-2}$Emissions (changes in levels)		0.00172		
		(0.00756)		
$\Delta_{t,t-1}$Emissions (changes in percents)			-0.130***	
			(0.0347)	
$\Delta_{t,t-2}$Emissions (changes in percents)				-0.0874***
				(0.0256)
Year FE	Yes	Yes	Yes	Yes
N	9887	8568	9887	8568
R^2	0.155	0.167	0.156	0.169

This table shows how the relation between ESG environmental ratings and emissions varies depending on how changes in emissions are measured. The dependent variable is the MSCI ESG environmental score. All columns control for the level of emissions. Columns (1) and (2) control for the one- and two-year change in the level of emissions, respectively. Columns (3) and (4) control for the one- and two-year percentage change in emissions, respectively. All columns include year fixed effects. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Table 8: Impact investing and changes in emissions by firm type

	Overweight in green funds		Environmental score	
	(1)	(2)	(3)	(4)
Emissions	-0.00633**	-0.00608**	-0.0161***	-0.0155***
	(0.00308)	(0.00303)	(0.00444)	(0.00464)
Brown × $\Delta_{t,t-1}$Emissions (changes in percents)	-0.0142		-0.133*	
	(0.0734)		(0.0698)	
Neutral × $\Delta_{t,t-1}$Emissions (changes in percents)	-0.161***		-0.121**	
	(0.0553)		(0.0538)	
Green × $\Delta_{t,t-1}$Emissions (changes in percents)	-0.199***		-0.146**	
	(0.0574)		(0.0589)	
Brown × $\Delta_{t,t-2}$Emissions (changes in percents)		0.00814		-0.0836
		(0.0473)		(0.0526)
Neutral × $\Delta_{t,t-2}$Emissions (changes in percents)		-0.0908**		-0.0673*
		(0.0385)		(0.0364)
Green × $\Delta_{t,t-2}$Emissions (changes in percents)		-0.169**		-0.147***
		(0.0704)		(0.0489)
Type FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	24345	21118	9887	8568
R^2	0.0110	0.0119	0.159	0.173

This table shows how impact investing differentially reacts to percentage changes in firm emissions depending on whether the firm is brown, netural, or green. The dependent variable in Columns (1) and (2) is the stock's overweight in green investment funds as defined in Table 6. The dependent variable in Columns (3) and (4) is the stock's MSCI KLD environmental rating as defined in Table 8. All columns include year fixed effects. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Online Appendix for

Counterproductive Impact Investing

Samuel M. Hartzmark   Kelly Shue

Table A1: Absolute change in emissions scope 1 intensity by quintile

	Absolute changes in emissions		
	(1)	(2)	(3)
Quintile 2	-148.8***	-113.9***	-45.30***
	(7.901)	(11.62)	(5.589)
Quintile 3	-156.9***	-126.1***	-61.62***
	(7.880)	(10.84)	(5.633)
Quintile 4	-158.1***	-126.6***	-72.22***
	(7.858)	(10.75)	(5.802)
Quintile 5	-159.1***	-127.9***	-79.87***
	(7.870)	(10.78)	(6.044)
Year FE	Yes	Yes	No
Year $\times$ SIC2 FE	No	No	Yes
Value-weighted	No	Yes	No
Within SIC2 industry	No	No	Yes
N	24345	24330	24280
R^2	0.259	0.253	0.385

This table shows year-on-year variation in emissions by quintiles representing the level of emissions. The dependent variable is $|e_{t+1} - e_t|$, the absolute change in scope 1 greenhouse gases emissions, where emissions is measured in tons of CO_2 equivalents emitted per million dollars of revenue. We regress the dependent variable on quintile dummies of emissions intensity in year t. Estimated coefficients represent the difference in absolute year-on-year variation in emissions relative to the omitted category, quintile 1 representing brown firms. Standard errors are clustered at the firm-level. In columns (1) and (2), quintiles are computed within each fiscal year, while in column (3), quintiles are computed within each fiscal year$\times$SIC2 industry. Columns (1) and (2) include year fixed effects, and column (3) includes year and SIC2 industry fixed effects. In column (2) regressions are value-weighted where each observation is weighted by the firm's market capitalization as a fraction of total CRSP market size in the same year. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.

Table A2: Absolute change in emissions scope 2 intensity by quintile

	Absolute changes in emissions		
	(1)	(2)	(3)
Quintile 2	-17.13***	-23.23***	-11.89***
	(1.154)	(3.464)	(0.991)
Quintile 3	-20.05***	-25.71***	-13.83***
	(1.159)	(3.452)	(0.999)
Quintile 4	-20.88***	-26.35***	-14.56***
	(1.164)	(3.522)	(1.005)
Quintile 5	-21.42***	-27.27***	-15.86***
	(1.150)	(3.457)	(1.021)
Year FE	Yes	Yes	No
Year × SIC2 FE	No	No	Yes
Value-weighted	No	Yes	No
Within SIC2 industry	No	No	Yes
N	24345	24330	24280
R^2	0.156	0.213	0.253

This table shows year-on-year variation in emissions by quintiles representing the level of emissions. The dependent variable is $|e_{t+1} - e_t|$, the absolute change in scope 2 greenhouse gases emissions, where emissions is measured in tons of CO_2 equivalents emitted per million dollars of revenue. We regress the dependent variable on quintile dummies of emissions intensity in year t. Estimated coefficients represent the difference in absolute year-on-year variation in emissions relative to the omitted category, quintile 1 representing brown firms. Standard errors are clustered at the firm-level. In columns (1) and (2), quintiles are computed within each fiscal year, while in column (3), quintiles are computed within each fiscal year×SIC2 industry. Columns (1) and (2) include year fixed effects, and column (3) includes year and SIC2 industry fixed effects. In column (2) regressions are value-weighted where each observation is weighted by the firm's market capitalization as a fraction of total CRSP market size in the same year. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses and are clustered at the firm-level.