Carbon Burden

Lubos Pastor
Robert F. Stambaugh
Lucian A. Taylor*

January 19, 2026
First draft: October 24, 2024

Abstract

We quantify the U.S. corporate sector's future carbon damages by computing its "carbon burden"—the present value of social costs of its future carbon emissions. Our baseline estimate of the carbon burden is 131% of total corporate equity value. Even with indirect emissions excluded, 13% of firms have carbon burdens exceeding their market capitalizations. The 30 largest emitters account for all the decarbonization of U.S. corporations predicted by 2050. Predicted emission reductions, and even firms' targets, fall short of the Paris Agreement. Carbon burden is priced: firms with higher burdens have higher costs of capital, even controlling for past emissions.

JEL classifications: D62, G30, G38, Q51, Q54
Keywords: Externality, carbon emissions, carbon burden

---

*Pastor is at the University of Chicago Booth School of Business, NBER, and CEPR. Stambaugh and Taylor are at the Wharton School of the University of Pennsylvania and the NBER. Pastor also serves as an independent director and trustee of Vanguard. The views expressed here do not necessarily reflect those of Vanguard or its funds. We are grateful for comments from our discussants Antonio Coppola, Matthew Gustafson, Yrjo Koskinen, Tuomas Tomunen, Shaojun Zhang; from Jaroslav Borovička, Frank Diebold, Gene Fama, Xavier Gabaix, Michael Gerrard, Marcin Kacperczyk, Christian Leuz, Eric Orts, Michael Roberts, Johannes Stroebel, Jeff Wurgler, Yao Zeng, Luigi Zingales; from conference participants at the 2026 AFA, 2025 WFA, 2025 Booth Global Economy and Financial Stability Conference, 2025 Kellogg Innovations in Sustainable Finance Conference, 2025 Stanford Sustainability Conference on Capital Finance, 2025 Lisbon Sustainability Week, 2025 Sir Oliver Hart Lecture at the LSE; and seminar participants at Chicago, Colorado, Harvard, Notre Dame, NYU, Ohio State, Wharton, and Wisconsin. We thank Vitor Furtado Farias, Vanessa Hu, Aria Liu, and Chris Yaoyao Zhu for excellent research assistance. We also thank the MSCI Sustainability Institute and the Rustandy Center for Social Sector Innovation for providing data. This research was supported by the Fama-Miller Center for Research in Finance and the Robert King Steel Fellowship at the University of Chicago Booth School of Business.

1. Introduction

How valuable are firms to society? Firms create value not only for shareholders but also for consumers, employees, and other stakeholders. Importantly, a firm's value to society includes any externalities produced by the firm. These can be positive, such as technological spillovers from R&D investment, or negative, such as environmental damage.

How big are corporate externalities? The magnitude of an externality can be helpful information to many. Policymakers can use it to design more effective regulations, taxes, or subsidies. Companies can use it in their sustainability efforts and risk management practices. Knowing the scale of corporate externalities can also influence consumer behavior and help investors make more informed investment decisions. From the academic perspective, the size of corporate externalities speaks to the debate about the famous doctrine of Friedman (1970). Friedman's position that companies should essentially just maximize market value becomes controversial in the presence of externalities (e.g., Hart and Zingales, 2017). Maximizing market value can then conflict with maximizing the welfare of shareholders who also have social and ethical concerns. This conflict is particularly strong when the externalities' social costs or benefits are large relative to a firm's market value.

In this paper, we explore the size of one externality: damages from corporate emissions of greenhouse gases. This "carbon externality" is clearly important given the severity of the climate crisis. Key to measuring this externality is recognizing its future dimensions. First, emissions in any given period have climate consequences for many years. Second, emissions are expected to remain high for many years, and the future path of emissions will be crucial in determining climate change. Our contribution is to quantify the economic value of damages produced by future emissions.

To value these damages, we propose a metric that we call "carbon burden." We define a firm's carbon burden as the present value of the social costs associated with its future greenhouse gas (GHG) emissions, which we refer to simply as "carbon emissions" or just "emissions." Key to the carbon burden is the social cost of carbon (SCC), the dollar cost of societal damages resulting from the emission of one additional ton of carbon into the atmosphere. For a ton emitted τ years from now, let SCC_τ denote the net present value, as of that emission year, of the resulting damages in that year and all subsequent years. Let C_τ denote a firm's expected carbon emissions τ years from now. We define the firm's carbon burden as

$$\text{Carbon burden} = \sum_{\tau=1}^{T} (1+\rho_\tau)^{-\tau} \times C_\tau \times SCC_\tau \,, \tag{1}$$

where ρ_τ is a discount rate that potentially includes a risk premium. We set $\rho_\tau = \rho$ and consider a range of values for ρ. For SCC_τ, we use estimates recently released by the U.S. Environmental Protection Agency (EPA). We use emission forecasts, C_τ, at both the aggregate level and the firm level. Our forecasts of aggregate U.S. carbon emissions come from U.S. government agencies. Our firm-level emission forecasts come from MSCI, a leading data provider. All of these forecasts are undoubtedly imprecise, but they come without confidence intervals, precluding us from quantifying the precision of the carbon burden estimates. The estimates must therefore be interpreted with caution.

We focus primary attention on the carbon burden imposed by emissions in all future years (i.e., $T = \infty$), but we consider finite horizons as well. With an infinite horizon, the concept of carbon burden is similar in spirit to that of market value, in that both are present values of infinite streams of estimated future dollar values. For example, the market value of a firm's equity is the present value of its future dividends, whereas a firm's carbon burden is the present value of the social costs from the firm's future emissions. The two concepts measure different dimensions of a firm's value to society, with market value belonging to shareholders and carbon burden representing a negative value borne by all. Both market value and carbon burden are measured in dollars, and we compare them in our analysis.

We equate aggregate corporate emissions with total U.S. emissions, because virtually all emissions are related to the emissions of some company, directly or indirectly.[1] Of course, responsibility for corporate emissions does not rest solely with corporations. Households, for example, surely share this responsibility, but quantifying the corporate externality in a manner that accounts for responsibility seems infeasible.

At the aggregate level, we analyze the total U.S. carbon burden as of year-end 2023. Applying our baseline discount rate of $\rho = 2\%$ to emission forecasts for all future years, we estimate the U.S. carbon burden to be \$87 trillion, which is 131% of the total value of U.S. corporate equity. The burden is large also when computed using other discount rates and when compared to the total value of U.S. corporate equity and debt (93%) and to U.S. national wealth (61%). While carbon damages are clearly large, their negative value is more than offset by the positive value of the corporate sector's consumer surplus, estimated by Pellegrino (2025). Both of these components of firms' value to society are large when compared to firms' value just to shareholders.

After quantifying the aggregate U.S. carbon burden, we analyze its potential reductions under the 2015 Paris Agreement, in which U.S. participation has been sporadic. If the U.S.

[1]For simplicity, we use "corporate emissions" to refer to the emissions of any company or business, not just companies that are incorporated.

were to participate and fully comply, U.S. emissions would fall at least 50% by 2030, relative to the 2005 level. Again applying the 2% discount rate to all future years, we find that meeting those goals would reduce the U.S. carbon burden substantially, by either 21% or 32%, depending on the projected emission path beyond 2030. We also show that achieving the Paris goals would require major emission reductions by the largest emitters. However, the largest emitters' targeted emission reductions fall well below the Paris goals, even if we take those targets at face value. When we replace firms' targets by emission forecasts from MSCI, the shortfall relative to Paris widens further.

We find high dispersion across firms in the ratios of carbon burden to market value. These ratios are smaller than 0.05 for the majority (55%) of firms. However, for 13% of firms, which represent 10% of total market capitalization, these ratios are greater than one, meaning those firms' carbon burdens exceed their market capitalization. These estimates are based on direct (scope 1) emissions, which are emissions from sources owned by the firm. We also consider indirect emissions from the consumption of purchased energy (scope 2) and indirect emissions incurred in the firm's entire value chain (scope 3).[2] Based on total (scope 1+2+3) emissions, 77% of firms, representing 50% of total market capitalization, have carbon burdens exceeding their market capitalization.

The ratios of carbon burden to market value also differ greatly across sectors. Based on direct emissions, these ratios are as high as 7 and 3 for typical firms in the utilities and energy sectors, respectively, and as low as 0.01 for a typical financial firm. When we add all indirect emissions, the ratio of carbon burden to market value grows to 66 for a typical energy firm, and there are four other sectors in which this ratio exceeds 10. One of these is financials—a typical financial firm's ratio of 17 for total emissions stands in stark contrast to its aforementioned 0.01 ratio for direct emissions.

We also examine the ratio of a firm's carbon burden from all future years' emissions to its burden from a single year's emissions in 2023. This ratio varies substantially across firms, as a result of a large dispersion in MSCI's forecasts of future emission growth rates, which range from -100% to $+33\%$ when cumulated between 2023 and 2050. Given this large dispersion, it is not sufficient to look at firms' recent emissions when judging carbon damages. For example, suppose two firms had the same emissions recently, but the first firm has a credible decarbonization plan whereas the second firm does not. The first firm's carbon damages are then lower. If firms' carbon burdens were widely reported, they could incentivize firms to develop credible emission reduction strategies.

[2]These scope definitions come from the Greenhouse Gas Protocol, https://ghgprotocol.org. Among the three measures, scope 3 emissions are generally the hardest to quantify and least likely to be reported.

Expected future emission growth rates differ significantly across sectors. For example, the cumulative direct emission growth from 2023 to 2050 is −37% for a typical utility but −6% for a typical financial firm, based on MSCI's forecasts. The corresponding growth rates based on firms' own reported emission targets are much more negative, ranging from −47% for a typical nondurables firm to −92% for a typical utility. The former growth rates are less negative because MSCI views firms' own emission reduction targets as too optimistic.

We find a negative cross-sectional relation between firms' recent emissions and forecasted future emission growth rates. For example, for the top 5% of emitters, their forecasted cumulative growth rate of direct emissions from 2023 to 2050 is −14%, but for the bottom 5% of emitters, it is +25%. This negative relation is so strong that the 30 largest emitters are expected to account for the entire drop in aggregate direct U.S. corporate emissions by year 2050. Between 2023 and 2050, the aggregate emissions are expected to decline from 2.0 billion to 1.5 billion metric tons. Over the same period, the emissions of the 30 largest emitters are also expected to decline by 0.5 billion tons, whereas the emissions of the remaining 2,411 firms in our sample are expected to change little. Strikingly, all of the decarbonization of the U.S. corporate sector, as measured by direct emissions, is expected to come from only 30 firms.

Besides recent emissions, a few other firm characteristics, namely investment, climate score, and the book-to-market ratio, help explain the cross section of forecasted emission growth rates. Emissions are expected to grow faster for firms that invest more, firms with lower climate scores, and value firms, though these relations are not always significant.

Future emissions could be priced in firms' current market values. For example, expected cash flows could be reduced by potential carbon taxes or tort awards for the emissions' damages, and discount rates could be affected by carbon-related systematic risk. As of this writing, there is no nationwide carbon tax in the U.S., but some states and municipalities have levied taxes, and some have filed tort suits against energy companies. Future emissions that are already priced in market values cannot be termed an externality. Carbon burden measures the externality gross of those pricing effects. While recognizing this distinction, we sometimes refer to carbon burden as an externality, for simplicity.

We find that firms with higher carbon burdens do have higher discount rates. Specifically, firms with higher ratios of carbon burden to market capitalization have higher expected stock returns, as proxied by the implied cost of capital (ICC). Moreover, while the ICC is also higher for firms with higher past emissions relative to market cap, we find this relation becomes insignificant, and even flips sign, in the presence of carbon burden. That is, expected

return relates positively to future carbon rather than past carbon. This evidence supports the notion that credible decarbonization plans are associated with lower costs of capital. A potential explanation is that firms with higher carbon burdens are more exposed to the risk of future carbon taxes or related policies, and this risk carries a positive premium.

We are not the first to relate the cross section of stock returns to carbon emissions (e.g., Bolton and Kacperczyk, 2021, 2023, Aswani, Raghunandan, and Rajgopal, 2024, Zhang, 2024, Eskildsen et al., 2025), but we are the first, to our knowledge, to relate it to forecasts of future emissions. The literature also examines the carbon exposures of institutional investors' equity portfolios (e.g., Bolton and Kacperczyk, 2021, Atta-Darkua et al., 2023, and Bolton, Eskildsen, and Kacperczyk, 2024). Institutional investors perceive regulatory climate-related risks as financially material and already affecting portfolios (Krueger, Sautner, and Starks, 2020), consistent with a risk-based interpretation of our ICC results. Also supporting the pricing of climate-regulation risk, Ilhan, Sautner, and Vilkov (2021) and Sautner et al. (2023) find that firms' exposure to climate-policy risk is related to options-market risks and risk premiums. Given their forward-looking nature, our carbon burden measures could also be helpful to investors interested in constructing net-zero portfolios (e.g., Cenedese, Han, and Kacperczyk, 2023). A forward-looking perspective is also present in the hypothetical emission futures contracts that van Binsbergen and Brogger (2022) propose as a way of assessing the impact of firms' environmental initiatives. Pedersen (2025) proves that a firm's carbon burden, as we define it, is positively related to the firm's hypothetical discount rate that would substitute for taxing emissions at the SCC. Chittaro et al. (2025) and Kacperczyk (2026) make related theoretical contributions. Garleanu and Pedersen (2025) build a macro-finance model in which the premium associated with carbon-tax risk can be positive or negative. Our evidence is more consistent with a positive premium.

Greenstone, Leuz, and Breuer (2023) introduce the concept of corporate carbon damages. For a given firm, they compute these damages as the product of the firm's direct emissions in 2019 and the SCC (also obtained from the EPA), divided by the firm's profit or sales in 2019. For the average U.S. firm, these damages represent 18.5% of profit and 2% of revenue. The main difference between our studies is that they study past emissions, whereas we study future emissions. As noted earlier, future emissions are crucial to gauging the carbon externality. To give an extreme example, if emissions were widely expected to drop to zero next year and remain zero forever, any past ratio of emission damages to profits would presumably be of significantly less interest. Unlike Greenstone et al., we describe patterns in forecasted future emissions, compute their present values, compare them to firms' market values, and relate them to the cross section of expected returns. In addition to the historical emissions data they use, we also use emission forecasts, compare them to firms' emission

reduction targets, and look at not only direct but also indirect emissions, which account for over half of aggregate emissions. Finally, whereas our focus is on measurement and pricing, theirs is on disclosure and the desirability of mandatory emissions reporting.

Our emission forecast data, which come from MSCI, are informed by firms' emission reduction targets. The usefulness of those targets is supported by the evidence of Bolton and Kacperczyk (2023) and Ramadorai and Zeni (2024), who find that the firms that commit to reducing their carbon emissions indeed tend to do so subsequently. These studies use data from CDP, and the former study also uses data from the Science Based Targets initiative (SBTi). Our data are richer, because when constructing its emission forecasts, MSCI uses data not only from CDP and SBTi but also from firms' annual reports, sustainability reports, investor presentations, and regulatory filings.

A few large emitters account for the bulk of U.S. corporate emissions, consistent with right skewness in the distribution of emissions across firms (e.g., Hartzmark and Shue, 2023). Our finding of a big role for large emitters in decarbonization is consistent with the evidence of Cohen, Gurun, and Nguyen (2024) that energy producers, which tend to be large emitters, are key green innovators. The result also complements that of Berg, Ma, and Streitz (2024), who find that large emitters have reduced their emissions faster than other public firms, especially since 2015, and especially due to divestment of pollutive assets. We contribute by studying the future, showing, for example, that just the top 30 emitters fully account for the predicted decarbonization of U.S. corporations.

This paper contributes not only to the climate finance literature, but also to the broader literature on corporate externalities, which is too large to summarize here.[3] A related strand of this literature focuses on environmental damages, such as the consequences of pollution (e.g., Graff Zivin and Neidell, 2012, and Hanna and Oliva, 2015). The literature also analyzes the effects of environmental policies on technological innovation (e.g., Acemoglu et al., 2016, and Aghion et al., 2016) as well as on the behavior of firms (e.g., Greenstone, 2002, Fowlie et al., 2016, and Bartram, Hou, and Kim, 2022), consumers (Busse, Knittel, and Zettelmeyer, 2013), and the workforce (Walker, 2013).

This paper is organized as follows. Section 2 explains how we compute the carbon burden. Sections 3, 4, and 5 compute carbon burdens at the aggregate, industry, and firm levels, respectively. Section 6 analyzes the relation between carbon burden and expected

[3]For example, the literature examines effects of externalities resulting from corporate activities such as R&D (e.g., Jaffe, 1986, Jaffe, Trajtenberg, and Henderson, 1993, Audretsch and Feldman, 1996, and Bloom, Schankerman, and Van Reenen, 2013), foreign direct investment (e.g., Aitken and Harrison, 1999, Javorcik, 2004, and Blalock and Gertler, 2008), and bankruptcy (Bernstein et al., 2019).

return. Section 7 concludes.

2. Computing the carbon burden

This section explains our methodology for computing the carbon burden. Section 2.1 describes the SCC values we use. Section 2.2 discusses how we discount to the present. In subsequent sections we combine these components with forecasts of carbon emissions to compute corporate carbon burdens at the aggregate, industry, and firm levels.

2.1. Social costs of GHG emissions

As noted earlier, key inputs to the carbon burden in equation (1) are the values of SCC_τ, the dollar cost of societal damages per additional CO_2-equivalent ton of GHG emitted in τ years. Various SCC estimates exist, and their collection is evolving. Many such estimates pertain just to emissions at the present time. We use the U.S. government's latest SCC estimates as of this writing (U.S. Environmental Protection Agency, 2023).[4] The EPA provides estimates of the social cost per ton of CO_2 emitted in each future year through 2080.

The EPA explains that the values of SCC_τ are estimates of certainty-equivalent costs produced by combining four modules, each with uncertainty considered, including the compounding of uncertainty across modules. The modules rely on prominent and widely used approaches, including recommendations made by the National Academies of Science, Engineering, and Medicine. The first module, addressing socioeconomics and emissions, projects future population, income, and GHG emissions. The projections take into account the likelihood of future emissions mitigation policies and technological developments. The second module, on climate, captures the relationships among GHG emissions, atmospheric GHG concentrations, and global mean surface temperature. The outputs of the first two modules are inputs to the third one, on damages, which estimates monetized future damages from climate change by combining three damage functions (subnational, country-level, and meta-analytical). The fourth module addresses discounting. The EPA provides a series of SCC_τ for three discount rates: 1.5%, 2.0%, and 2.5% per year. We briefly postpone a discussion of discounting until the next subsection.

The values of SCC_τ are increasing in τ and decreasing in the discount rate. For example,

[4]These SCC estimates lie well within the range of the SCC estimates reported in the recent meta-analyses of Tol (2023, 2025).

when the discount rate is 2.5%, SCC_τ increases from \$128 in 2024 to \$284 in 2080. When the discount rate is 1.5%, the SCC_τ values are much higher, equal to \$356 in 2024 and increasing to \$601 in 2080. Figure 1 plots the SCC_τ values through 2080, when the EPA series end. To obtain values for subsequent years, we extend each series along a linear projection through the values for 2060 and 2080. In the plots, the SCC_τ values between those years grow virtually linearly, so we simply extend those linear trends.

The EPA estimates SCC_τ as the marginal social cost of an incremental unit of GHG emitted τ years from now, relative to the simulated baseline path of global emissions. When computing an entity's carbon burden in equation (1), we multiply SCC_τ by C_τ, the forecast of the entity's emissions τ years ahead. The smaller is C_τ, the more appropriate it is to apply the marginal cost, SCC_τ.[5] Even when C_τ represents emissions of the entire U.S. corporate sector, however, applying SCC_τ seems reasonable because U.S. emissions in any given year are small relative to the stock of carbon in the Earth's atmosphere. For example, in 2022, U.S. CO_2 emissions were just 0.16% of the CO_2 then present in the atmosphere.[6] This fraction is small because the amount of carbon emitted globally in any given year is small relative to the atmospheric stock, and also because U.S. emissions account for only 17% of global carbon emissions, based on CO_2 equivalents in 2022.[7]

An entity whose carbon burden is computed using SCC_τ is implicitly treated as the marginal emitter relative to all others. For example, our estimated U.S. carbon burden views the U.S. as the marginal emitting country with respect to the rest of the world. From a U.S. standpoint, this perspective seems sensible: it takes the rest of the world's emissions as given and asks how much additional damage U.S. emissions will inflict. Computing an individual firm's carbon burden using SCC_τ asks an analogous question from the firm's standpoint. This marginal perspective on an entity's externality is standard in public economics, where taxing an externality at its marginal social cost traces back to Pigou (1920). Summing direct-emission carbon burdens, each computed using SCC_τ, across all entities in the world

[5]Technically, with damages convex in global emissions, SCC_τ applies best to the last ton of emissions in C_τ, because every entity's C_τ is (in theory) part of the future global emission paths the EPA simulates before injecting an emission "pulse" in τ years to estimate SCC_τ. If C_τ is a small fraction of global GHGs, the marginal social cost of the first ton of emissions in C_τ is only slightly lower than the marginal cost of the last ton, a difference we approximate by zero.

[6]The National Oceanic and Atmospheric Administration (noaa.gov) reports that the deseasonalized December 2022 average CO_2 in the atmosphere reached 419.74 parts per million (PPM). Using conversion factors provided by NOAA, multiplying PPM by 2.12 converts to billions of tons of carbon, and then further multiplying by 3.67 converts to tons of CO_2, yielding a total of 3.288 trillion tons of atmospheric CO_2. The U.S. CO_2 emissions of 5.1 billion tons (see Section 3.1) represent 0.16% of this total.

[7]According to the Global Carbon Budget (globalcarbonbudget.org), global carbon emissions in 2022 totaled 10.14 billion tons, which is 37.15 billion equivalent tons of CO_2 (the conversion factor is 3.664). The U.S. GHG emissions of 6.40 billion tons (see Section 3.1) represent 17% of this global total.

would overstate global damages, but such an exercise lies well outside our analysis.[8] We focus on just the U.S. corporate sector. Some studies apply the marginal perspective at an even broader level. For example, although they do not analyze future years, Greenstone, Leuz, and Breuer (2023) multiply an EPA-estimated SCC by the sum of scope 1 emissions in 2019 for nearly 15,000 firms across many countries.

2.2. Discounting

Two discounting operations underlie the carbon burden in equation (1). First, when computing SCC_τ, the EPA discounts all future damages arising from a ton of carbon emitted τ periods from now back to that period. Second, we discount $C_\tau \times SCC_\tau$, a quantity applying τ periods ahead, back to the present.[9]

As noted earlier, the EPA computes its SCC estimates for each of three annual discount rates: 1.5%, 2.0%, and 2.5%. The EPA treats these as initial discount rates that could prevail in τ periods. It then allows discount rates for subsequent periods to comove with aggregate consumption growth, effectively using a consumption-based stochastic discount factor that implicitly recognizes emissions are likely to be high when consumption is high.

We discount $C_\tau \times SCC_\tau$ to the present using a discount rate ρ_τ, as shown in equation (1). What value for ρ_τ is appropriate? To consider this question, recall that C_τ denotes expected emissions in τ periods. Define $\tilde{C}_\tau$ as actual emissions, with $C_\tau = \mathrm{E}(\tilde{C}_\tau)$. If $\tilde{C}_\tau$ is treated as known, i.e., $\tilde{C}_\tau = C_\tau$, then the EPA advises setting ρ_τ to the τ-period real riskless rate. Doing so essentially treats SCC_τ as known also, or at least having estimation risk that does not command a risk premium. We follow the EPA's treatment of SCC_τ in this respect.

In general, $\tilde{C}_\tau$ differs from the forecast, C_τ. How should we account for the risk in $\tilde{C}_\tau - C_\tau$ when discounting $C_\tau \times SCC_\tau$? Discounting at the corporate cost of capital would be inappropriate because the risk profiles of corporate profits and emissions are quite different. For example, consider the capital asset pricing model (CAPM) and three annual U.S. time

[8] That summation is equivalent to computing equation (1) with C_τ set to predicted global emissions. Computing global damages using that equation would additionally require replacing each marginal cost, SCC_τ, by a corresponding average cost that is lower than SCC_τ, because, as widely agreed, damages are convex in aggregate emissions.

[9] Others, of course, have addressed the problems of monetizing environmental costs and discounting them to the present. Public economists have long advocated for computing similar present values in the context of natural capital accounting, a systematic way of measuring the economic value of natural resources to society. The international standard for natural capital accounting is the United Nations' System of Environmental-Economic Accounting. This framework does not mandate a specific valuation method or discount rate. For a recent corporate-finance textbook exposition, see Schoenmaker and Schramade (2023).

series: the excess stock market return (from Ken French's website), log changes in total after-tax corporate profits (from FRED), and log changes in total emissions (from Section 3.1), over the period of 1990 (the first year that emissions data are available) through 2022 (the last year that profit data are available). While log changes in profits exhibit a significant 34% correlation with the market return, log changes in emissions exhibit an insignificant correlation very close to zero (−4.9%). So, the CAPM implies a lower discount rate for emissions damages than for corporate profits. Also, log changes in emissions and profits are barely correlated—when we regress one on the other, the slope is insignificant and the R-squared only 0.07. Given their low correlation, emissions and profits can have very different covariances with the stochastic discount factor, regardless of which factor one picks.

Given the current state of the climate finance literature, it is not clear what the best approach to discounting future emissions damages is. States of the world with unexpectedly high emissions could be good or bad, depending on what agents care about. On one hand, emissions tend to be high in periods of strong economic growth, which are generally good states of the world. (This is the mechanism behind the EPA's discounting approach in constructing SCC_τ.) On the other hand, emissions can also be high in bad states of the world, such as when technological innovation fails to make progress toward renewables, or when unexpectedly high emissions cause climate-related economic disruptions. In Stroebel and Wurgler (2021)'s survey of 861 finance academics and professionals, most respondents believe that realizations of climate risk are uncorrelated with economic conditions. More research is needed to figure out the appropriate way of discounting future emissions.[10]

Meanwhile, to make progress on the question at hand, we take a simple approach to specifying ρ_τ. At the end of 2023, the date at which we compute carbon burdens, Treasury par real yields range from 1.72% at 5 years to 1.90% at 30 years.[11] Given this rather flat yield curve at levels just below 2%, one specification we choose, especially since we extend τ well beyond 30, is $\rho_\tau = 2\%$ for all τ. At that baseline value, ρ_τ includes virtually no premium for the risk associated with $\tilde{C}_\tau - C_\tau$. As discussed above, the sign of any risk premium seems ambiguous, so we also entertain both positive and negative values for the premium: 0.5% and −0.5% on top of the 2% baseline. We thus entertain three values for ρ_τ: 1.5%, 2.0%, and 2.5%.

Only a partial coincidence is that our three ρ_τ values coincide with the EPA's initial discount rates used in constructing their three SCC_τ series. We could of course specify

[10] Joint modeling of economic dynamics and the dynamics of climate change is beyond the scope of this paper. See Giglio, Kelly, and Stroebel (2021) for a discussion of some of the challenges in figuring out the risk premium associated with climate damages, including whether its sign is positive or negative.

[11] See the "Data" menu at https://home.treasury.gov/.

other risk premia as deviations from a 2% riskless rate, but we avoid doing so to simplify the analysis and give readers just three rates to digest. Still, with three SCC_τ series and three ρ_τ values, there are nine possible pairings of an SCC_τ series with a ρ_τ value. To simplify the presentation further, we report carbon burdens for just three of the pairings: (1.5%, 1.5%), (2.0%, 2.0%), and (2.5%, 2.5%). The middle combination is reasonably viewed as the baseline case, while the first and third produce the highest and lowest values of the carbon burden. Recall from Figure 1 that SCC_τ is decreasing in the corresponding discount rate, and of course the discount factor in equation (1) is decreasing in ρ_τ.

The issues of uncertainty about ρ_τ and $\tilde{C}_\tau$ would be simultaneously solved by the existence of emissions futures contracts similar to those proposed by van Binsbergen and Brogger (2022). Imagine a contract paying SCC_τ dollars for each ton of emissions that a firm emits τ years from now, where SCC_τ is an SCC forecast agreed upon today. If we had such contracts' market prices for each future τ, we could sum those prices across $\tau = 1, ..., \infty$ to obtain the market's assessment of the firm's carbon burden, conditional on the SCC forecasts. In such an imaginary world, carbon burden estimates would be more precise.

The range of values for ρ_τ, 1.5% to 2.5%, is supported by expert views. Drupp et al. (2018) survey economists who are experts on social discounting, having published at least one paper on this topic in a leading economics journal between 2000 and 2014. The distribution of the risk-free social discount rates across over 200 survey responses has a median of 2% and a mean of 2.3%. There is "a surprising degree of consensus among experts," with 77% of experts finding the median discount rate of 2% acceptable, and 92% of them being comfortable with the discount rate somewhere between 1% and 3%. The same median and mean, 2% and 2.3%, emerge also from an independent survey of Howard and Sylvan (2020), who poll all authors who had published at least one article related to climate change in a top-25 economics journal or top-six environmental economics journal since 1994, obtaining 216 valid responses. The EPA's discount rates lie between the 1.4% used by Stern (2006) and the 2.6% found by Giglio, Maggiori, and Stroebel (2015) as the long-run discount rate for real estate cash flows. Giglio et al. (2021) argue that the 2.6% value provides an upper bound on the discount rates for long-term cash flows from investments in climate change abatement.

3. The aggregate U.S. carbon burden

We use data on forecasts of U.S. GHG emissions (Section 3.1) to assess the carbon burden for the U.S. corporate sector as a whole (Section 3.2). Recall that we equate corporate emissions

with total U.S. emissions, given that virtually all emissions are either direct (scope 1) or indirect (scopes 2 and 3) emissions of some company. We also interpret the burden's magnitude (Section 3.3) and consider its potential reductions from the country's past commitment to the Paris Agreement (Section 3.4).

3.1. Forecasts of U.S. GHG emissions

To estimate carbon burdens as of year-end 2023, we first obtain forecasts of emissions in the U.S. for 2024 and beyond. We construct aggregate GHG emissions by adding up three types of emissions: energy-related CO_2, non-energy-related CO_2, and non-CO_2 GHGs.

The first type, energy-related CO_2, accounts for the largest fraction of GHG emissions, by far. The U.S. Energy Information Administration (EIA) provides annual forecasts of U.S. energy-related CO_2 emissions through 2050. The forecasts come from the EIA's National Energy Modeling System, which takes a general equilibrium approach to modeling U.S. energy markets and projecting production, imports, exports, conversion, consumption, and energy prices (U.S. Energy Information Administration, 2023b). The system has 14 modules devoted to separate sources of supply and demand, conversion, and various economic and policy channels. We use the EIA's reference-level forecasts for 2024 through 2050.[12]

The second type, non-energy-related CO_2, is the smallest part of GHG emissions. Non-energy-related emissions come from sources such as agriculture, industrial processes, and waste. Lacking forecasts for this emission type, we approximate them based on historical CO_2 emission breakdown data.[13] Averaging across 1990 through 2022, non-energy-related CO_2 emissions account for 3.6% of total CO_2 emissions. Assuming this share remains unchanged going forward, we apply it to the EIA's forecasts of energy-related CO_2 emissions to obtain annual non-energy-related CO_2 emission forecasts through 2050.

The third type of emissions includes non-CO_2 gases such as methane and nitrous oxide. The EPA provides forecasts of U.S. non-CO_2 GHG emissions from all sources, both related and unrelated to energy, through 2050. To construct its forecasts, the EPA combines historical emissions data and trends based on projected activity.[14] We use linear interpolation to

[12]The data can be obtained via the EIA website (eia.gov), searching first for "Annual Energy Outlook 2023" and then selecting Table 18. The total CO_2 values provided there are plotted and identified as the "reference" case in the publication, U.S. Energy Information Administration (2023a).

[13]See U.S. Environmental Protection Agency (2024). These data, which track U.S. emissions by source back to 1990, can be obtained via the EPA's Greenhouse Gas Inventory Data Explorer website.

[14]See U.S. Environmental Protection Agency (2019) for more detail on the EPA's methodology. The data can be obtained via the EPA's Non-CO_2 Greenhouse Gas Data Tool website.

convert the forecasts from their five-year frequency to an annual series.

We sum up the forecasts across the three emission types to compute aggregate U.S. GHG emission forecasts through 2050. Beyond 2050, we project the same annual growth rate as in the aggregate emission forecasts from 2023 to 2050, which is -0.458%. The solid line in Figure 2 plots our resulting reference forecasts of U.S. aggregate GHG emissions.

3.2. The U.S. carbon burden

We compute the aggregate U.S. carbon burden by setting the values of C_τ in equation (1) equal to the forecasted GHG emissions plotted in Figure 2. We report the carbon burdens associated with three future periods, all beginning in 2024. The first period ends in 2050, the second in 2080, and the third covers all future years. Recall that 2050 is when our emission forecasts end, and 2080 is when our social cost estimates end, so the periods with those ending dates avoid one or both of the approaches we take to extend the two series.

Panel A of Table 1 displays the U.S. carbon burden in dollar terms. The values cover a wide range, from \$17.4 trillion, for the shortest period and highest discount rate, to \$178.8 trillion, for the entire future and the lowest discount rate. When pairing all future years with the 2% discount rate, our baseline value, the U.S. carbon burden is \$87.1 trillion.

To begin putting these dollar amounts into perspective, we divide them by the total value of U.S. corporate equity as of year-end 2023, which is equal to \$66.4 trillion.[15] Panel B of Table 1 shows that these ratios range from 26% to 269%. For the 2% discount rate, the U.S. carbon burden for all future years is 131% of total U.S. corporate equity value. Even the burden for just the shortest future period ending in 2050, which relies on neither of our series-extension procedures, is 44% of equity value. In brief, the U.S. carbon burden is large.

In addition to comparing the carbon burden to the value of corporate equity, we also compare it to the combined value of equity and debt. The value of U.S. corporate debt at year-end 2023 is about \$27.3 trillion, which includes \$12.1 trillion of bonds and \$15.2 trillion of loans.[16] Adding this value to the \$66.4 trillion value of equity, total value of equity and debt is about \$93.7 trillion. Our baseline estimate of the U.S. carbon burden, \$87.1 trillion,

[15]This amount equals the value of total issues net of holdings of foreign equities by U.S. residents. It includes both publicly traded equity and closely held equity, where the latter includes both S and C corporations. See Table L.224 of Board of Governors of the Federal Reserve System (2024).

[16]These amounts come from Tables L.213 and L.214 of Board of Governors of the Federal Reserve System (2024). The bond value equals total liabilities net of holdings of foreign bonds by U.S. residents. The loan value equals total liabilities minus those of households, governments, and foreign entities.

thus represents 93% of total value of corporate equity and debt.

The corporate sector's carbon burden is substantial compared not just to corporate wealth but also to total wealth. The Federal Reserve computes total U.S. net wealth as the value of tangible assets controlled by the household, nonprofit, business, and government sectors of the U.S. economy, net of U.S. financial obligations to the rest of the world. At year-end 2023, total U.S. wealth is about $143.6 trillion (see Table B.1 of Board of Governors of the Federal Reserve System, 2024). The carbon burden of $87.1 trillion thus constitutes 61% of U.S. national wealth—a very substantial fraction.

How does the corporate carbon burden compare to the consumer surplus produced by companies, which is a positive component of their value to society? We do not have a present value of consumer surpluses in future years, but we can compare emission damages to consumer surplus in a single year. Pellegrino (2025) estimates a U.S. consumer surplus of $11.1 trillion in 2021. In that year, the U.S. emitted 6.35 billion tons of carbon (GHG in CO_2-equivalent tons), and the EPA's baseline SCC estimate was $197 per ton. Those values imply a social cost of $1.25 trillion for 2021 emissions, which is 11% as large as the consumer surplus. The U.S. corporate sector thus produces a consumer surplus far larger than its emissions damages. This conclusion accords with Allcott et al. (2025), who find that consumer surplus well exceeds emission damages in a sample of 74 large companies.

Another single-year comparison reveals that emission damages are also modest relative to U.S. output. For example, in 2023, when its GDP was $27.4 trillion, the U.S. emitted 6.28 billion tons of carbon. Multiplying the latter by that year's SCC estimate of $204 per ton gives a social cost of $1.28 trillion, which is 4.7% of the 2023 U.S. GDP. One might be struck by how modest this fraction is when compared, for example, to the large ratio of carbon burden to equity market value noted earlier, equal to 131% at the baseline 2% discount rate. The gap between the latter discount rate and the cost of equity, along with the corporate profit margin, can account for the difference (as explained in the Appendix).

3.3. Interpreting the burden's magnitude

Dollar values of damages from carbon emissions are easier to interpret when compared to meaningful benchmarks. There are various choices for the latter, as illustrated above, but the remainder of the study will focus on the benchmark used in Panel B of Table 1, dividing an entity's carbon burden by its market value of equity. Our analysis at the firm level helps guide this choice, given that a firm's equity holders are more directly connected than other

stakeholders to management decisions affecting emissions. Moreover, the firm-level analysis includes asset pricing implications, with future carbon taxes entertained as a priced risk to which firms with high carbon burdens are more exposed. Dividing carbon burden by equity value translates such exposure to rates of return earned by shareholders.

While the U.S. carbon burden is large when compared to the value of corporate equity, readers should bear several points in mind when interpreting the numbers. First, the carbon burdens we compute are most reasonably viewed as status-quo estimates that exclude future changes in policy. Recall from Section 3.1 that our calculations are based on emission forecasts from the EIA and EPA. The EPA's "projections include the impact of existing GHG reduction policies to the extent they are reflected in historical data but exclude additional GHG reductions" (U.S. Environmental Protection Agency, 2019). Similarly, the EIA's forecasts incorporate "only current laws and regulations" as opposed to "targets associated with yet-to-be-developed policy" (U.S. Energy Information Administration, 2023a). One potential future policy is a carbon tax. As noted earlier, the carbon burden measures the corporate sector's externality in the absence of such a tax. If a carbon tax is imposed, future emissions could well be reduced below the reference forecasts.

Absent such reductions, our results show that if carbon is taxed at a rate equal to the SCC, the present value of the future taxes (i.e., the carbon burden) would be a substantial fraction of corporate equity. The tax would not reduce corporate equity value by the full carbon burden, however, because some of the tax's incidence would fall on consumers rather than equityholders. In particular, consumers would likely bear much of the incidence of a tax on GHGs emitted in producing goods having inelastic demand.

Measuring the U.S. carbon burden as a fraction of total corporate equity should not be construed as assigning responsibility for the burden to the corporate sector. Responsibility for the carbon burden is shared more broadly. Consider a country's choice between generating electricity using nuclear plants versus burning fossil fuels, which has first-order implications for carbon emissions. Countries differ in this choice; for example, nuclear power plants generated 68% of France's electricity in 2021, whereas the U.S. fraction is only 19%, and Germany no longer operates any nuclear reactors.[17] It seems difficult to say how much of the choice can be attributed to a country's corporate sector, let alone its electric utilities, as opposed to the country's body politic. Similarly, it seems difficult to say how much responsibility for the combustion of gasoline lies with the corporate sector, let alone its automobile and oil companies, as opposed to the household sector.[18] At the same time, within

[17] See https://www.eia.gov/todayinenergy/detail.php?id=55259.

[18] With less nuance, Callahan and Mankin (2025) assign complete responsibility to major fossil fuel com-

the corporate sector, identifying large sources of emissions is potentially useful information for a country seeking to reduce its carbon burden. Therefore, in subsequent sections, we analyze carbon burdens at the industry and firm levels as well.

3.4. Potential reductions under the Paris Agreement

The Paris Agreement is an international treaty adopted in 2015 that calls for substantial reductions in global GHG emissions. U.S. participation in the agreement was withdrawn in 2017, reinstated in 2021, and withdrawn again in 2025. Further reversals seem difficult to rule out, given the politics. We analyze the potential reductions in the U.S. carbon burden that would be achieved by participating and fully meeting the agreement's U.S. emission targets versus not doing so. For the targets, we take the country's most recent pre-withdrawal commitments under the agreement (known as "nationally determined contributions"). When it rejoined the agreement in 2021, the U.S. targeted reductions in its emissions, relative to the 2005 level, of at least 26% by 2025 and 50% by 2030. As noted earlier, our emission forecasts, which we plot in Figure 2 and use as our reference levels, do not include changes in emission targets yet to be implemented. In particular, those forecasts appear not to incorporate the cuts targeted under Paris: the forecast for 2030 is only 25% below the 2005 level, compared to a reduction of at least 50% targeted by Paris. We therefore interpret the difference between our forecasts and the levels targeted by Paris as the potential reductions implied by the agreement.

We consider two Paris scenarios for emission levels beyond 2030. Both scenarios have emissions relative to the 2005 level be 26% lower in 2025 and 50% lower in 2030.[19] The 2005 level is 7.4 billion CO_2-equivalent tons, so a 50% reduction implies a 2030 level of 3.7 billion tons, which is 2/3 (67%) of the reference-level forecast of 5.5 billion tons in that year. In the first scenario, this 2/3 ratio is maintained in all subsequent years, and the resulting emission levels are plotted as "Paris scenario 1" in Figure 2. Our second Paris scenario, more conservative, merely accelerates reductions that are forecast to occur later otherwise. That is, emissions remain at 3.7 billion tons in the years following 2030 until that level exceeds the reference level, at which point the scenario follows the same path as the reference level. The resulting forecasts are plotted as "Paris scenario 2" in Figure 2.

Table 2 reports the estimated reductions in the U.S. carbon burden, measured at year-end 2023, under the first Paris scenario. Panel A reports the dollar amounts, Panel B divides

panies for damages associated with the companies' past emissions, direct and indirect.

[19] We linearly interpolate from the current level to those points, consistent with the plot in the U.S. submission to the United Nations registry of national contributions (unfccc.int/NDCREG).

those amounts by the value of U.S. corporate equity, and Panel C divides the dollar amounts by the corresponding U.S. carbon burdens reported in Panel A of Table 1. We see from Panel C that adherence to the Paris Agreement would reduce the U.S. carbon burden by between 29% and 32% across the three discount rates and three future periods.

Table 3 reports reductions under the second scenario. Panel B shows that the bulk of reductions occur by year 2080, not surprisingly given that this scenario simply front-loads reductions otherwise occurring later. Even in this more conservative scenario, Panel C shows that the Paris Agreement reduces carbon burdens by 28% through both 2050 and 2080, for all three discount rates. Using the 2% discount rate and all future years, the reduction is 21%. All of these reductions are substantial.

4. Carbon burdens across industries

This section analyzes carbon burdens across industry sectors. We use firm-level emission forecasts from MSCI, which we describe in Section 4.1, to compute the carbon burden of a typical firm in each sector. We analyze those burdens in Section 4.2. For firms that have targets for future emissions, we compare those targets to MSCI's forecasts in Section 4.3.

4.1. MSCI firm-level emission forecast data

We downloaded the MSCI Climate Change Metrics data from the MSCI ESG Manager in 2024, as soon as they were made available to the academic community by the newly established MSCI Sustainability Institute through its Climate Data Knowledge Program. Our primary interest is in MSCI's forecasts of individual firms' future emissions, which we use not only in this section but also in Section 5. These forecasts are unique and valuable for the computation of firm-level carbon burdens, which are inherently forward-looking. We obtain MSCI's historical emissions data from the same source.

MSCI provides firm-level forecasts of scope 1, 2, and 3 emissions for each year from 2023 through 2050. To construct its forecasts, MSCI collects firms' decarbonization plans and evaluates them, including their credibility. To collect data on firms' future emission targets, MSCI studies firms' publicly available documents, such as annual reports, sustainability reports, CDP reports, the Science Based Targets initiative, Forms 10-K and 20-F, and investor presentations. MSCI allows firms to verify or amend their targets, and even input new ones, through a dedicated platform. MSCI also uses natural language processing

software to identify new target announcements for its biweekly data updates.

Among the 2,851 U.S. firms in its sample, MSCI identifies 798 firms, including most large emitters, as having emission targets. For the firms with targets, MSCI offers two types of emission projections: target-based and credibility-adjusted. The former projections take firms' emission reduction targets at their face value. The latter projections make adjustments after assessing the targets' credibility. These credibility-adjusted projections represent MSCI's forecasts. If a firm has no future emissions target, MSCI assumes that its emissions will grow at the business-as-usual rate of 1% per year.[20]

To compute the target-based projections, MSCI assumes that firms will meet their future targets exactly and uses interpolation. First, MSCI interpolates emissions linearly between the firm's most recent emissions and the first target emission value. If the firm has multiple future targets, MSCI interpolates linearly between each pair of subsequent targets. After the last target year, MSCI assumes zero growth in emissions until 2050.

To compute the credibility-adjusted projections, MSCI adjusts the target-based projections after performing a target credibility assessment. The purpose of this assessment is to penalize stated decarbonization trajectories that lack credibility. For example, MSCI assigns low credibility to plans setting scope 3 net-zero targets in the distant future with no interim targets. Faced with a target that it does not view as fully credible, MSCI projects higher future emissions compared to target-based values. Specifically, MSCI computes its credibility-adjusted emissions forecast for firm n in future year T as follows:

$$\text{Forecast}_{n,T} \;=\; w_n \times \text{Target}_{n,T} \;+\; (1 - w_n) \times \text{Base}_{n,T}\,, \tag{2}$$

where $\text{Target}_{n,T}$ is the target-based forecast of firm n's emissions in year T and $\text{Base}_{n,T}$ is the forecast of the firm's emissions assuming 1% annual emissions growth between today and year T. MSCI chooses the firm's "credibility weight" w_n after evaluating the firm's decarbonization plan in terms of its ambition, comprehensiveness, and feasibility. Larger w_n's are more likely to go to firms that have, for example, at least one short-term target, at least one externally validated target, a track record of achieving past targets, and a current trajectory to meet their targets.

MSCI's emissions forecasts go out to year 2050, as do the aggregate forecasts used in Section 3. To extend the firm-level forecasts beyond 2050, we follow the same procedure as in Section 3, extrapolating the (negative) growth trend in the aggregate forecasts from 2023 to 2050 and then applying that trend to each firm.

[20]To explain this choice, MSCI notes that 1% is the annual global emissions growth rate from 2009 to 2019, adjusted for GDP, according to the 2020 United Nations Environment Programme Emissions Gap Report.

4.2. Carbon burdens by industry

We compute firm-level carbon burdens as of the end of 2023 by substituting MSCI's forecasts of firms' future emissions into equation (1). As before, we use the EPA's SCC estimates and three discount rates. We analyze U.S. firms in the intersection of the MSCI and CRSP/Compustat databases. We scale each firm's carbon burden by the firm's market cap, denoting the resulting ratio by CB/M. For firms with multiple common share classes, we aggregate them to compute firm-level market cap. We assign firms to 12 industries following the SIC-code classifications of Fama and French, which we obtain from Ken French's website. For each industry, we define a typical firm's CB/M as the weighted average of the CB/M ratios across all firms in this industry, using market capitalization weights.

Table 4 reports properties of carbon burdens for the 12 industries. In Panel A, we compute carbon burdens for future years through 2050, while Panel B includes all future years. We use our baseline discount rate of 2%. The first three columns report the CB/M ratios for a typical firm in each sector. The first column considers just scope 1 emissions, the second column adds scope 2, and the third column sums all three scopes. The values reported in the second and third columns must be interpreted with caution. They correctly represent CB/M for a typical firm in the sector, but they overstate CB/M for the sector as a whole, due to double-counting. For example, most scope 2 emissions are scope 1 emissions for utilities, and the same ton of carbon can be included in scope 3 emissions of multiple firms.[21] In contrast, the values reported in column 1 can also be interpreted as CB/M for each sector as a whole because there is no double-counting of scope 1 emissions.

Table 4 shows that carbon burdens differ greatly across industries. For scope 1, a typical utility has a carbon burden through 2050 that is 2.70 times its market cap, and a typical energy firm has the second-largest ratio at 1.06. In contrast, for six industries, the typical firm's ratio is 0.05 or less. Adding the later years more than doubles the largest values, with utilities and energy increasing to 6.94 and 2.95, but there are still five industries at 0.05 or less. Adding scope 2 changes the picture very little, unlike adding scope 3.[22] The energy

[21] The double-counting of emissions occurs only across firms, because for any given firm, scopes 1, 2, and 3 are mutually exclusive. Also, there is no double-counting of emissions in our aggregate analysis in Section 3 because we do not add up firm-level forecasts; instead, we use U.S. agencies' forecasts of aggregate emissions.

[22] Although scope 3 greatly double-counts emissions, it captures more than half of aggregate emissions not captured by firms' scopes 1 and 2. We can estimate those aggregate emissions captured by scope 3 by subtracting the corporate sector's scope 1 emissions from total U.S. GHG emissions, because virtually all of the latter are part of at least one firm's scope 3 emissions. For 2023, that calculation gives $6,277 - 2,814 = 3,463$ million tons, or 55% of total U.S. emissions. Our calculation does not subtract scope 2 (in addition to scope 1) from total U.S. emissions, because total scope 2 is already counted in total scope 1, as noted earlier. The value of 2,814 is the sum of all 2023 scope 1 emissions across all firms in the MSCI database. Summing firm-by-firm scope 3 emissions, even if they were accurately measured, would not produce a meaningful

industry's scope 3 emissions subsume much of aggregate emissions, so a typical energy firm's carbon burden is the largest by far, 66 times its market cap when including all future years. The carbon burdens in other industries also become much larger when including scope 3. For example, typical firms in four other industries have ratios of 10 or higher in Panel B. One of them is the financial industry, whose ratio for scope 1 is just 0.01.

Financial firms' direct emissions are small, given the sector's service-based nature. Their scope 3 emissions are large, however, because they include the emissions of companies and projects financed by financial institutions. For example, scope 3 emissions are high for banks lending to fossil fuel companies and investment funds holding shares in high-emitting industries. The GHG Protocol includes these "financed emissions" as part of scope 3. While financial firms have little control over their current financed emissions, they have more control over future emissions, because they provide financing to replace emitting real assets when those assets depreciate. An emitter's inability to externally finance investments in emitting assets could potentially restrict the emitter's future emissions.

The last three columns of Table 4 report the "future/present" ratio—the ratio of the carbon burden from emissions in all future years to the burden from a single year's emissions in 2023—for a typical firm in each sector. We define a typical firm's ratio as the weighted average of the corresponding ratios across all firms in the given sector, using the burdens from 2023 emissions as weights.[23] The future/present ratio generally ranges in the mid-20s when including emissions through 2050, and it is roughly three times larger when including all future years. In the latter case, the future/present ratio is akin to a price/dividend ratio, which divides total discounted expected future dividends (price) by last year's dividend. Instead of dividends, here we have social costs, discounted at a rate conceptually distinct from the cost of capital used to discount dividends. The discount rate and the social costs per ton of future carbon are common across firms, so differences across industries in the future/present ratios in Table 4 arise just from differences in forecasts of emissions growth.

The future/present ratio exhibits notable variation across industries. For example, in Panel B, the ratio for scope 1 ranges from 49 for telecom to 80 for retail (shops), a value 63% higher. When including scopes 2 and 3, the ratio ranges from 61 for business equipment to

aggregate quantity, because scope 3 inherently double-counts emissions across firms, unlike scope 1.

[23]This definition differs slightly from the typical firm's definition in the first three columns of Table 4, on purpose. Both definitions allow us to interpret the values for scope 1 as pertaining to a typical firm in the sector as well as to the sector as a whole. Generically, an X-weighted average of firms' Y/X is equal to the sum of Y across firms divided by the sum of X across firms. Given that there is no double-counting of scope 1 emissions across firms, the latter ratio represents the sector's Y/X ratio. In all columns, Y is the carbon burden from emissions in all future years. In the first three columns, X is market capitalization; in the last three columns, X is carbon burden from 2023 emissions.

92 for energy, 51% higher. In short, computing carbon burdens of just last year's emissions tells an incomplete story. Not only are such carbon burdens much lower than when including the future, but they also omit differences in forecasts of emissions growth.

Differences in emission-growth forecasts are apparent from the "forecast" columns in Table 5, which report MSCI's forecasts of cumulative growth rates in each industry's emissions through 2050. We compute these industry-level growth rates from MSCI's firm-level forecasts. For scope 1, the two industries with especially large carbon burdens, utilities and energy, have forecasted growth rates that differ substantially: -37% versus -24%. When all three scopes are included, MSCI predicts that six of the industries will increase their emissions through 2050, whereas the other six will reduce their emissions.

In the Appendix, we report each industry's carbon burden as a fraction of the total burden across industries. For example, based on direct emissions, utilities account for 37% of the total, and energy accounts for 20%. Five industries have shares below 1%: business equipment, durables, health, money, and telecom.

4.3. Emission targets versus forecasts

The "target" columns in Table 5 report the targeted emission growth rates for firms that have emission targets according to MSCI, industry by industry. For the same firms within each industry, we compute their total 2050 targeted emissions and divide them by the 2050 emissions implied by MSCI forecasts. The resulting value appears in the "ratio" columns of Table 5. In essence, the closer the ratio is to 1, the more realistic the target, because a small gap between the target and the forecast implies the target is unlikely to be missed by much.

For scope 1, all of the ratios are well below 1, indicating targets that are too optimistic. All industries target substantial emission reductions, but the 92% and 89% reductions targeted by utilities and chemicals seem the least realistic, with target-to-forecast ratios of just 0.11 and 0.12, respectively. The non-durable sector's targeted reduction of 47% is the most modest, but it also seems the most realistic, with a ratio of 0.5. As in Table 4, adding scope 2 makes little difference, but things change when adding scope 3. First, the targeted reductions become less ambitious. Second, the targets become more realistic, in that the target-to-forecast ratio increases for every industry. The most realistic industry, energy, has a ratio of 0.70, far above its scope 1 ratio of 0.16. One interpretation is that firms set more realistic targets for emissions that they are less able to control.

5. Carbon burdens across firms

In this section, we analyze the cross section of carbon burdens for U.S. firms. We compute firms' carbon burdens as of the end of 2023, as described in Section 4.2. Moving beyond the industry-level analysis in Section 4 seems useful because firm-level carbon burdens exhibit substantial intra-industry variation. To demonstrate this fact, we show that the cross-sectional variation in firms' carbon burdens is far from explained by industry fixed effects. Specifically, we run cross-sectional regressions of firm-level log carbon burdens on industry fixed effects, both with and without controlling for the firm's log market capitalization. We consider three dependent variables, all in logs: unscaled carbon burden, CB/M, and carbon burden divided by the burden from the firm's emissions in year 2023 only. Carbon burdens are based on the 2% discount rate and emission forecasts for all future years.

Table 6 shows adjusted R-squareds from these regressions. Panel A (B) reports the R-squareds for specifications in which industry fixed effects are computed based on the Fama-French industry classification covering 49 (12) industries. All R-squareds in the table are far below 1, peaking at 0.655. Most R-squareds are well below 0.5, especially when carbon burdens are scaled. The relatively low R-squareds indicate substantial intra-industry variation in firms' carbon burdens. In addition, the R-squared values in Panel A are only modestly larger than those in Panel B, indicating that 12 industries do a decent job in capturing the industry-level variation in carbon burdens. This fact provides support for our results in Section 4, in which we use only 12 industries, for ease of exposition.

5.1. Magnitudes of firms' carbon burdens

Figure 3 plots the distribution of CB/M across firms. There are four panels, as we consider two emissions categories (scope 1 and scope 1+2+3) and two ways of computing the carbon burden (based on all future years and only through 2050). Each panel plots the cumulative distribution function of CB/M, weighting each firm equally. That is, for any given value of CB/M, we plot the fraction of firms whose CB/M is smaller than that value.

Panel A of Figure 3 shows that the CB/M ratios vary greatly across firms. For most firms, the carbon burden associated with their direct (scope 1) emissions represents only a small fraction of the firm's market capitalization. For example, 55% of firms have CB/M ratios smaller than 0.05 under the baseline 2% discount rate. However, the distribution of CB/M is heavily right-skewed, and some firms' CB/M ratios are very large. For example, 13% of firms have CB/M ratios greater than 1. These firms' carbon burdens exceeds their market

capitalizations; that is, the present value of their future carbon costs to society exceeds the present value of their future dividends to shareholders. Of course, firms with large carbon burdens are not necessarily undesirable from a social planner's perspective, as such firms can also provide society with large benefits such as consumer surplus.

Not surprisingly, carbon burdens are larger when the discount rate is smaller, and vice versa. For example, when the discount rate is 2.5%, only 10% of firms have CB/M > 1, but when the rate is 1.5%, we observe CB/M > 1 for 18% of firms. For all three discount rates, there are many firms whose carbon burden exceeds their market capitalization.

Firms' carbon burdens are clearly larger when we consider not only direct but also indirect emissions. Panel C of Figure 3 plots the distribution of CB/M based on total (scope 1+2+3) emissions. For the 2% discount rate, 77% of firms have CB/M ratios greater than 1. The proportion is 66% for $\rho = 2.5\%$ and 87% for $\rho = 1.5\%$. We thus see that, based on total emissions, most firms' carbon burdens exceed the firms' market capitalizations. Of course, these percentages must be interpreted with the understanding that a given ton of carbon can appear in multiple firms' total emissions, due to double counting across firms.

Figure 4 is a value-weighted counterpart of Figure 3. Whereas Figure 3 plots the fraction of firms whose CB/M is below each x-axis value, Figure 4 plots the fraction of total market capitalization belonging to firms whose CB/M is below each x-axis value. The fractions in Figure 4 are larger than in Figure 3. This is not surprising, because the largest firms at the end of 2023 are mostly technology firms, which are relatively light emitters. For example, for scope 1 and the 2% discount rate, 75% of total market capitalization belongs to firms with CB/M < 0.07. Nonetheless, the cross-sectional dispersion in CB/M is large, and 10% of total market capitalization belongs to firms with CB/M > 1.

When we consider not only direct but also indirect emissions, the proportion of total market capitalization belonging to firms with CB/M > 1 is quite a bit larger. For example, based on total emissions and the 2% discount rate, half of total market capitalization belongs to firms whose carbon burdens exceed their market capitalizations.

5.2. Future versus present emissions

Carbon emissions are persistent: high emitters today are likely to be high emitters tomorrow. As a result, high emitters today tend to have high carbon burdens. When assessing a firm's carbon externality, is it necessary to consider the firm's future emissions, or could we simply look at its recent emissions? Put differently, do MSCI's emission forecasts contain much

information that is not already contained in firms' recent emissions?

To answer these questions, we compute each firm's future/present ratio, as analyzed previously at the industry level. The numerator of this ratio is the carbon burden computed from future emission forecasts through 2050, and the denominator is the burden from the firm's emissions in year 2023 only. If the ratio turns out to be equal across firms, then MSCI's emission forecasts do not add information beyond recent emissions.

Figure 5 plots the distribution of the future/present ratio across firms. To avoid spikes in the histograms, we exclude firms that either do not have an emission target or have a target that MSCI deems uninformative; recall that for such firms, MSCI forecasts a 1% emissions growth per year. In Panel A, which focuses on direct emissions, the sample includes 696 firms; in Panel B, which focuses on total emissions, it includes 353 firms. In both panels, the future/present ratio is quite dispersed across firms, taking on values as low as 0.5 and as high as 30. Therefore, while recent emissions contain significant information about a firm's carbon externality, they do not paint the full picture.

The future/present ratios are dispersed across firms because MSCI's forecasts of future emission growth are quite dispersed. Figure 6 plots the cross-sectional distribution of firms' cumulative forecasted emissions growth rates, computed as the forecast of the firm's emissions in 2050 divided by the firm's emissions in 2023, minus 1. As in Figure 5, we exclude firms for which MSCI forecasts 1% emissions growth. The figure shows a wide distribution of growth rates, ranging from -100% to +33%. For most firms, emissions are predicted to fall by 2050, in some cases to zero. For some firms, they are predicted to rise. The wide distribution in Figure 6 helps us understand the wide distribution in Figure 5.

5.3. Determinants of future emission growth

Do the forecasted emission growth rates differ between high and low emitters? To answer this question, Figure 7 shows a binscatter plot of firms' cumulative future emissions growth, computed as in Figure 6, against the firms' "current" emissions in 2023, measured in logs. For both direct and total emissions, we observe a strong, negative relation between current emissions and future emission growth rates. Higher emitters have lower forecasted emissions growth rates. For direct emissions, this growth rate is -14% for the top 5% of emitters but $+25\%$ for the bottom 5% of emitters. The latter growth rate is positive because Figure 7 includes all firms, including those for which MSCI forecasts 1% annual growth. If we exclude those firms, the relation remains negative. In that smaller set of firms, the future growth

rate of direct emissions is −47% for the top 5% of emitters but −17% for the bottom 5% of emitters (see the Appendix). The negative cross-sectional relation between current emissions and future emission growth rates is clearly economically significant.

The relation is also statistically significant. This is clear from Table 7, which reports results from cross-sectional regressions of future emission growth rates on current emissions and other firm characteristics. The dependent variable is the annualized growth rate of a firm's emissions from 2023 to 2050, computed from MSCI forecasts. The independent variables include the log of current emissions, the book-to-market ratio, investment, climate score, and revenue growth, whose definitions are in the caption of Table 7. We measure all regressors at the end of 2023. Because the indirect emissions that constitute scope 3 are especially difficult to quantify, including those emissions as an independent variable risks an error-in-variable problem.[24] Therefore, we run these regressions only for scopes 1 and 1+2, with and without industry fixed effects. In all four specifications, current emissions enter with a significantly negative slope, with t-statistics ranging from −8.87 to −12.53.

The other four regressors exhibit weaker relations to forecasted emission growth rates. Book-to-market enters with a positive slope, indicating larger emission increases for value firms, but the relation is only marginally significant. Investment also enters with a positive slope, significant in the last specification, pointing to larger emission increases for firms that invest more. Only the climate score enters with consistent significance across the four specifications. Its slope estimate is always negative, with t-statistics ranging from −3.26 to −3.40, indicating larger emission declines for "greener" firms. This association could well be reverse-causal, in that firms with more ambitious emission targets could be rewarded by MSCI with higher climate scores. We do not analyze causality; we are simply trying to explain the variation in MSCI's forecasted emission growth rates. We explain relatively little of it: adjusted R-squareds range from 10% to 12.2%.[25] Clearly, MSCI's approach to forecasting emissions is more sophisticated than a linear regression with five regressors.

Both Figure 7 and Table 7 show that future emissions are expected to decline markedly for high-emitting firms. This result is so strong that a handful of the largest emitters are responsible for the entire drop in emissions expected in the U.S. corporate sector, as we show in Figure 8. This figure plots the time series of direct emissions aggregated within two subsets of firms: the 30 largest emitters as of 2022 and the 2,411 remaining firms. We also

[24]MSCI estimates a firm's scope 3, rather than simply taking the firm's reported value, but the inherent nature of scope 3 surely makes the estimate substantially noisier than scopes 1 and 2.

[25]The sample behind Table 7 includes also firms for which MSCI forecasts 1% annual growth. If we exclude those firms, the results look similar—both current emissions and the climate score retain significantly negative slopes in all four specifications, and the other regressors are almost never significant. See the Appendix.

plot the total emissions of all 2,441 firms. In years through 2022, emissions are historical values from MSCI; after 2022, emissions are from MSCI's forecasts.

Figure 8 shows that aggregate corporate emissions have declined from 2.7 to 2.1 billion metric tons between 2008 and 2022, and that they are expected to decline further to 1.5 billion metric tons by 2050. This steady decline is not surprising, given the ongoing decarbonization of the U.S. economy. What is more surprising is the outsized role of the top 30 emitters. First, these emitters account for a substantially larger share of aggregate emissions than the remaining 2,411 firms. Second, the top 30 emitters account for just about all of the expected aggregate decline in emissions by 2050. Essentially no decline is expected for the other 2,411 firms. The disproportionate influence of the top 30 emitters is apparent also from pre-2022 historical emissions. In short, all of the decarbonization of the U.S. corporate sector by 2050 is expected to come from the 30 largest emitters.[26]

As an alternative to firm-level emission forecasts from MSCI, we also consider forecasts from a simple vector autoregression (VAR) model that uses historical emissions data from MSCI and Trucost to forecast individual firms' future emissions. The resulting estimates of carbon burdens are similar to their MSCI-forecast-based counterparts for large emitters, and they tend to be even larger for low emitters. Similar to the results in Table 7, emissions are predicted to grow faster for firms that invest more and firms with lower climate scores. VAR-based results also support our conclusion from Figure 8 that all of U.S. decarbonization is expected to come from the 30 largest emitters. See Section A.4 for details.

The above analysis also reveals substantial discrepancies between the emissions data from MSCI and Trucost. These discrepancies are larger for smaller emitters and firms that do not disclose their emissions. See Section A.5 for details.

5.4. Paris redux

As noted in Section 3.4, U.S. participation in the Paris Agreement has been on and off. Suppose nevertheless that the U.S. were to meet its goals as of 2021 under that agreement. Carbon emissions in 2030 would then have to be 41% lower than in 2023, declining from 6.3 to 3.7 billion tons. How does this 41% reduction compare to targets for 2030 emissions that U.S. firms have set, or to MSCI's forecasts of firms' emissions?

Panel A of Table 8 summarizes the targeted and forecasted emission reductions for the

[26]The top 10 emitters as of 2022, based on scope 1 emissions, are Exxon Mobil, Vistra, Southern, Duke Energy, Berkshire Hathaway, Chevron, American Electric Power, Nextera Energy, AES, and Entergy.

798 U.S. firms that have emission targets, as identified by MSCI. Results are shown for each of four emission categories: scope 1, scope 2, scope 3, and their sum. The median targeted reductions for scopes 1 and 2, at 33% and 32% respectively, are moderately below the 41% Paris reduction. For scope 3, however, the median targeted reduction is actually negative, at −7%, implying an increase rather than a reduction. The equal-weighted and emission-weighted average scope 3 reductions are positive but still quite low, at 12% and 8%. Recall that scope 3 captures over half of U.S. aggregate emissions not captured by firms' scopes 1 and 2. If scope 3 falls only modestly, let alone increases, the U.S. would fall well short of the 2021 Paris goal.

MSCI's forecasts of firms' emissions tell yet a worse story than firms' targets. Even for scopes 1 and 2, the medians of the forecasted reductions are just 17% and 15%, and the averages are similarly low. For scope 3, the forecasts darken the already bleak picture from the targets. In all emission categories, MSCI is rather pessimistic about firms' meeting their targeted reductions, predicting reductions often two or three times smaller than targeted.

The highest corporate emitters are pivotal in the nation's decarbonization efforts, because they account for a large fraction of aggregate emissions, as earlier noted. For example, in 2023, the top 10% of scope 1 emitters account for 96% of U.S. firms' scope 1 emissions. Most firms in the top 10% have emission targets.[27] Panel B of Table 8 repeats the analysis in Panel A for the firms in the top 10% of U.S. emitters in each category. Panel B delivers the same messages as Panel A. The emission-weighted averages are essentially identical, given the dominance of large emitters, but the medians and equal-weighted averages are also similar to Panel A, for both targets and forecasts. The overall message, reinforced by the high emitters, is that the U.S. corporate sector is far from the Paris-level trajectory that would produce the potential reductions in carbon burden estimated in Tables 2 and 3.

6. Carbon burden and expected stock returns

In this section, we relate carbon burden to the cross section of expected stock returns. Prior studies, cited in the introduction, link expected returns to firms' past carbon emissions. We instead focus on forecasts of future emissions. We find that expected returns are significantly related to forecasted emissions, even controlling for past emissions and other firm characteristics.

[27]Emission targets are much more prevalent among large emitters. Among the top 10% of emitters, 65% to 74% have targets, depending on the emission category, whereas among the other 90% of emitters in any category, fewer than 24% have targets.

We estimate each firm's expected stock return by its implied cost of capital (ICC)—the discount rate that equates the firm's market value of equity to the present value of its expected future cash flows. We compute the ICC for each stock at year-end 2023 using the method of Hou, van Dijk, and Zhang (2012), which builds on Gebhardt, Lee, and Swaminathan (2001) but replaces analysts' earnings forecasts with regression-based forecasts. This approach delivers the most precise expected return estimates in the cross section among all ICC methods evaluated by Lee, So, and Wang (2021).[28] Prior studies relate the cross section of ICCs to various measures of greenness (e.g., Chava, 2014; Pastor, Stambaugh, and Taylor, 2022), including past carbon emissions (Eskildsen et al., 2025), but, to our knowledge, we are the first to relate it to future emissions.

We compute each firm's carbon burden using a 2% discount rate and scope 1 emissions from all future years. We compute these burdens as of 2023 year-end, capitalizing emissions from 2024 forward. We divide each firm's carbon burden by its 2023 year-end market value of equity, M, and denote the ratio by CB/M, as before. We also compute C/M, the firm's ratio of year-2023 scope 1 emissions to M. In addition to CB/M and C/M, we include three 2023 year-end control variables: market beta and the logs of M and the book-to-market ratio. We estimate market betas following Fama and French (1992): each June, we sort stocks into ten portfolios by M, estimate each portfolio's beta from a time-series regression over the prior 120 months, and assign that beta to the stocks in the corresponding portfolio.

Table 9 reports results from cross-sectional regressions of firm-level ICC on CB/M, C/M, and the three controls. Our main finding is a positive and significant relation between ICC and CB/M: firms with higher carbon burdens tend to have higher costs of capital. This relation holds both with and without controls, with t-statistics ranging from 3.23 to 7.65 across specifications. Importantly, future carbon (CB/M) remains positively related to expected returns even after controlling for past carbon (C/M). This result suggests that market participants look beyond past emissions and consider firms' decarbonization plans when pricing stocks. These findings are consistent with the view that firms can lower their costs of capital by committing to credible decarbonization plans.

6.1. Economic significance

The positive relation between ICC and future carbon is also economically significant. Consider two hypothetical firms: a heavy emitter and a light one. The heavy emitter's CB/M equals the average among the top 10% of firms sorted by CB/M, and the light emitter's

[28] For details on our ICC computations, see Pastor, Stambaugh, and Taylor (2022).

CB/M is the average among the bottom 10%. Estimates in column 1 of Table 9 imply that the heavy emitter's ICC exceeds the light emitter's by 1.7% per year—a large difference. The implied difference remains substantial, 0.9% per year, even when we use the smallest estimate of the CB/M slope (from column 4 of Table 9).[29] The relation between ICC and CB/M is therefore clearly economically significant.

Another example illustrating the economic significance involves two utilities: Vistra Energy and American Electric Power (AEP). Both are among the largest emitters: in 2023, Vistra has the highest scope 1 emissions in the utility sector and the second highest in our entire sample, whereas AEP ranks fourth among utilities and sixth overall. Vistra's ratio of carbon burden to recent emissions (CB/C) is close to the median among the top-10 utility emitters, indicating a typical forecasted rate of decarbonization. Among those top-10 emitters, AEP has the lowest CB/C, implying the fastest decarbonization. If Vistra were expected to decarbonize as fast as AEP, holding constant Vistra's market cap, what would be the estimated effect on Vistra's ICC? Adjusting Vistra's carbon burden to match AEP's decarbonization rate, we estimate that Vistra's ICC would be lower by 1.7% to 4.4% per year, depending on the set of controls.[30] This is a large effect.

Repeating this calculation for each utility firm, we estimate how much ICCs would change if all utilities decarbonized at the same rate as AEP—that is, if they shared AEP's CB/C ratio. We find a value-weighted average decrease in ICC between 0.2% and 0.6% per year, depending on the controls. For utilities in the top quartile of CB/M, the average decrease in ICC is even larger, between 0.7% and 2.0%.

[29]The heavy emitter's CB/M is 14.42 and the light emitter's is 0.0005, a difference of 14.419. Multiplying 14.419 by 1.15, the slope estimate in column 1 of Table 9, and undoing the division of CB/M by 1000 before running the regression, gives $14.419 \times 1.15/1000 = 1.66\%$. Multiplying 14.419 by 0.617, the estimate in column 4, gives $14.419 \times 0.617/1000 = 0.89\%$. Also note that 14.419 is close to twice the cross-sectional standard deviation of CB/M (which is 7.15), so the heavy-light comparison can also be interpreted as a two-standard-deviation change in CB/M.

[30]Vistra's actual CB/M is 65.2. Its counterfactual CB/M is 38.2, the product of its C/M (0.0067682) and AEP's CB/C ratio (5642). Matching AEP's decarbonization rate thus reduces Vistra's CB/M by $27 = 65.2 - 38.2$. To compute the implied change in ICC, we multiply 27 by the estimated coefficient on CB/M in Table 9 and divide by 1000 (to undo the scaling of CB/M before running the regressions). Across the columns of Table 9, the estimated coefficients range from 0.617 to 1.648. Using the smallest estimate, Vistra's ICC would be lower by 1.7 percentage points per year ($0.017 = 0.617 \times 27/1000$). Using the largest estimate, the reduction would be 4.4 percentage points ($0.044 = 1.648 \times 27/1000$). Note that C in the above CB/C and C/M ratios denotes emissions in tons of carbon.

6.2. Future versus past carbon

We relate the ICC not only to future carbon but also to past carbon. In the cross-sectional regression of ICC on C/M, we find a positive and significant relation, whether or not we include the three controls (columns 2 and 5 of Table 9). This finding supports the evidence of Eskildsen et al. (2025), who also document a positive and significant cross-sectional relation between ICC and scaled recent carbon emissions. However, the ICC-C/M relation disappears once we control for CB/M (columns 3 and 6). In fact, including CB/M flips the sign of the estimated slope on C/M from significantly positive to insignificantly negative. In other words, past carbon enters positively when included alone but is driven out by the inclusion of future carbon. In contrast, future carbon retains its positive and significant coefficient even after controlling for past carbon, as discussed earlier. In the horserace between past and future carbon, future carbon wins.

Why are expected returns more closely related to future carbon than to past carbon? A potential answer is that investors require compensation for carbon-tax risk—the risk that a carbon tax, or a similar policy, will be imposed in the future. This risk is systematic, and firms with greater future emissions are more exposed to it. The risk premium could be positive or negative. On one hand, as often argued, a carbon tax can benefit investors and their descendants by mitigating climate change. On the other hand, the tax reduces corporate profits, because part of its burden falls on firms, and investors are harmed unless the redistributed tax proceeds fully compensate them for the accompanying loss of equity value. In addition, a carbon tax shifts investors' consumption toward greener alternatives that they would otherwise less desire.[31] Our results accord with the negative effects on investors dominating, generating a positive carbon-tax risk premium. Recall that firms with lower carbon burdens have lower ICCs, even holding recent carbon constant. This finding suggests that firms with more credible decarbonization plans are less exposed to carbon-tax risk and are rewarded with lower costs of capital.

Our results are less supportive of investors requiring compensation for "tort risk"—the risk of having to pay damages awarded by courts. In recent years, several heavy U.S. emitters have been sued for damages by certain states and municipalities. While tort risk is similar in spirit to carbon-tax risk, the latter is about future carbon whereas the former is more about past carbon (as it would seem frivolous to sue a company for future emissions that may or may not materialize). In that sense, our evidence is consistent with investors being more concerned about carbon-tax risk than about tort risk.

[31] See Garleanu and Pedersen (2025) for a recent theoretical analysis of the risk premium associated with a potential carbon tax.

The carbon-tax perspective also helps motivate our use of CB/M rather than alternative scalings of carbon burden. Consider an example in which a permanent carbon tax is unexpectedly imposed at a rate equal to the corresponding year's SCC, the tax is fully borne by the firm, emission forecasts remain unchanged, and the firm's cost of capital is ρ_τ from equation (1). Under these assumptions, the firm's carbon burden equals the present value of its tax liabilities. The tax therefore reduces the firm's market value by the amount of its carbon burden, generating a stock return of $-$CB/M. The CB/M ratio thus perfectly captures shareholders' return exposure to carbon-tax risk in this example.

The same intuition extends to a more general setting, which we present in Section A.2. This setting relaxes several restrictive assumptions from the above example, allowing, for instance, the carbon tax rate to differ from the SCC and emission forecasts to change following the imposition of the tax. We show that even in this setting, a firm's carbon burden is proportional to the risk of future carbon-tax payments. Specifically, an unexpected increase in the probability of a future carbon tax leads to a stock return proportional to $-$CB/M. Carbon burden captures the tax-related cash-flow risk. Scaling it by market equity, as in CB/M, translates this cash-flow risk into return risk faced by shareholders.

Carbon-tax risk is not the only possible interpretation of our evidence. Another plausible mechanism involves investor preferences (e.g., Pastor, Stambaugh, and Taylor, 2021, Pedersen, Fitzgibbons, and Pomorski, 2021). Investors may have tastes for low-carbon firms, particularly those with credible decarbonization plans and thus lower carbon burdens. Distinguishing among competing explanations is challenging, especially given that we observe only a single cross section of carbon burdens, but it is a task worthy of future research.

6.3. Robustness

Our calculation of the carbon burden relies on the EPA's forecasts of the SCC, which are subject to considerable uncertainty. The SCC forecasts, however, have no material impact on our regression results. To demonstrate this, we recompute each carbon burden by setting $SCC_\tau = 1$ for all τ in equation (1), thereby removing the SCC from the calculation. We then rerun the regressions from Table 9 using this adjusted version of CB/M.

Table 10 reports the results. As before, we find a positive and significant relation between ICC and CB/M, regardless of which controls are included. We also find a positive and significant relation between ICC and C/M, but that relation vanishes once we control for CB/M. These results are very similar to those in Table 9.

As in Table 9, the results in Table 10 are again quite economically significant. Consider the example of Vistra vs. AEP from Section 6.1. If Vistra were expected to decarbonize as quickly as AEP, its ICC would be lower by 1.6% to 4.6%, depending on the specification. Across all utilities, the value-weighted average decrease in ICC ranges from 0.2% to 0.6%, and among utilities in the top quartile by CB/M, the average decrease ranges from 0.7% to 2.0%. These magnitudes closely mirror those obtained from Table 9. Hence, the SCC plays no role in our conclusions about the pricing of future versus past carbon. Instead, the conclusions are driven by firms' differing paths of future emissions.

Our results hold even at the level of the industry. We conduct two analyses to assess the role of industry grouping. First, we repeat the analysis at the industry level, using the Fama-French 49 classification. After aggregating firm-level variables to the industry level using value-weighted averages, we rerun the regressions from Table 9 across the 49 industries. The results closely resemble those in Table 9: the coefficient on CB/M is positive and significant with or without controls, and the coefficient on C/M flips from positive to negative once CB/M is included. Second, we repeat the firm-level regressions with industry fixed effects to isolate within-industry variation. The results are similar but weaker: after controlling for C/M, the coefficient on CB/M remains positive but loses statistical significance. See Table A.8. We conclude that across-industry variation plays a larger role than within-industry variation in explaining our results.

Finally, we examine the robustness of our results to three design modifications. First, we recompute CB/M and C/M using scope 1+2 instead of scope 1 emissions. Second, we winsorize CB/M and C/M at the 99th percentiles to remove outliers. Third, we rescale both variables by firm revenue, using CB/R and C/R instead of CB/M and C/M. In all three cases, the results look very similar to those in Table 9. See Tables A.9 through A.11.

We conclude that our asset pricing results are strong and robust. Future carbon is priced in the cross section of expected stock returns, even after controlling for past carbon.

7. Conclusion

We estimate carbon burdens, novel measures of future carbon damages, for U.S. corporations. We find these burdens to be large. Based on our year-end 2023 baseline estimates, the aggregate U.S. carbon burden is $87 trillion, which equals 131% of the total value of corporate equity. Carbon burdens vary greatly across industries, from 1% of market value for a typical financial firm to 694% for a typical utility, based on direct emissions. When

indirect emissions are added, the utility's carbon burden more than doubles, but the financial firm's burden grows a thousandfold. For 13% of firms, which represent 10% of total market capitalization, their direct carbon burdens exceed their market values. For these firms, the present value of their carbon costs to society exceeds the present value of their dividends to shareholders. Firms' large carbon externalities suggest that a continued debate regarding the Friedman (1970) doctrine, according to which firms should focus solely on maximizing profits, is warranted.

We estimate that if the U.S. stuck to its 2021 goals under the Paris Agreement, its carbon burden would fall by 21% to 32%. Key to achieving those goals are the emission reductions of the largest emitters. Promisingly, the largest emitters have the most negative expected future emission growth rates, as the cross-sectional relation between current emissions and future emission growth rates is strongly negative. The relation is so strong that all of the decarbonization of the U.S. corporate sector by 2050 is expected to come from the 30 largest emitters. However, the largest emitters' emission reduction targets fall well below the country's 2021 Paris goals, even if we take those targets at face value.

Our carbon burden estimates come with a fair amount of imprecision that is hard to quantify. All three building blocks of the carbon burden—emission forecasts, forecasts of the SCC, and the discount rate—are imprecise, to an uncertain degree.[32] We consider three discount rates, but we are unable to compute standard errors because the forecasts we obtain from the MSCI, EIA, and EPA come without confidence bands. Nevertheless, in all scenarios we consider, the corporate sector's carbon burden is large.

As argued earlier, it would be naive to assign full responsibility for the aggregate carbon burden to the corporate sector, because how much carbon a country emits depends to a large extent on household demand and politics. Similarly, it is unclear how to allocate responsibility across firms, given their symbiotic relationships. For example, it would be simplistic to hold utilities fully accountable for their direct emissions, since the demand for their power comes from other sectors. Carbon burden is inherently shared, and assigning responsibility for it to individual firms is somewhat arbitrary. Determining the extent to which firms are responsible for emissions is beyond the scope of this study.

The capital market evidently cares about firms' emissions. We find that firms with larger carbon burdens are priced lower, in that their equities have higher expected returns. For example, a firm whose carbon burden divided by market cap is in the top decile has an annual expected return roughly 1% to 2% higher than a firm in the bottom decile. Moreover, we

[32]For example, Barnett et al. (2025) emphasize the central role of uncertainty in climate policy.

find that future rather than past emissions are more relevant for equity pricing.

Designing policies that reduce the aggregate carbon burden fairly, efficiently, and significantly is an important task for scholars and policymakers alike. To improve the way we discount future carbon emissions, we need more research into their risk profile. Finally, moving beyond carbon, future research should try to quantify other externalities, positive and negative, that corporations impose on society.

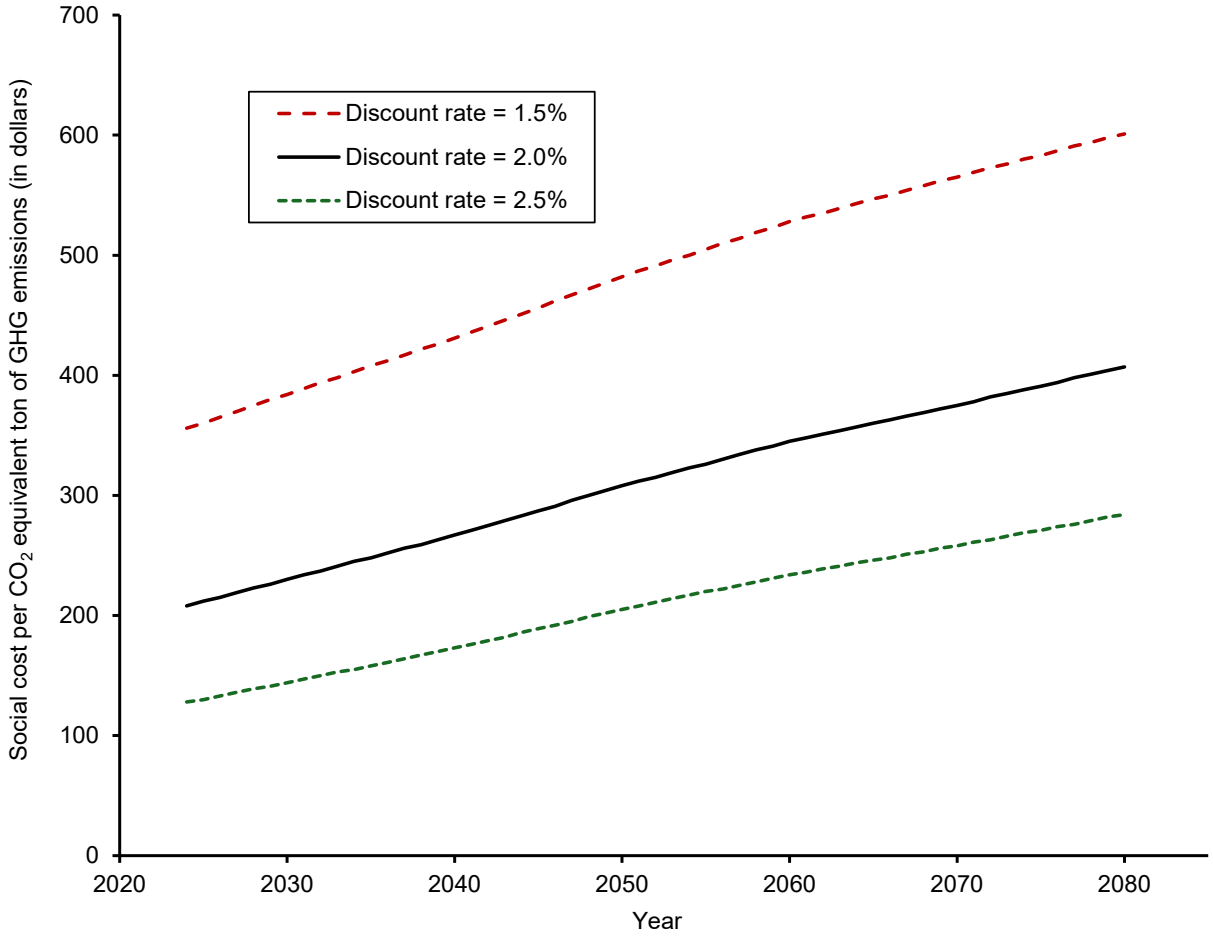


Figure 1. Social costs of GHG emissions. The figure plots EPA estimates of the social cost per CO_2-equivalent ton of GHGs emitted in a given future year. The EPA provides the costs through 2080 that are associated with each of three discount rates: 1.5% (long dashes), 2.0% (solid line), and 2.5% (short dashes).

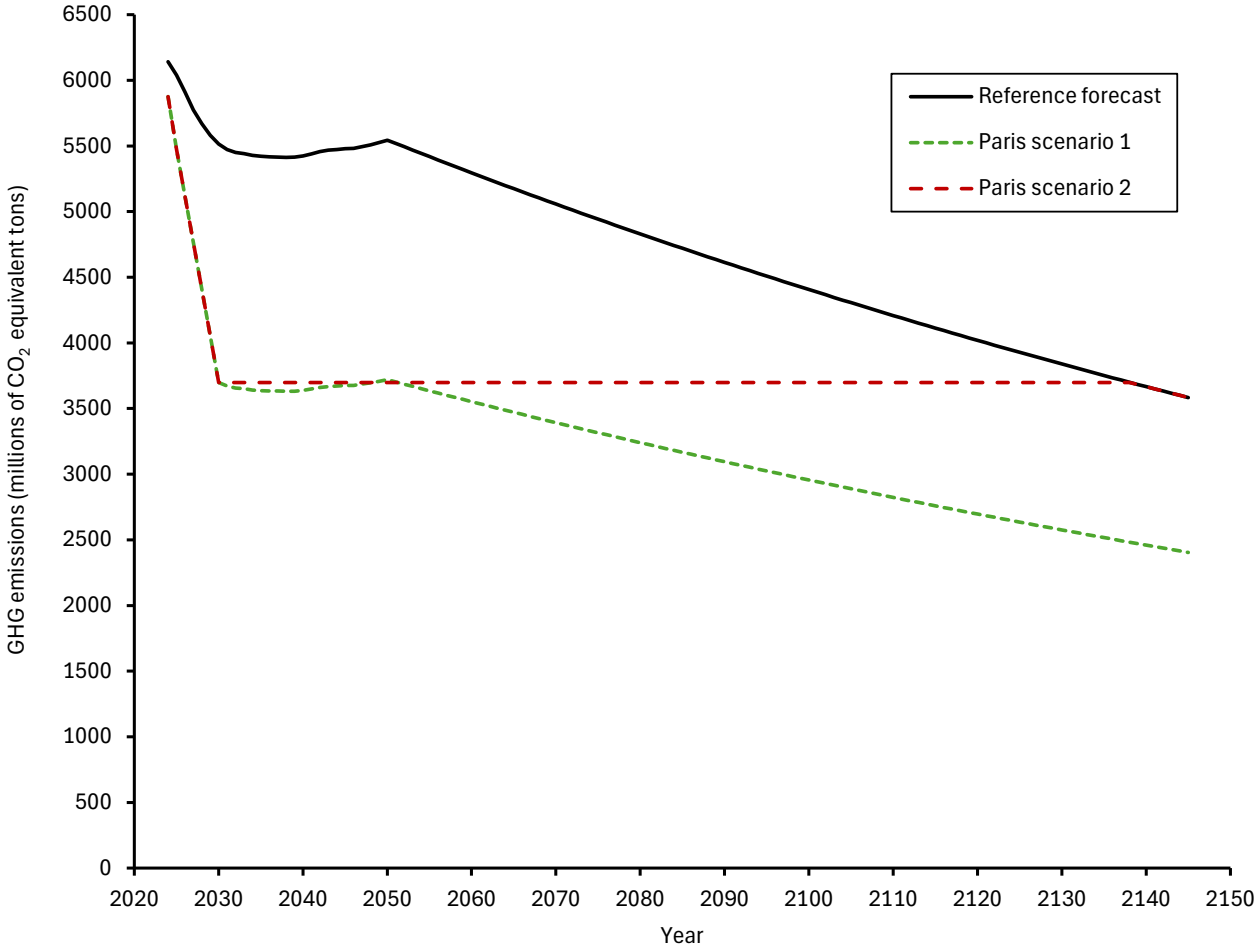


Figure 2. Forecasts of U.S. GHG emissions. The figure plots the reference forecasts (solid line) as well as forecasts under two scenarios for the Paris agreement. In the first Paris scenario (short dashes), the ratio of emissions to reference-level forecasts is maintained at the agreement's 2030 level in all later years. In the second Paris scenario (long dashes), no additional reductions relative to the reference level occur after 2030. The plot truncates the forecast time horizon, which technically extends to infinity.

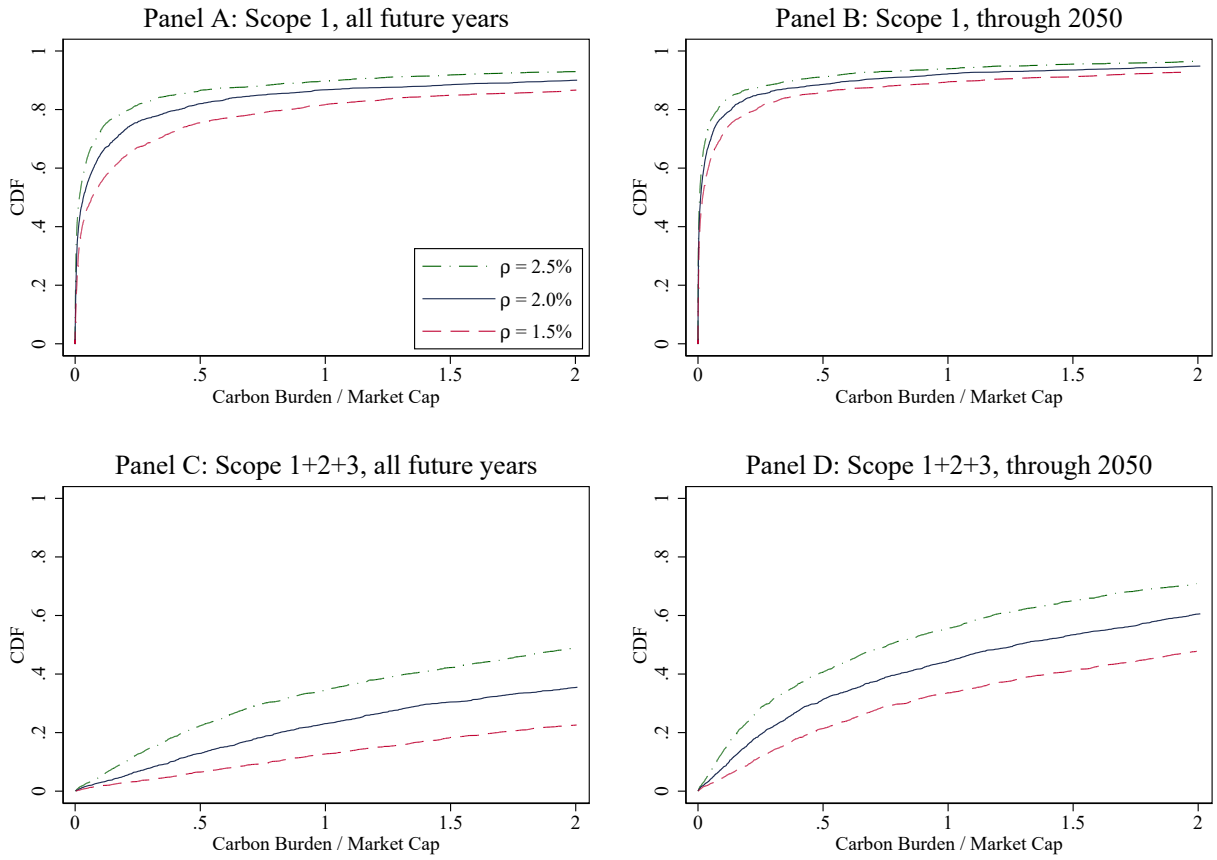


Figure 3. Distribution of firms' carbon burden as a fraction of market cap. This figure shows cumulative distribution functions (CDFs) of the ratio of carbon burden to market cap, computed in the cross section of firms in 2023. Carbon burdens are computed using MSCI's forecasts. The CDFs weight each firm equally.

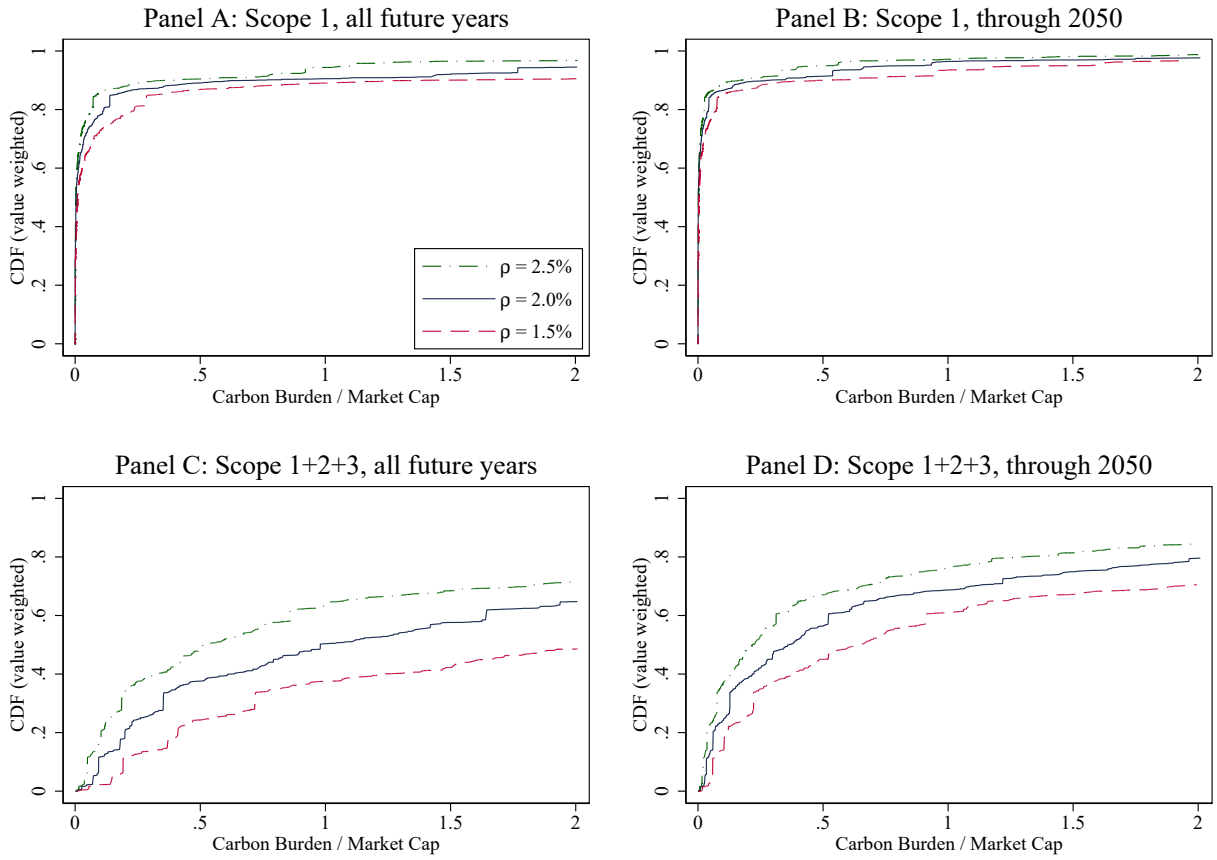


Figure 4. Value-weighted version of previous figure. Whereas the previous figure plots the fraction of firms below each x-axis value, this figure plots the fraction of aggregate market cap belonging to firms below each x-axis value.

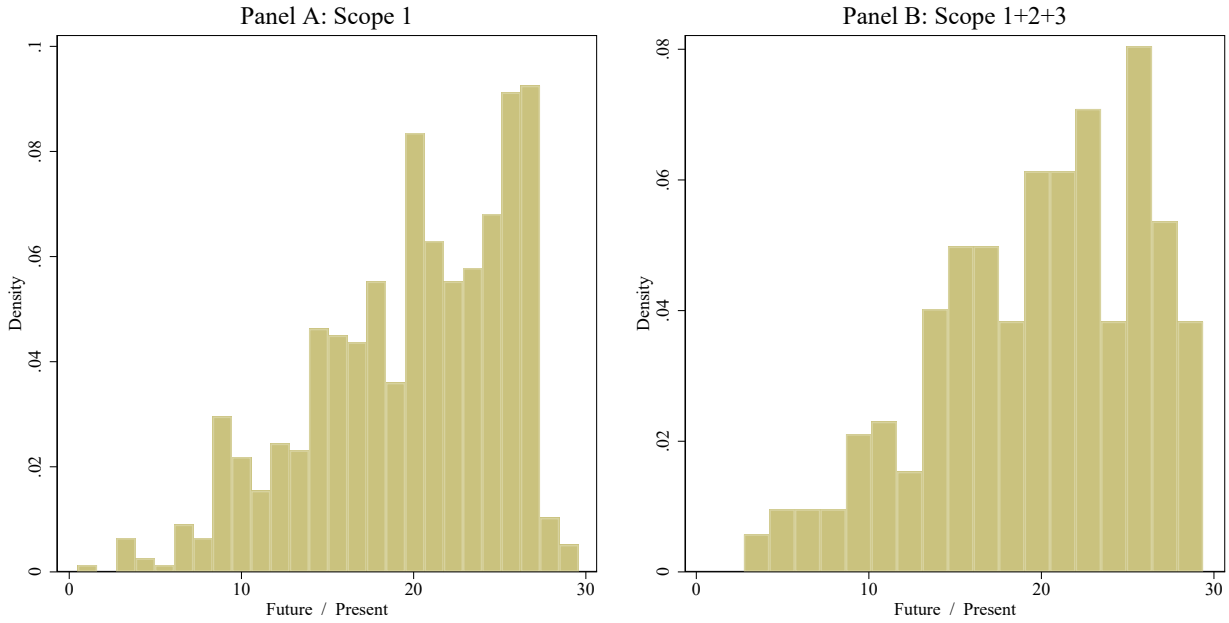


Figure 5. Firms' carbon burdens: Future years vs. current year. We compute each firm's ratio of carbon burden through 2050 to carbon burden from 2023. The figure plots this ratio's distribution across firms. Carbon burdens are computed using MSCI's emissions forecasts, with $\rho = 2\%$. Panel A (B) excludes firms with scope 1 (1, 2, or 3) growth rate equal to 1%; these excluded firms either do not have a target or have a target that MSCI deems uninformative. Panel A (B) includes 696 (353) firms in total.

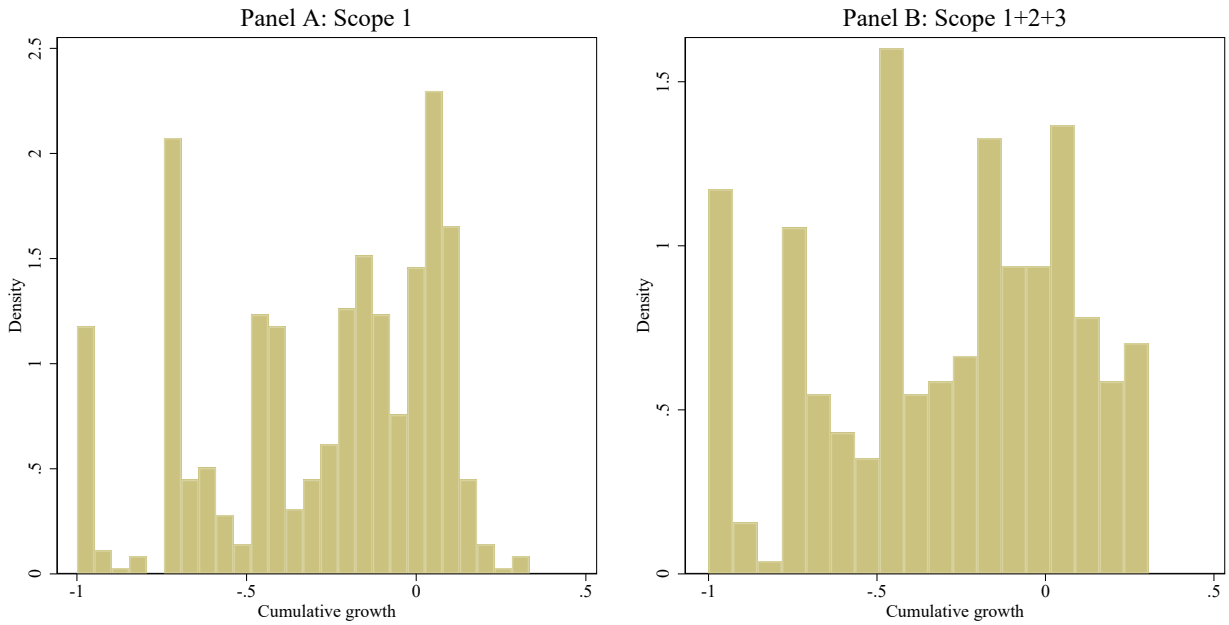


Figure 6. Firms' forecasted emissions growth. This figure plots the distribution of firms' cumulative forecasted emissions growth rates, computed as the fraction change in emissions from 2023 to 2050. Emissions forecasts are from MSCI. Panel A (B) excludes firms with scope 1 (1, 2, or 3) annual growth rate equal to 1%; these excluded firms either do not have a target or have a target that MSCI deems uninformative. Panel A (B) includes 696 (353) firms in total.

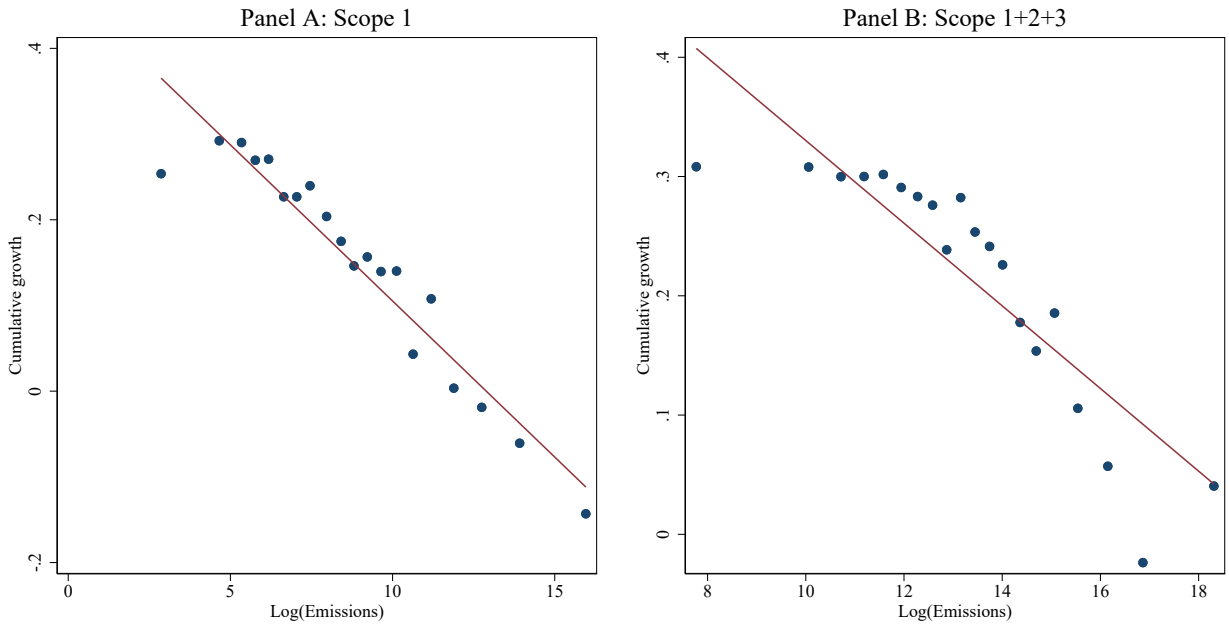


Figure 7. Current emissions and forecasted emissions growth. This figure shows the binscatter plots of firms' cumulative forecasted emissions growth rates, computed as the firm's fraction change in emissions from 2023 to 2050, against the log emissions in 2023. Emissions forecasts are from MSCI. Panel A (B) includes 2,543 (2,574) firms in total.

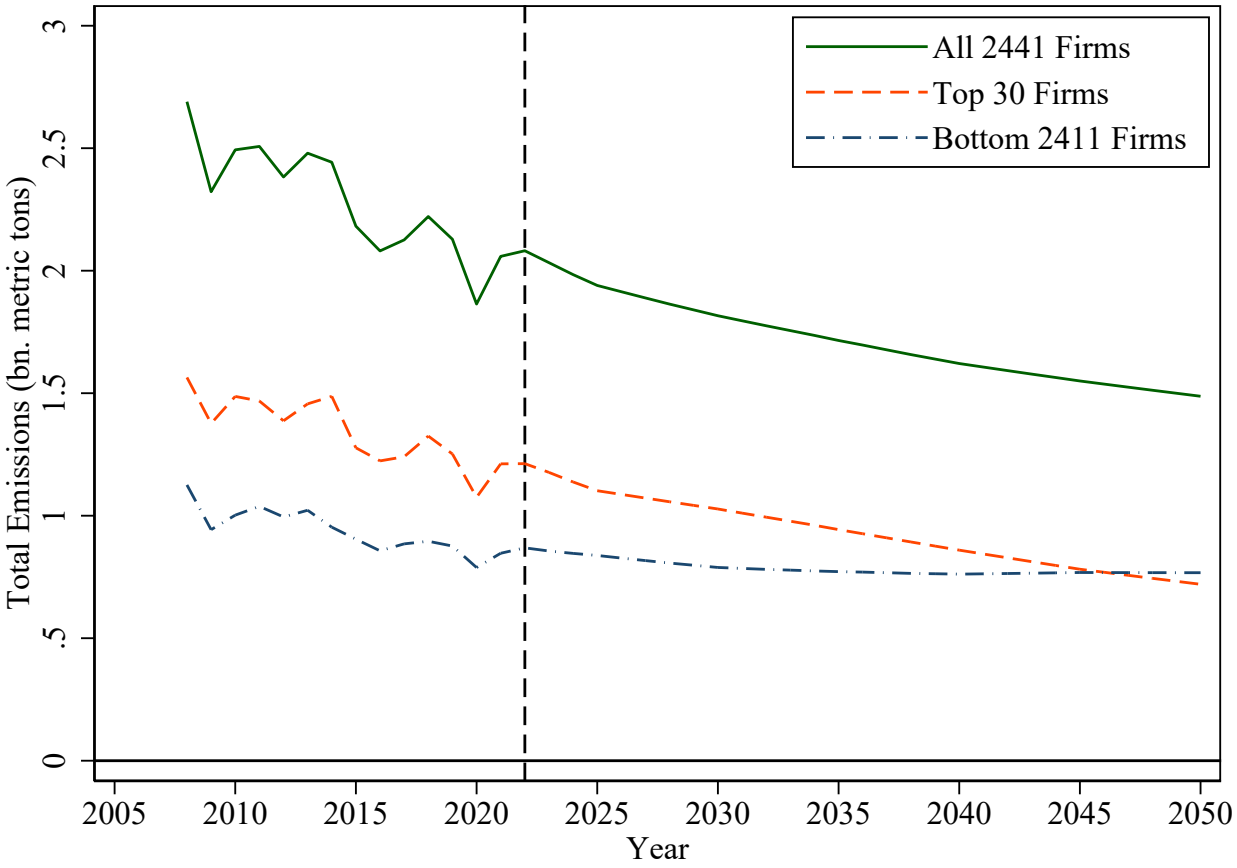


Figure 8. Past and future emissions. This figure shows past and future scope 1 emissions, using MSCI data on historical emissions and forecasts. The sample includes firms that have non-missing emission scope 1 forecasts for 2023 and historical emissions for 2022. We rank firms based on their emissions in 2022. The figure shows the sum of emissions, in billions of metric tons, each year within groups of firms ranked by their emissions in 2022. For example, "Top 30 Firms" includes the 30 firms with the highest scope 1 emissions in 2022. In years after 2022, emissions are from MSCI forecasts. In years ≤ 2022, emissions are the actual historical emissions. In years $t < 2022$, historical emissions are divided by an annual factor equal to (1) the year-2022 emissions aggregated across subsample firms with non-missing year-t emissions divided by (2) the year-2022 emissions aggregated across all subsample firms. For example, suppose 25 of the top 30 firms were operating in 2020, and these 25 firms accounted for 90% of the 30 top firms' emissions in 2022. To adjust the 2020 emissions, we divide the total emissions of these 25 firms by a factor of 0.9, which increases their year-2020 emissions by 1.111. The purpose of this adjustment is to correct for an upward trend in data coverage before 2022. Without this adjustment, we would impute zeros for missing firms' emissions, which would bias the historical emissions downward.

Table 1
Total U.S. carbon burden

Panel A shows the estimated total social costs of U.S. GHG emissions over various future periods beginning in 2024. Results are based on the reference forecasts of U.S. GHG emissions and are shown for three values of the discount rate. Panel B shows each amount in Panel A as a fraction of the total value of U.S. corporate equity at year-end 2023.

	Discount rate		
Future period	2.5%	2.0%	1.5%
Panel A. Trillions of dollars			
Through 2050	17.35	28.98	50.64
Through 2080	30.87	53.21	95.81
All future years	45.61	87.09	178.84
Panel B. Fraction of U.S. corporate equity value			
Through 2050	0.261	0.436	0.763
Through 2080	0.465	0.801	1.443
All future years	0.687	1.312	2.693

Table 2
U.S. carbon burden reductions under the Paris Agreement (Scenario 1)

Panel A shows the estimated reductions in social costs of U.S. GHG emissions under the first scenario for the Paris Agreement. In this scenario, the fraction of reference-level emissions in later years is maintained at the agreement's 2030 level. Reductions are relative to the reference forecasts of U.S. GHG emissions and are shown for three values of the discount rate and over various future periods beginning in 2024. Panel B shows each amount in Panel A as a fraction of the total value of U.S. corporate equity at year-end 2023. Panel C shows each amount in Panel A as a fraction of the corresponding U.S. carbon burden reported in Panel A of Table 1.

	Discount rate		
Future period	2.5%	2.0%	1.5%
Panel A. Trillions of dollars			
Through 2050	4.95	8.29	14.52
Through 2080	9.40	16.27	29.39
All future years	14.25	27.42	56.72
Panel B. Fraction of U.S. corporate equity value			
Through 2050	0.075	0.125	0.219
Through 2080	0.142	0.245	0.443
All future years	0.215	0.413	0.854
Panel C. Fraction of U.S. carbon burden			
Through 2050	0.285	0.286	0.287
Through 2080	0.305	0.306	0.307
All future years	0.312	0.315	0.317

Table 3
U.S. carbon burden reductions under the Paris Agreement (Scenario 2)

Panel A shows the estimated reductions in social costs of U.S. GHG emissions under the second scenario for the Paris Agreement. In this scenario, no additional reductions relative to the reference level occur after achieving the agreement's 2030 level. Reductions are relative to the reference forecasts of U.S. GHG emissions and are shown for three values of the discount rate and over various future periods beginning in 2024. Panel B shows each amount in Panel A as a fraction of the total value of U.S. corporate equity at year-end 2023. Panel C shows each amount in Panel A as a fraction of the corresponding U.S. carbon burden reported in Panel A of Table 1.

	Discount rate		
Future period	2.5%	2.0%	1.5%
Panel A. Trillions of dollars			
Through 2050	4.87	8.15	14.27
Through 2080	8.75	15.10	27.19
All future years	10.34	18.31	33.88
Panel B. Fraction of U.S. corporate equity value			
Through 2050	0.073	0.123	0.215
Through 2080	0.132	0.227	0.410
All future years	0.156	0.276	0.510
Panel C. Fraction of U.S. carbon burden			
Through 2050	0.280	0.281	0.282
Through 2080	0.284	0.284	0.284
All future years	0.227	0.210	0.189

Table 4
Carbon burden across industries

This table shows the ratios of carbon burden to market cap and of future to present carbon burden for a typical firm in each industry. Carbon burden is computed using MSCI forecasts through 2050 in Panel A and in all future years in Panel B, with $\rho = 2\%$. "Carbon Burden: Future / Present" equals the ratio of the industry's carbon burden from future years to its burden from 2023 emissions only. We use the Fama-French 12 industry classification. Industry "Other" includes Mines, Construction, Building Materials, Transportation, Hotels, Business Services, and Entertainment.

	Carbon Burden / Market Cap			Carbon Burden: Future / Present		
Industry	Scope 1	Scope 1+2	Scope 1+2+3	Scope 1	Scope 1+2	Scope 1+2+3
Panel A: Through 2050						
NoDur	0.12	0.18	2.09	22.74	22.28	22.87
Durbl	0.02	0.05	3.52	19.84	19.32	24.68
Manuf	0.26	0.37	5.56	22.33	22.41	26.22
Enrgy	1.06	1.19	20.47	22.05	22.12	28.47
Chems	0.45	0.60	3.00	22.41	22.08	25.52
BusEq	0.00	0.01	0.24	19.19	18.76	20.16
Telcm	0.01	0.07	0.66	16.54	16.26	21.40
Utils	2.70	2.80	5.53	20.35	20.46	23.61
Shops	0.04	0.08	1.88	25.27	24.29	27.79
Hlth	0.01	0.02	0.46	21.88	22.46	23.43
Money	0.00	0.01	5.28	23.51	20.51	26.63
Other	0.37	0.41	1.69	22.24	22.33	24.00
Panel B: All future years						
NoDur	0.36	0.53	6.20	68.48	66.47	67.81
Durbl	0.05	0.15	10.98	57.83	55.75	77.31
Manuf	0.73	1.08	17.42	64.09	64.89	82.33
Enrgy	2.95	3.33	65.89	61.34	61.84	92.19
Chems	1.23	1.63	9.17	60.96	59.91	78.25
BusEq	0.01	0.03	0.71	52.50	53.91	60.80
Telcm	0.04	0.20	2.04	48.96	47.56	66.46
Utils	6.94	7.24	15.89	52.95	53.59	68.66
Shops	0.14	0.26	6.03	80.01	76.12	89.53
Hlth	0.03	0.07	1.41	66.66	68.79	71.15
Money	0.01	0.04	16.60	72.01	61.97	83.82
Other	1.09	1.21	5.02	65.21	65.57	71.69

Table 5
Targets vs. forecasts

This table shows the cumulative growth rate in each industry's aggregate emissions. Cumulative growth rate is the fraction change between the industry's aggregate 2023 and 2050 emissions. Column "Forecast" uses MSCI forecasts. Column "Target" uses firms' targets. Targets are available for fewer firms than forecasts are. "Ratio" is the industry's sum of 2050 emissions targets divided by the industry's sum of 2050 emissions forecasts, using only firms for which both targets and forecasts are available.

	Scope 1			Scope 1+2			Scope 1+2+3		
Industry	Forecast	Target	Ratio	Forecast	Target	Ratio	Forecast	Target	Ratio
NoDur	-0.11	-0.47	0.50	-0.14	-0.51	0.47	-0.12	-0.45	0.54
Durbl	-0.26	-0.64	0.35	-0.29	-0.63	0.40	0.02	-0.19	0.66
Manuf	-0.19	-0.55	0.47	-0.17	-0.52	0.49	0.09	-0.17	0.60
Enrgy	-0.24	-0.83	0.16	-0.23	-0.82	0.17	0.24	0.15	0.70
Chems	-0.25	-0.89	0.12	-0.26	-0.87	0.14	0.03	-0.35	0.53
BusEq	-0.35	-0.80	0.28	-0.32	-0.78	0.28	-0.21	-0.51	0.55
Telcm	-0.37	-0.80	0.29	-0.39	-0.74	0.38	-0.12	-0.37	0.61
Utils	-0.37	-0.92	0.11	-0.36	-0.91	0.11	-0.13	-0.58	0.38
Shops	0.07	-0.76	0.12	0.01	-0.73	0.15	0.20	-0.15	0.39
Hlth	-0.13	-0.61	0.28	-0.10	-0.68	0.20	-0.07	-0.31	0.57
Money	-0.06	-0.60	0.24	-0.19	-0.74	0.22	0.11	-0.21	0.51
Other	-0.17	-0.80	0.16	-0.16	-0.79	0.16	-0.07	-0.66	0.25

Table 6
Explaining variation in firms' carbon metrics

This table reports adjusted R-squared values from cross-sectional regressions of carbon-burden metrics (denoted in column headers) on industry fixed effects and/or log of market cap. "CB" represents the carbon burden from all future years. "M" denotes market capitalization. "Future / Present" refers to carbon burden from all future years divided by the burden from 2023 emissions only. Carbon burdens are computed from MSCI emissions forecasts as of the end of 2023, with $\rho = 2\%$. Firms are classified into Fama-French 49 (12) industries in Panel A (B).

	Dependent Variable (log)					
Scope	CB		CB/M		Future/Present	
Panel A: Using the Fama-French 49 industries						
Scope 1	0.521	0.643	0.522	0.579	0.063	0.243
Scope 1+2	0.432	0.591	0.415	0.504	0.061	0.242
Scope 1+2+3	0.367	0.655	0.393	0.458	0.051	0.155
Panel B: Using the Fama-French 12 industries						
Scope 1	0.412	0.550	0.419	0.467	0.035	0.216
Scope 1+2	0.330	0.507	0.320	0.399	0.035	0.217
Scope 1+2+3	0.272	0.589	0.294	0.350	0.028	0.131
Industry FEs	Y	Y	Y	Y	Y	Y
Log(M)		Y		Y		Y

Table 7
Firm characteristics and forecasted emissions growth

This table shows estimates from cross-sectional regressions with dependent variable equal to the annualized emission growth rate of a firm's emissions from 2023 to 2050, based on MSCI forecast data. We set the growth rate to zero for firms with 2023 and 2050 emissions equal to zero. For other firms with 2050 emissions equal to zero, we set 2050 emissions to 1% of the 2023 emissions level so that we can compute an annualized growth rate. All regressors are measured at the end of 2023. B/M is the book-to-market ratio. Investment is the one-year fraction change in book assets. Climate Score is computed from MSCI's ESG ratings and is defined as $-(10 - Climate_score_{i,t-1}) \times Climate_weight_{i,t-1}/100$, similar to Pástor, Stambaugh, and Taylor (2022). *Climate_score* is "Climate Change Theme Score," a number between zero and 10 measuring a company's resilience to long-term risks related to climate change. *Climate_weight* is "Climate Change Theme Weight," a number between zero and 100 measuring the importance of climate change relative to other ESG issues in the company's industry. Revenue Growth is the one-year fraction change in revenue. B/M, Investment, and Revenue Growth are winsorized at the 1st and 99th percentiles. The bottom rows specify the emissions scope considered and whether we include fixed effects for Fama-French 12 industries. In parentheses, we report t-statistics clustered by industry. We multiply slope coefficients by 1,000.

	(1)	(2)	(3)	(4)
Log(Emissions)	-2.343	-2.959	-2.688	-3.211
	(-9.52)	(-12.33)	(-8.87)	(-12.53)
B/M	1.246	1.320	1.229	1.240
	(1.73)	(1.84)	(1.70)	(1.82)
Investment	2.077	2.768	2.302	2.921
	(1.27)	(1.79)	(1.59)	(2.09)
Climate Score	-7.327	-7.050	-9.193	-8.850
	(-3.39)	(-3.30)	(-3.40)	(-3.26)
Revenue Growth	-0.547	-1.007	-1.009	-1.372
	(-1.29)	(-1.86)	(-2.33)	(-2.49)
Constant	0.020	0.029	0.017	0.025
	(7.35)	(11.08)	(4.88)	(8.33)
Observations	2191	2213	2191	2213
Adjusted R^2	0.100	0.107	0.118	0.122
Scopes	1	1+2	1	1+2
Industry FE			Y	Y

Table 8
Firms' targeted and forecasted emission reductions

The table reports firms' targeted and forecasted percentage emission reductions through 2030, the horizon aligned with the Paris agreement. The reductions are stated relative to 2023 as the base year. For each emission category, Panel A includes all firms identified by MSCI as reporting an emission target, while Panel B includes only those firms that are also among the category's top 10% of emitters in 2023. The forecasted reductions reflect MSCI's emission forecasts for each firm, computed as in equation (2). The emission-weighted averages weight each firm's percentage reduction by the firm's 2023 emissions.

	Emission categories			
	Scope 1	Scope 2	Scope 3	All
Panel A: All firms with emission targets				
Number of firms	798	798	798	798
Median targeted reduction	33	32	-7	4
Equal-weighted average targeted reduction	38	36	12	16
Emission-weighted average targeted reduction	28	33	8	11
Median forecasted reduction	17	15	-7	-1
Equal-weighted average forecasted reduction	21	19	5	7
Emission-weighted average forecasted reduction	14	20	2	4
Panel B: Above firms in each category's top 10% of 2023 emitters				
Number of firms	184	211	190	200
Median targeted reduction	26	31	3	11
Equal-weighted average targeted reduction	31	36	16	18
Emission-weighted average targeted reduction	28	33	8	10
Median forecasted reduction	11	19	-1	3
Equal-weighted average forecasted reduction	16	21	8	9
Emission-weighted average forecasted reduction	14	20	2	3

Table 9
ICCs and carbon measures

This table shows estimates from cross-sectional regressions with dependent variable equal to the firm's implied cost of capital. All variables are measured at the end of 2023. CB/M denotes Carbon Burden divided by M, the market value of equity. Carbon Burden is calculated using a 2% discount rate and includes scope 1 emissions from all future years. C/M is scope 1 emissions divided by M. We estimate stocks' market betas following Fama and French (1992): at the end of every June, we assign stocks to ten portfolios based on market equity, estimate the portfolios' betas from time-series regressions using a 120-month window, then set each stock's market beta to that of its corresponding portfolio. B/M is the book-to-market ratio. In the regressions, CB/M is divided by 1000. t-statistics are reported in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
CB/M	1.150		1.648	0.617		1.346
	(7.65)		(3.31)	(4.94)		(3.23)
C/M		17.188	-8.573		8.555	-12.571
		(6.96)	(-1.05)		(4.15)	(-1.83)
Market Beta				0.059	0.060	0.059
				(4.95)	(5.00)	(4.88)
log(B/M)				0.027	0.027	0.027
				(25.79)	(25.78)	(25.87)
log(M)				0.004	0.004	0.004
				(4.23)	(4.18)	(4.32)
Constant	0.074	0.074	0.074	-0.008	-0.009	-0.008
	(67.20)	(66.95)	(67.12)	(-0.38)	(-0.39)	(-0.36)
Observations	1990	1990	1990	1925	1925	1925
Adjusted R^2	0.028	0.023	0.028	0.314	0.312	0.315

Table 10
Version of previous table with SCC = 1

In this table, we repeat the analysis from Table 9, except we compute firms' carbon burdens while setting SCC = 1, and we do not divide CB/M by 1000.

	(1)	(2)	(3)	(4)	(5)	(6)
CB/M	0.408		0.626	0.218		0.518
	(7.60)		(3.23)	(4.87)		(3.19)
C/M		17.188	-10.418		8.555	-14.370
		(6.96)	(-1.17)		(4.15)	(-1.92)
Market Beta				0.060	0.060	0.059
				(4.95)	(5.00)	(4.89)
log(B/M)				0.027	0.027	0.027
				(25.79)	(25.78)	(25.87)
log(M)				0.004	0.004	0.004
				(4.23)	(4.18)	(4.32)
Constant	0.074	0.074	0.074	-0.008	-0.009	-0.008
	(67.17)	(66.95)	(67.11)	(-0.38)	(-0.39)	(-0.36)
Observations	1990	1990	1990	1925	1925	1925
Adjusted R^2	0.028	0.023	0.028	0.314	0.312	0.315

REFERENCES

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr, 2016, Transition to clean technology, *Journal of Political Economy* 124, 52–104.

Aghion, Philippe, Antoine Dechezlepretre, David Hemous, Ralf Martin, and John Van Reenen, 2016, Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry, *Journal of Political Economy* 124, 1–51.

Aitken, Brian J., and Ann E. Harrison, 1999, Do domestic firms benefit from direct foreign investment? Evidence from Venezuela, *American Economic Review* 89, 605–618.

Alcott, Hunt, Giovanni Montanari, Bora Ozaltun, and Brandon Tan, 2025, An economic view of corporate social impact, *Journal of Finance*, forthcoming.

Aswani, Jitendra, Aneesh Raghunandan, and Shiva Rajgopal, 2024, Are carbon emissions associated with stock returns? *Review of Finance* 28, 75–106.

Atta-Darkua, Vaska, Simon Glossner, Philipp Krueger, and Pedro Matos, 2023, Decarbonizing institutional investor portfolios: Helping to green the planet or just greening your portfolio? Working paper.

Audretsch, David B., and Maryann P. Feldman, 1996, R&D spillovers and the geography of innovation and production, *American Economic Review* 86, 630–640.

Barnett, Michael, William A. Brock, Lars Peter Hansen, and Hong Zhang, 2025, Uncertainty, social valuation, and climate change policy, Working paper.

Bartram, Sohnke M., Kewei Hou, and Sehoon Kim, 2022, Real effects of climate policy: Financial constraints and spillovers, *Journal of Financial Economics* 143, 668–696.

Berg, Tobias, Lin Ma, and Daniel Streitz, 2024, Out of sight, out of mind: Divestments and the global reallocation of pollutive assets, Working paper.

Bernstein, Shai, Emanuele Colonnelli, Xavier Giroud, and Benjamin Iverson, 2019, Bankruptcy spillovers, *Journal of Financial Economics* 133, 608–633.

van Binsbergen, Jules H., and Andreas Brogger, 2022, The future of emissions, Working paper.

Blalock, Garrick, and Paul J. Gertler, 2008, Welfare gains from foreign direct investment through technology transfer to local suppliers, *Journal of International Economics* 74, 402–421.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, 2013, Identifying technology spillovers and product market rivalry, *Econometrica* 81, 1347–1393.

Board of Governors of the Federal Reserve System, 2024, *Financial Accounts of the United States*.

Bolton, Patrick, Marc Eskildsen, and Marcin Kacperczyk, 2024, Carbon home bias, Working paper.

Bolton, Patrick and Marcin Kacperczyk, 2021, Do investors care about carbon risk? *Journal of Financial Economics* 142, 517–549.

Bolton, Patrick and Marcin Kacperczyk, 2023, Global pricing of carbon-transition risk, *Journal of Finance* 78, 3677–3754.

Bolton, Patrick and Marcin Kacperczyk, 2023, Firm commitments, Working paper.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer, 2013, Are consumers myopic? Evidence from new and used car purchases, *American Economic Review* 103, 220–256.

Callahan, Christopher W., and Justin S. Mankin, 2025, Carbon majors and the scientific case for climate liability, *Nature* 640, 893–901.

Cenedese, Gino, Shangqi Han, and Marcin Kacperczyk, 2023, Carbon-transition risk and net-zero portfolios, Working paper.

Chava, Sudheer, 2014, Environmental externalities and cost of capital, *Management Science* 60, 2223–2247.

Chittaro, Lautaro, Monika Piazzesi, Marcelo Sena, and Martin Schneider, 2025, Asset returns as carbon taxes, Working paper.

Cohen, Lauren, Umit G. Gurun, and Quoc H. Nguyen, 2024, The ESG-innovation disconnect: Evidence from green patenting, Working paper.

Drupp, M.A., M.C. Freeman, B. Groom, and F. Nesje, 2018, Discounting disentangled, *American Economic Journal: Economic Policy* 10, 109–34.

Eskildsen, Marc, Markus Ibert, Theis Ingerslev Jensen, and Lasse Heje Pedersen, In search of the true greenium, Working paper, 2025.

Fama, Eugene F. and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.

Friedman, Milton, 1970, The social responsibility of business is to increase its profits, *The New York Times*.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan, 2016, Market-based emissions regulation and industry dynamics, *Journal of Political Economy* 124, 249–302.

Garleanu, Nicolae, and Lasse Heje Pedersen, 2025, Climate risk pricing, Working paper.

Gebhardt, William R., Charles MC Lee, and Bhaskaran Swaminathan, 2001, Toward an implied cost of capital, *Journal of Accounting Research* 39, 135–176.

Giglio, Stefano, Bryan Kelly, and Johannes Stroebel, 2021, Climate finance, *Annual Review of Financial Economics* 13, 15–36.

Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber, 2021, Climate change and long-run discount rates: Evidence from real estate, *Review of Financial Studies* 34, 3527–3571.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel, 2015, Very long-run discount rates, *Quarterly Journal of Economics* 130, 1–53.

Graff Zivin, Joshua, and Matthew Neidell, 2012, The impact of pollution on worker productivity, *American Economic Review* 102, 3652–3673.

Greenstone, Michael, 2002, The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 Clean Air Act amendments and the census of manufactures, *Journal of Political Economy* 110, 1175–1219.

Greenstone, Michael, Christian Leuz, and Patricia Breuer, 2023, Mandatory disclosure would reveal corporate carbon damages, *Science* 381, 837–840.

Hanna, Rema, and Paulina Oliva, 2015, The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City, *Journal of Public Economics* 122, 68–79.

Hart, Oliver, and Luigi Zingales, 2017, Companies should maximize shareholder welfare not market value, *Journal of Law, Finance, and Accounting* 2, 247-–274.

Hartzmark, Samuel M., and Kelly Shue, 2022, Counterproductive sustainable investing: The impact elasticity of brown and green firms, Working paper.

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

Howard, P.H. and D. Sylvan, 2020, Wisdom of the experts: using survey responses to address positive and normative uncertainties in climate-economic models, *Climatic Change* 162, 213–232.

Jaffe, Adam B., 1986, Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value, *American Economic Review* 76, 984–1001.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson, 1993, Geographic localization of knowledge spillovers as evidenced by patent citations, *Quarterly Journal of Economics* 108, 577–598.

Javorcik, Beata S., 2004, Does foreign direct investment increase the productivity of domestic firms? In search of spillovers through backward linkages, *American Economic Review* 94, 605–627.

Kacperczyk, Marcin, 2026, What is the carbon premium a premium on? Working paper.

Lee, Charles M.C., Eric C. So, and Charles C.Y. Wang, 2021, Evaluating firm-level expected-return proxies: Implications for estimating treatment effects, *Review of Financial Studies* 34, 1907–1951.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2021, Sustainable investing in equilibrium, *Journal of Financial Economics* 142, 550–571.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2022, Dissecting green returns, *Journal of Financial Economics* 146, 403–424.

Pedersen, Lasse Heje, 2025, Carbon pricing versus green finance, *Journal of Finance*, forthcoming.

Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski, 2021, Responsible investing: The ESG-efficient frontier, *Journal of Financial Economics* 142, 572–597.

Pellegrino, Bruno, 2025, Product differentiation and oligopoly: A network approach, *American Economic Review* 115, 1170–1225.

Pigou, Arthur C., 1920, *The Economics of Welfare*, Macmillan.

Ramadorai, Tarun, and Federica Zeni, 2024, Climate regulation and emissions abatement: Theory and evidence from firms' disclosures, *Management Science*, forthcoming.

Schoenmaker, Dirk, and Willem Schramade, 2023, *Corporate finance for long-term value*, Springer Nature.

Stern, Nicholas, 2006, *Stern review on the economics of climate change*. London, UK: Her Majesty's Treasury.

Stroebel, Johannes, and Jeffrey Wurgler, 2021, What do you think about climate finance? *Journal of Financial Economics* 142, 487–498.

Tol, Richard S.J., 2023, Social cost of carbon estimates have increased over time, *Nature Climate Change* 13, 532–536.

Tol, Richard S.J., 2025, Meta-emulation: An application to the social cost of carbon, *Energy Economics* 152: 109026.

U.S. Energy Information Administration, 2023a, *Annual Energy Outlook*.

U.S. Energy Information Administration, 2023b, *The National Energy Modeling System: An Overview*.

U.S. Environmental Protection Agency, 2019, *Global Non-CO2 Greenhouse Gas Emission Projections & Marginal Abatement Cost Analysis: Methodology Documentation*.

U.S. Environmental Protection Agency, 2023, *Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances*.

U.S. Environmental Protection Agency, 2024, *Inventory of U.S. Greenhouse Gas Emissions and Sinks, 1990–2022*.

Walker, W. Reed, 2013, The transitional costs of sectoral reallocation: Evidence from the Clean Air Act and the workforce, *Quarterly Journal of Economics* 128, 1787–1835.

Zhang, Shaojun, 2024, Carbon returns across the globe, *Journal of Finance*, forthcoming.

Appendix

A.1. Reconciling ratios

Here we explain why the modest ratio of annual emission damages to GDP is consistent with the large ratio of carbon burden to equity market value. We first provide the basic intuition in a simple no-growth framework. We then present a somewhat richer framework that models production in a traditional way and endogenizes the profit margin.

A.1.1. Simple framework

Suppose the corporate sector produces output whose value in period t is given by

$$Y_t = Y + \epsilon_t \,, \tag{A.1}$$

where Y is expected output and ϵ_t is a zero-mean random component, which makes corporate ownership risky. Producing output generates, as a by-product, a negative externality whose value is the fraction f of expected output in each period:

$$\mathcal{E}_t = fY \,. \tag{A.2}$$

Since there is no growth, the present value of all future externalities in perpetuity—the carbon burden—is simply equal to

$$\text{CB} = \frac{fY}{r} \,, \tag{A.3}$$

where r is the riskless rate. The corporate sector's dividends, equal to net profit (consistent with no growth), are given by a constant fraction of output:

$$D_t = hY_t \,, \tag{A.4}$$

where h denotes the profit margin. The market value of the corporate sector is the present value of all expected future dividends, discounted at the cost of capital r_S, which is equal to r plus a risk premium that reflects the risk in ϵ_t:

$$M = \frac{D}{r_S} \,, \tag{A.5}$$

where $D = hY$ is the expected dividend in each period. Combining equations (A.3) and (A.5), the ratio of the carbon burden to market value is

$$\frac{\text{CB}}{M} = f \, \frac{1}{h} \, \frac{r_S}{r} \,. \tag{A.6}$$

This equation helps us understand how CB/M can be large even when f is small. There are two reasons. First, the profit margin, h, is much smaller than one, making $1/h$ large. For

example, in 2023, the net profit margin of the U.S. corporate sector was about 10%, resulting in $1/h = 10$.[33] Second, the corporate cost of capital exceeds the riskless rate due to a risk premium, so that $r_S/r > 1$. For example, suppose $r = 2\%$, which is the baseline value in our empirical analysis, and $r_S = 6\%$, whose reciprocal, 16.7, is close to the historical average price-earnings ratio. We then obtain $r_S/r = 3$. Plugging these values into equation (A.6) along with $f = 4.7\%$, the ratio of annual emission damage to GDP calculated in Section 3.2, we obtain $\text{CB}/M = 1.41$. That is, the carbon burden in this example is equal to 141% of equity value, which is not far off our baseline estimate of 131% in Table 1.

A.1.2. Endogenizing the profit margin

In the above framework, production is not modeled explicitly and the corporate profit margin is specified exogenously. In this section, we present a somewhat richer framework that models production in a traditional way and endogenizes the profit margin. The model remains very simple, with no growth, no frictions, and a single consumption good. As above, total output is given by

$$Y_t = Y + \epsilon_t \,, \tag{A.7}$$

but here we model expected output Y explicitly as

$$Y = K^\alpha L^{1-\alpha} \,, \tag{A.8}$$

where K is capital, L is labor, and $0 < \alpha < 1$. Denoting the marginal products of capital and labor by r_K and w, respectively, we have $Y = r_K K + wL$. The share of capital is $r_K K/Y = \alpha$ and the labor share is $wL/Y = 1 - \alpha$. A by-product of production is an externality that reduces the utility value of consumption by fraction f of expected output in each period:

$$\mathcal{E}_t = fY \,. \tag{A.9}$$

Denoting the real riskless rate by r and recognizing that there is no growth, the present value of all future externalities in perpetuity—the carbon burden—is simply equal to

$$\text{CB} = \frac{fY}{r} \,. \tag{A.10}$$

Capital evolves as $K_{t+1} = (1 - \delta)K_t + I_t$, where δ is a positive depreciation rate and I_t is investment. We assume $I_t = \delta K_t$, so that $K_t = K$ and $I_t = I$ for all t (no growth). The investment is made by the corporate sector from its gross capital revenue, $r_K K + \epsilon_t$. Therefore, the owners of capital receive dividends D_t equal to $r_K K - I + \epsilon_t$, so that

$$D_t = (r_K - \delta)K + \epsilon_t \,. \tag{A.11}$$

[33] Aggregate U.S. after-tax corporate profits in 2023 are \$2.673 trillion, according to Table 9 in the March 28, 2024 news release from the Bureau of Economic Analysis. Dividing this figure by U.S. GDP of \$27.4 trillion yields 0.098, or approximately 10%.

The expected dividend in each period is $D = (r_K - \delta)K$. The market value of the corporate sector is the present value of all expected future dividends:

$$M = \frac{D}{r_S} = \frac{r_K - \delta}{r_S} K \,, \tag{A.12}$$

where r_S is the cost of capital, which differs from r by a risk premium reflecting the risk embeded in ϵ_t. In a frictionless economy, a unit of capital can be costlessly transformed into a unit of the consumption good. Capital stock adjusts so that Tobin's Q is equal to 1 and $M = K$. This condition and equation (A.12) pin down the marginal product of capital:

$$r_K = \delta + r_S \,. \tag{A.13}$$

Given equations (A.10), (A.12), and (A.13), the ratio of the carbon burden to market value is

$$\frac{\text{CB}}{M} = \frac{fY}{rK} = \frac{Y}{r_K K} \frac{r_K}{r} f = \frac{1}{\alpha} \frac{\delta + r_S}{r} f \,. \tag{A.14}$$

This equation helps us understand how CB/M can be large even when f is small. The first term on the right-hand side of equation (A.14) is the inverse of the capital share of GDP, so its value is close to 3. The second term, $(\delta + r_S)/r$, is also always greater than 1, because $\delta > 0$ (positive depreciation) and $r_S > r$ (positive risk premium). Using the same values of $r_S = 6\%$ and $r = 2\%$ as above and choosing a round value of $\delta = 10\%$, the value of the second term is 8, and equation (A.14) then implies $\text{CB}/M = 24f$. When again $f = 4.7\%$, we have $\text{CB}/M = 113\%$. If we increase δ slightly to 12%, we obtain $\text{CB}/M = 127\%$, which is close to our baseline estimated value of 131%.

Even though this model is richer than the simpler framework, it remains too simple for full calibration. A proper calibration would require adding realistic features such as economic growth, gradual decarbonization, debt financing, and frictions. In a model with all these features, the intuition would be far less transparent. In contrast, our equation (A.14) makes it clear that CB/M is much larger than f, for three reasons.

First, the capital share of GDP, α, is smaller than 1 (historically about one third). M is the market value of capital, whereas the externality underlying the CB is proportional to all of GDP, including the labor component. The lower the capital share, the larger the CB/M ratio, holding f constant.

Second, $\delta > 0$. Maintaining the level of output requires investment to offset the depreciation of capital. Investment is financed by capital owners from gross capital revenue, which reduces dividends (see equation (A.11)), which in turn reduces the dividends' present value, M. In other words, keeping the expected level of output (and externality) constant requires ongoing dividend reductions, which reduce M relative to CB. The larger the depreciation rate, the larger the CB/M ratio, holding f constant.

Third, the corporate cost of capital exceeds the riskless rate due to a risk premium, so that $r_S/r > 1$. The larger this ratio, the larger the CB/M value, holding f constant.

Overall, the insights we obtain here are similar to those presented in the simpler framework. The first two reasons presented above are related to the corporate profit margin, whose expected value we endogenize here as $D/Y = r_S K/Y = \alpha r_S/(\delta + r_S)$.

A.2. Carbon-tax risk and carbon burden

Consider the potential levying of a tax on emissions in every future year. If such a tax is levied, assume the tax rate in year $t+k$ is

$$\tau_{t+k} = \delta \, SCC_{t+k}, \tag{A.15}$$

meaning the tax paid in year $t+k$ by a firm is τ_{t+k} times the firm's carbon emissions in that year. We assume that the firm can always pay this tax and that δ is positive and constant across years, for simplicity.[34]

Let θ denote the market's assessed probability in year t of the above tax being levied, and let $C_{i,t+k}$ denote the firm's expected emissions in the absence of a carbon tax. Conditional on the tax being levied, the firm's expected emissions are assumed to be

$$\tilde{C}_{i,t+k} = (1-b)C_{i,t+k} \tag{A.16}$$

with $0 \leq b < 1$. In other words, the firm is expected to emit less if the carbon tax is imposed. We specify the functional form in (A.16) to be the same across firms for simplicity. Firm i's unconditional expected emissions in year $t+k$ are then given by

$$\begin{aligned} \bar{C}_{i,t+k}(\theta) &= \theta \tilde{C}_{i,t+k} + (1-\theta) C_{i,t+k} \\ &= (1 - b\theta) C_{i,t+k} \ . \end{aligned} \tag{A.17}$$

A firm's expected carbon-tax payment in year $t+k$ is the probability of the tax, θ, times the expected tax payment conditional on the tax being levied, $\tau_{t+k}\tilde{C}_{i,t+k}$. This product is given by

$$\theta \tau_{t+k} \tilde{C}_{i,t+k} = \tau_{t+k} \theta \left(\frac{1-b}{1-b\theta} \right) \bar{C}_{i,t+k}(\theta), \tag{A.18}$$

using equations (A.16) and (A.17). Multiplying that quantity by $(1 + r_{i,t+k})^{-k}$ gives the present value in year t of that expected payment, where $r_{i,t+k}$ is a k-period discount rate that reflects any premium for risk in the tax payment. The risk includes whether the tax will be levied and, if so, what firm i's emissions will then be. We assume for simplicity that

$$r_{i,t+k} = \rho_k, \tag{A.19}$$

[34] It is likely that $\delta < 1$, for at least two reasons. First, the value of SCC_{t+k} that we use is estimated by the EPA under a scenario of essentially no future carbon taxes. A carbon tax would likely reduce emissions, thereby lowering the social cost of carbon in the presence of the tax, given a social damage function that is convex in emissions. A socially optimal tax would equate the tax rate to the social cost of carbon in the presence of the tax. Second, political frictions, which in part explain the current absence of a tax, could also constrain a future tax to be lower than is socially optimal.

a common discount rate across firms. The present value of firm i's future carbon taxes is then

$$\begin{aligned} A_{i,t} &= \sum_{k=1}^{\infty}(1+\rho_k)^{-k}\theta\tau_{t+k}\tilde{C}_{i,t+k} \\ &= \sum_{k=1}^{\infty}(1+\rho_k)^{-k}\theta\left(\delta\, SCC_{t+k}\right)\left[\left(\frac{1-b}{1-b\theta}\right)\bar{C}_{i,t+k}(\theta)\right] \\ &= \delta\theta\left(\frac{1-b}{1-b\theta}\right)\sum_{k=1}^{\infty}(1+\rho_k)^{-k}SCC_{t+k}\bar{C}_{i,t+k}(\theta) \\ &= \delta\theta\left(\frac{1-b}{1-b\theta}\right)B_{i,t}(\theta), \end{aligned} \tag{A.20}$$

where $B_{i,t}(\theta)$ is firm i's carbon burden, following the definition in equation (1). Differentiating A_{it} with respect to θ gives[35]

$$\frac{dA_{i,t}}{d\theta} = \delta\left(\frac{1-b}{1-b\theta}\right)B_{i,t}(\theta), \tag{A.21}$$

and note that the quantity multiplying $B_{i,t}(\theta)$ is positive (recalling $\delta > 0$ and $0 \leq b < 1$). We thus see that a potential carbon tax produces greater risk in the future *payments* of carbon taxes, relative to firm value $V_{i,t}$, for a firm with a larger value of $B_{i,t}/V_{i,t}$. Specifically, when the probability of a future carbon tax increases unexpectedly from θ by a small amount, the direct impact on the firm's value is proportional to $B_{i,t}$, resulting in an immediate stock return proportional to $-B_{i,t}/V_{i,t}$, or $-$CB/M, using the notation in Section 6.2.

Firm value is given by

$$V_{i,t} = X_{i,t} - A_{i,t}, \tag{A.22}$$

where $X_{i,t}$ is the present value of future profits before payments of carbon taxes. While it seems plausible that $dX_{i,t}/d\theta \leq 0$, at least for browner firms, it is not clear how $dX_{i,t}/d\theta$ should relate, if at all, to $B_{i,t}$ across firms. Deriving a firm's total carbon tax risk, $dV_{i,t}/d\theta$, which includes both the direct and indirect components (direct through $A_{i,t}$, indirect through $X_{i,t}$), would require a more extensive model of equilibrium responses, as in Garleanu and Pedersen (2025). This section clarifies why the direct component of a firm's carbon risk is proportional to the firm's carbon burden, and specifically to its share of the firm's market value (CB/M), consistent with our empirical analysis. The example presented in Section 6.2 is a special case of the framework presented here, with $\delta = 1$, $b = 0$, θ changing unexpectedly from 0 to 1, and $dX_{i,t}/d\theta = 0$.

[35] Differentiate using the product rule, with $A_{i,t} = f(\theta)g(\theta)$, $f(\theta) = \delta(1-b)\theta$, and $g(\theta) = (1-b\theta)^{-1}B_{i,t}(\theta)$. Note that $g'(\theta) = 0$, applying equation (A.17), so $\frac{dA_{i,t}}{d\theta} = f'(\theta)g(\theta)$, giving equation (A.21).

A.3. Carbon emissions data

We analyze emissions at three levels: scope 1, 2, and 3. In some cases, we also sum across scopes, computing scope 1+2 and scope 1+2+3. We recognize that emissions are double-counted when we sum scope 1+2 or 1+2+3 emissions across firms.

Our emissions data come primarily from MSCI. The forecast data are described in Section 4.1. The historical emissions data start in 2008. MSCI reports emissions by fiscal year. We assign fiscal years ending between January 1 and May 31 to the previous calendar year. For example, when a firm's fiscal year ends in February 2020, we take the calendar year to be 2019, but when the fiscal year ends in November 2020, we take the year to be 2020.

We also use historical emissions data from Trucost, which we obtain from WRDS. We use Trucost data from years 2016 to 2022 because data coverage before 2016 is low. Like MSCI, Trucost reports emissions by fiscal year. We assign them to calendar years in the same way. We also obtain several firm-level variables from CRSP and Compustat. We begin with the set of U.S. firms in the intersection of the MSCI and CRSP/Compustat databases, which we merge by CUSIP, and then merge in Trucost by GVKEY.

There are some extreme outliers in firms' fraction changes in emissions. Some of these appear to be data mistakes. To deal with these outliers, we apply a few filters to both the historical MSCI and Trucost data, for all three emission scopes (scope 1, 1+2, 1+2+3):

1. If the level of emissions and emissions/revenue both increase (decrease) by more than 9x over the previous year and then both decrease (increase) by more than 9x over the following year, then set the year's emissions (and all variables depending on it) to missing. This filter catches large spikes or troughs in emissions that are not accompanied by a spike or trough in revenues. We suspect these are data mistakes. The number 9 is chosen to catch decimal-place mistakes, which would change emissions by a factor of 10. In the MSCI data, this filter sets 15 (6) [24] observations to missing for scope 1 (1+2) [1+2+3]. In the Trucost data, this filter sets 7 (7) [7] observations to missing for scope 1 (1+2) [1+2+3].

2. If the level of emissions is more than 100x larger (smaller) than in the previous year, and if revenues are less than 10x larger (smaller) than the previous year, then set emissions in this year (and all variables depending on it) to missing. This filter catches large, non-reverting changes in emissions that are not accompanied by a similar change in revenues. In the MSCI data, this filter sets 53 (13) [25] observations to missing for scope 1 (1+2) [1+2+3]. In the Trucost data, this filter sets 41 (6) [26] observations to missing for scope 1 (1+2) [1+2+3]. One reason why the numbers of missing observations are higher for MSCI than Trucost is that we use more years of data from MSCI than

Trucost.[36]

[36]We correct one other data mistake in Trucost. We replace the Trucost 2016 Scope 1 emissions for Exxon with the corresponding value from MSCI. In Trucost, scope 1 emissions of Exxon spike roughly threefold in 2016. When asked about this spike, S&P Global, the owners of Trucost data, said they plan to rectify it future versions of their data.

Online Appendix

A.4. VAR-based emission forecasts

Our primary source of firm-level emission forecasts, used in Sections 4 to 6, is MSCI. In this section, we construct an alternative secondary source that does not use data on future emissions. Instead, we build a simple econometric model that uses data on historical emissions to forecast each firm's future emissions into perpetuity.

A.4.1. VAR methodology

We use a vector autoregression (VAR) to forecast firms' shares of aggregate emissions. Our forecast of each firm's future emissions is the product of the aggregate emissions forecast (from Section 3.1) and the firm's forecasted share (from our VAR model). We model firms' shares of aggregate emissions to ensure that our forecasts of firm-level emissions add up to a constant fraction of the aggregate forecasts, for consistency.

Let $\theta_{n,t}$ denote firm n's emissions in year t as a fraction of aggregate emissions. Let $Y_{n,t}$ denote the $1 \times K$ vector containing emission-relevant firm-level variables observable at the end of year t, with $K = 5$. The first element of $Y_{n,t}$ is $\log(\theta_{n,t})$, the main variable of interest. The remaining elements of $Y_{n,t}$ are the same four variables that we related to emission growth forecasts in Table 7: book-to-market, investment, climate score, and revenue growth. We estimate the following first-order VAR, pooled across firms and years:

$$Y_{n,t} = c + Y_{n,t-1}A + u_{n,t}\,, \tag{A.23}$$

where A is a $K \times K$ matrix of coefficients and c is a $1 \times K$ vector of constants. After estimating A and c, we obtain the forecast of $Y_{n,t+\tau}$ as of time t as

$$\mathrm{E}[Y_{n,t+\tau}|Y_{n,t}; c, A] = c\left(\sum_{s=0}^{\tau-1} A^s\right) + Y_{n,t}A^\tau\,. \tag{A.24}$$

We then isolate the element of $\mathrm{E}[Y_{n,t+\tau}|Y_{n,t}; c, A]$ corresponding to $\mathrm{E}[\log(\theta_{n,t+\tau})|Y_{n,t}; c, A]$, which is the firm's forecasted log share in year $t+\tau$. Let $\bar{C}_{t+\tau}$ denote the aggregate emissions forecasted for year $t+\tau$. Then, the emissions forecast for firm n in year $t+\tau$ is

$$\mathrm{E}[C_{n,t+\tau}|Y_{n,t}; c, A] = \bar{C}_{t+\tau}\,\mathrm{E}[\theta_{n,t+\tau}|Y_{n,t}; c, A]\,. \tag{A.25}$$

We substitute these forecasts into equation (1), along with the EPA's SCC forecasts, to compute firms' carbon burdens as of year-end 2022.[37]

[37] In previous sections, we compute carbon burdens as of year-end 2023. We switch to year-end 2022 when using the VAR approach because our historical emissions data end in 2022. Carbon burdens from the VAR approach include emissions forecasted from year 2023 into perpetuity.

One slight complication is that the VAR delivers a forecast of $\log(\theta_{n,t+\tau})$, not a forecast of $\theta_{n,t+\tau}$, which we need in equation (A.25). To go from the former to the latter, we need to make an adjustment for Jensen's inequality. If the VAR's error terms $u_{n,t}$ from equation (A.23) are normally distributed, then the properties of the lognormal distribution imply

$$\mathrm{E}[\theta_{n,t+\tau}|Y_{n,t}] = \exp\left(\mathrm{E}[\log(\theta_{n,t+\tau})|Y_{n,t}] + \frac{1}{2}\mathrm{Var}(\log(\theta_{n,t+\tau})|Y_{n,t})\right). \tag{A.26}$$

The term $\mathrm{E}[\log(\theta_{n,t+\tau})|Y_{n,t}]$ is easily extracted from the VAR, as explained above. If the error terms are i.i.d., then $\mathrm{Var}(\log(\theta_{n,t+\tau})|Y_{n,t})$ is a constant for each τ. Therefore, applying the Jensen's inequality adjustment amounts to adding a τ-specific constant to log shares, or, equivalently, multiplying forecasted non-log shares by a τ-specific constant.

A simple solution to this complication emerges as a byproduct of another fix. We find it desirable for firms' forecasted aggregate emissions shares to be in line with their historical values, but that feature need not obtain empirically without further adjustments. To deliver this feature, we scale the sum of forecasted shares across firms so that it equals the sum of historical shares. Specifically, let $S(\tau)$ denote the sum of $\mathrm{E}[\theta_{n,t+\tau}|Y_{n,t}]$ across firms n. For each τ, we replace $\mathrm{E}[\theta_{n,t+\tau}|Y_{n,t}]$ with $\mathrm{E}[\theta_{n,t+\tau}|Y_{n,t}] \times S(0)/S(\tau)$, which forces the sum of forecasted shares to match its value in $t = 2022$, namely, $S(0)$.[38] This adjustment requires multiplying shares by a τ-specific constant, similar to the adjustment for Jensen's inequality. Therefore, after rescaling shares in this way, we find the same forecasted shares whether or not we apply the Jensen's inequality adjustment in the previous step.

When estimating the VAR, we exclude observations in each year's lowest quartile of emissions, because those observations are the most likely to exhibit extreme, and likely erroneous, year-to-year changes in emissions. However, we apply the estimated VAR model to estimate carbon burdens for all firms, including those in the lowest quartile. We conduct the VAR estimation for scope 1 emissions only, for simplicity.

A.4.2. VAR-based carbon burden estimates

Table A.1 reports the slope estimates for the VAR equation in which the dependent variable is $\log(\theta_{n,t})$. All five independent variables are measured at the end of year $t-1$. The four columns correspond to four different samples: two using historical emissions data from MSCI (columns 1 and 3) and two using data from Trucost (columns 2 and 4). Columns 1 and 2 use as much data as possible from each database (starting in 2008 for MSCI and 2016 for Trucost). Columns 3 and 4 use observations present in both databases.

[38] This value is about 0.4. As noted earlier, direct (scope 1) corporate emissions account for less than half of total emissions.

Table A.1
Forecasting firms' emissions in historical data

This table shows estimates from panel regressions with dependent variable equal to the firm's log scope 1 emissions share in year t. All regressors are measured at the end of year $t-1$. The first two columns use as much data as possible from each database (starting in 2008 for MSCI and 2016 for Trucost). Columns 3 and 4 use observations present in both databases. All regressions exclude observations in the first quartile of emissions. Specifically, column 1 excludes observations in the first quartile of MSCI emissions; column 2 excludes observations in the first quartile of Trucost emissions; and columns 3 and 4 exclude observations in the first quartile of either MSCI or Trucost. In parentheses we show t-statistics double-clustered by firm and year.

	All Observations		Overlapping Observations	
	MSCI	Trucost	MSCI	Trucost
Log(Emissions Share)	0.990	0.987	0.988	0.983
	(342.58)	(852.61)	(293.48)	(427.40)
B/M	-0.030	-0.021	-0.037	-0.025
	(-4.10)	(-1.49)	(-3.65)	(-1.68)
Investment	0.154	0.144	0.132	0.153
	(4.58)	(5.46)	(3.76)	(3.80)
Climate Score	-0.027	-0.077	-0.043	-0.083
	(-2.10)	(-5.14)	(-2.63)	(-4.84)
Revenue Growth	-0.049	-0.102	-0.072	-0.138
	(-1.11)	(-2.55)	(-1.10)	(-2.03)
Constant	-0.099	-0.165	-0.119	-0.195
	(-2.82)	(-7.47)	(-3.68)	(-5.45)
Observations	12150	9820	7291	7291
Adjusted R^2	0.970	0.944	0.971	0.950

Table A.1 shows that the strongest predictor of $\log(\theta_{n,t})$ is its own lag, $\log(\theta_{n,t-1})$, with the slope of almost 1, indicating strong persistence in emissions. Investment also enters consistently with a positive slope, perhaps because firms that invest more subsequently grow more, thereby generating larger future emissions. This finding is present also in Table 7, to a weaker degree. Also similar to Table 7 is the consistently negative slope on the climate score. The estimated slopes on book-to-market and revenue growth are also negative but not always significant. The R-squareds are close to one, especially due to the inclusion of lagged emissions. The results are fairly similar across the four columns.

VAR-based carbon burden estimates differ greatly across firms, even more so than their counterparts based on MSCI's emission forecasts. This fact is apparent from the cross-sectional distributions of carbon burdens scaled by market cap, which we plot in Figures A.1 and A.2, analogous to Figures 3 and 4. Moreover, the VAR-based estimates tend to be larger. For example, using the 2% discount rate and MSCI data, 48% of firms have carbon burdens exceeding their market caps. The fraction is even larger, 62%, when we estimate the VAR based on Trucost data. In both datasets, the firms whose carbon burdens exceed their market caps represent more than 14% of total market cap—somewhat higher than the 10% observed earlier in Figure 4 based on MSCI forecasts. Even under the higher 2.5% discount rate, VAR-based carbon burdens exceed the market cap for 28% of firms based on MSCI historical emissions and for 39% of firms based on Trucost emissions, representing more than 7% of total market cap in both cases.

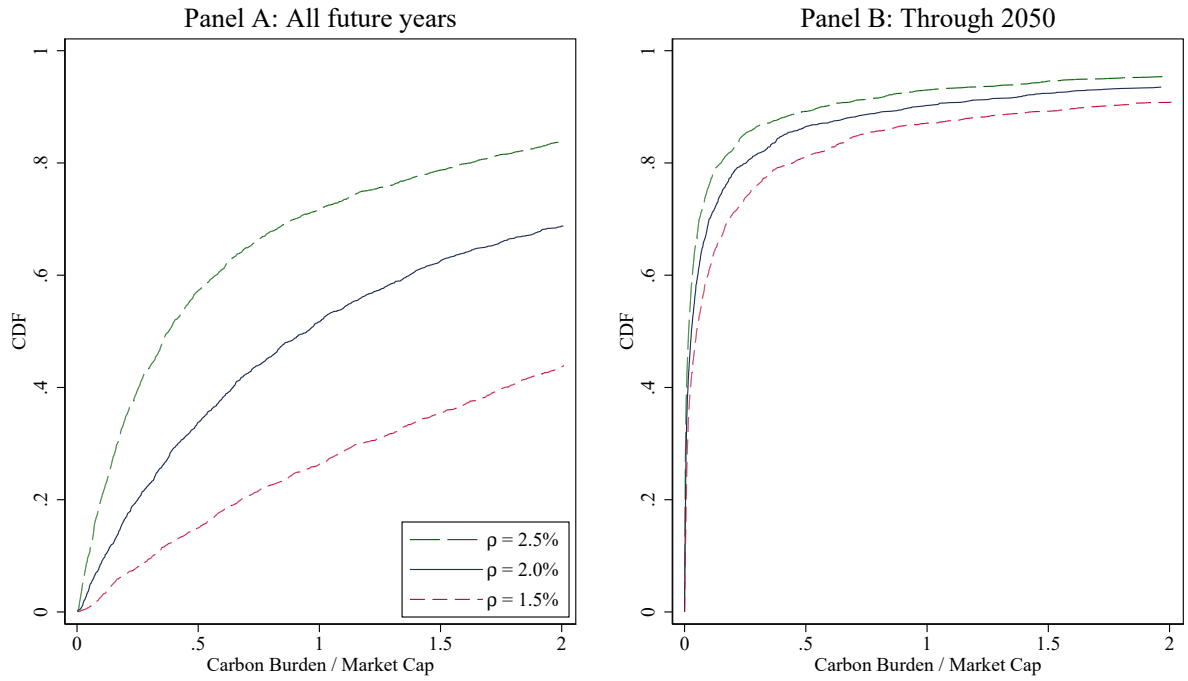


Figure A.1. Distribution of carbon burden / market cap from the VAR approach with MSCI data. This figure shows CDFs of carbon burden / market cap, computed from the VAR approach with historical scope 1 emissions data from MSCI. Carbon burden and market cap are both measured as of the end of 2022. The CDFs weight each firm equally.

The table below shows the fraction of companies with a Carbon Burden to Market Cap ratio greater than 1.

Discount Rate	All years (Panel A)	Through 2050 (Panel B)
2.5%	0.282	0.070
2.0%	0.482	0.097
1.5%	0.736	0.129

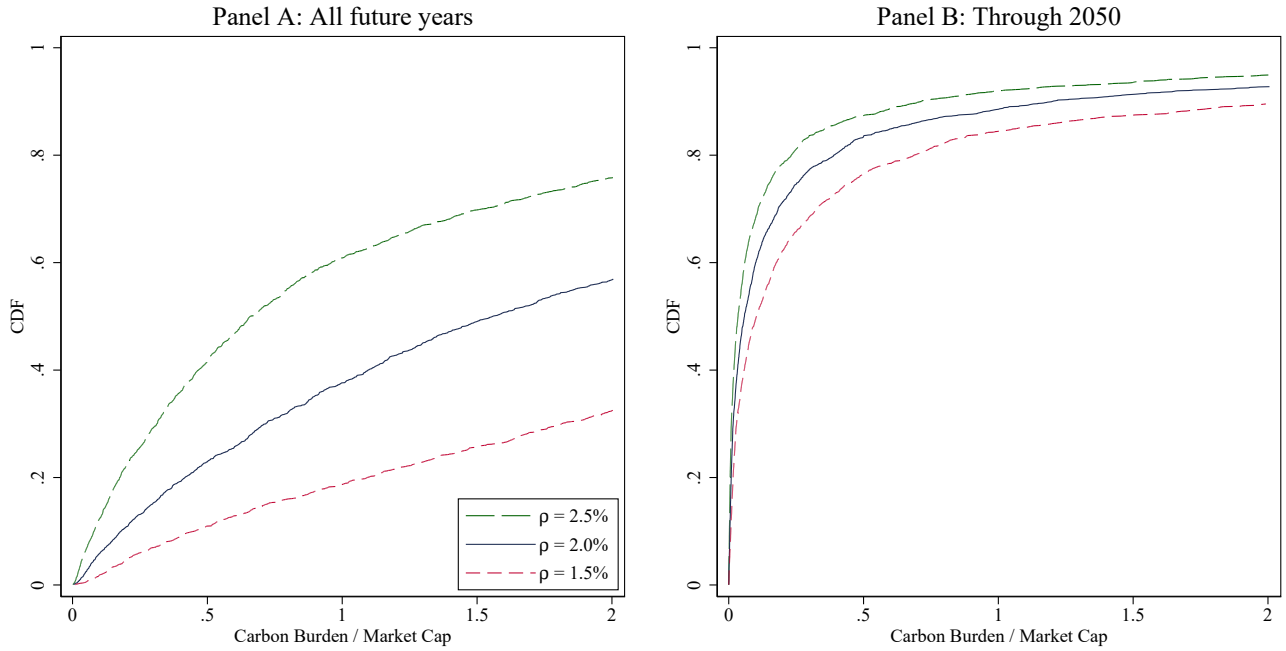


Figure A.2. Distribution of carbon burden / market cap from the VAR approach with Trucost data. This figure is the same as Figure A.1 but shows results based on historical Trucost emissions data.

The table below shows the fraction of companies with a Carbon Burden to Market Cap ratio greater than 1.

Discount Rate	All years (Panel A)	Through 2050 (Panel B)
2.5%	0.391	0.080
2.0%	0.623	0.115
1.5%	0.813	0.155

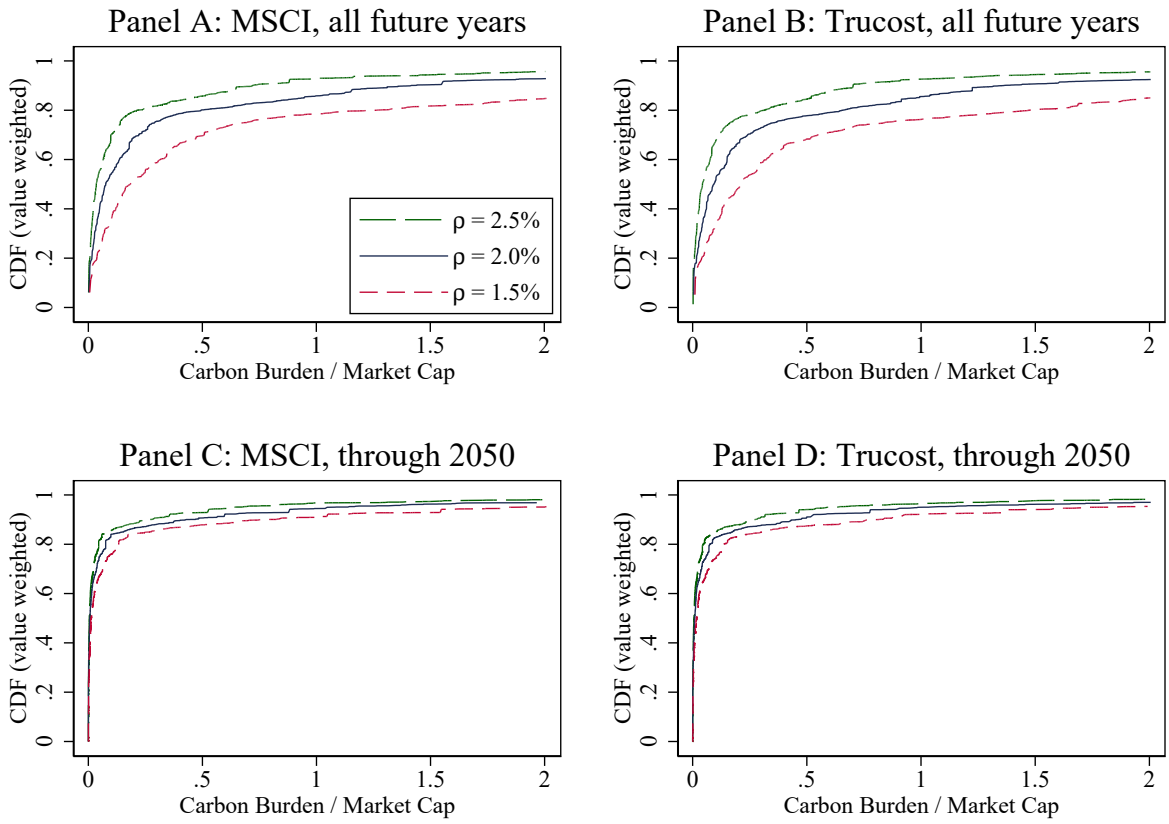


Figure A.3. Value-weighted distribution of carbon burden / market cap from the VAR approach. This figure shows cumulative distribution functions (CDFs) of carbon burden / market cap, computed in the cross section of firms in 2022. We use carbon burden computed using the VAR approach. The VAR model is estimated using all years' historical emissions from each database. The CDFs weight each dollar of market cap equally by plotting the fraction of aggregate market cap belonging to firms with carbon burden / market cap below the x-axis value.

The table below shows the fraction of market cap belonging to companies with a Carbon Burden to Market Cap ratio greater than 1.

	All future years		Through 2050	
	MSCI	Trucost	MSCI	Trucost
Discount Rate	(Panel A)	(Panel B)	(Panel C)	(Panel D)
2.5%	0.073	0.074	0.032	0.035
2.0%	0.142	0.144	0.055	0.050
1.5%	0.215	0.237	0.087	0.079

The previous paragraph suggests that the VAR-based carbon burdens estimated based on Trucost data tend to be larger than those estimated based on MSCI data. In Panel A of Figure A.4, we conduct this comparison more closely by showing a scatterplot of firms' Trucost-based VAR estimates of carbon burdens against MSCI-based VAR estimates. All of these estimates are computed from emissions in all future years and scaled by the firm's market cap. The scatterplot confirms that for most firms, Trucost-based VAR estimates are larger, but there are also many firms for which the opposite is true. The scatterplot is concentrated near the 45-degree line, indicating a fair amount of resemblance between the carbon burdens computed based on the two different data sources.

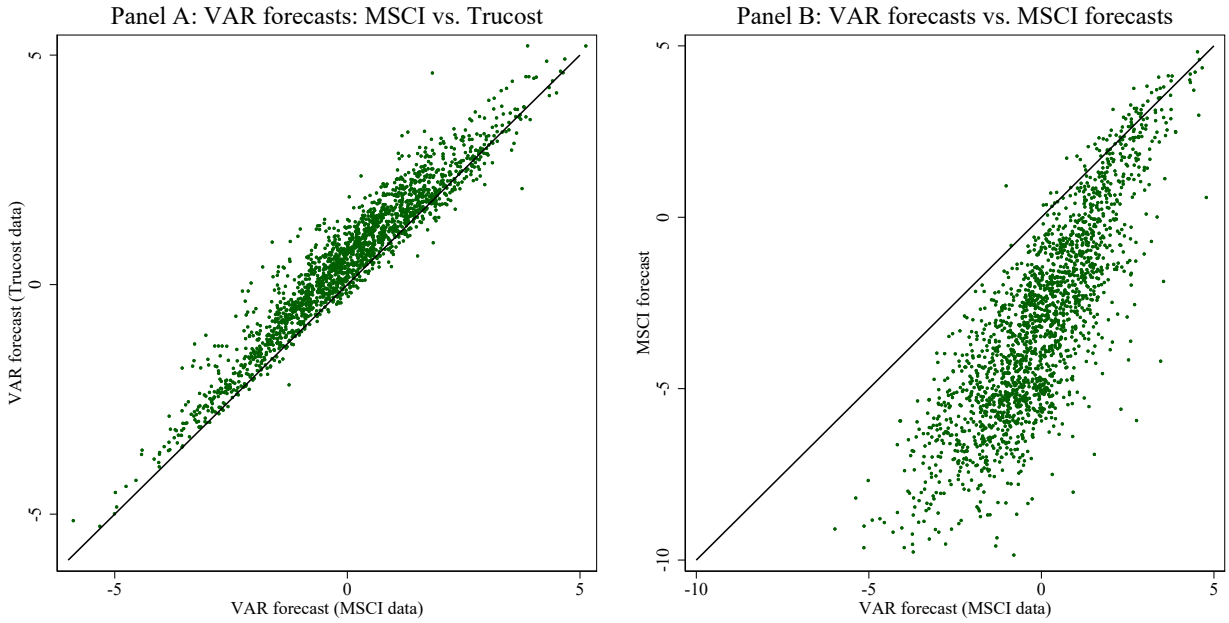


Figure A.4. Comparing carbon-burden estimates. This figure shows scatter plots of firms' ratios of carbon burden to market cap, comparing estimates from one method to another. In Panel A, both dimensions use VAR-based forecasts, but the y-axis is based on historical Trucost data and the x-axis is based on historical MSCI data. Panel A uses the overlapping sample of MSCI and Trucost data. In Panel B, the y-axis uses MSCI forecasts, and the x-axis uses VAR-based forecasts based on historical MSCI data. We use carbon burdens from all future years, with $\rho = 2\%$. All variables are in logs.

Motivated by the deviations from the 45-degree line, we analyze the discrepancies between MSCI's and Trucost's historical emissions data for the same firm in the same year. We conduct the analysis in Appendix A.5.

How do VAR-based carbon burdens compare to those computed based on MSCI forecasts in Section 5? In Panel B of Figure A.4, we produce a scatterplot analogous to that in Panel

A, except that on the y-axis, we replace Trucost-based VAR estimates with estimates based on MSCI forecasts. The plot shows a high degree of similarity between the two estimates for the highest emitters, but a low degree of similarity for the lowest emitters. For most firms, especially for low emitters, carbon burden estimates based on MSCI forecasts are lower than VAR-based estimates. There are at least two reasons. First, MSCI's forecasts reflect firms' forward-looking decarbonization targets (see Section 4.1), which are often more ambitious than the emission reductions that can be inferred from historical data. Second, our VAR approach implies that in an infinitely distant future, all firms' shares of aggregate carbon emissions will be the same. This implication is not unreasonable, given the large amount of long-run creative destruction in the economy. One corollary is that smaller emitters' emission shares are forecasted to grow faster, boosting such emitters' VAR-based carbon burden estimates. As noted earlier, we prioritize emission forecasts from MSCI and use VAR-based forecasts only for comparison.

Recall from Figure 8 that based on MSCI emission forecasts, the top 30 emitters account for essentially all of the expected aggregate decline in emissions by 2050. Figure A.5 shows that this result holds up, and is even stronger, based on VAR forecasts. According to our VAR estimates based on MSCI data, aggregate corporate emissions are expected to decline by 0.3 billion metric tons between 2022 and 2050. The emissions of the top 30 emitters are expected to decline by 0.4 billion tons over the same period, whereas those of the remaining firms are expected to increase by 0.1 billion tons. These results support the conclusion from Figure 8 that all of the decarbonization of the U.S. corporate sector in the coming decades is expected to come from the 30 largest emitters.

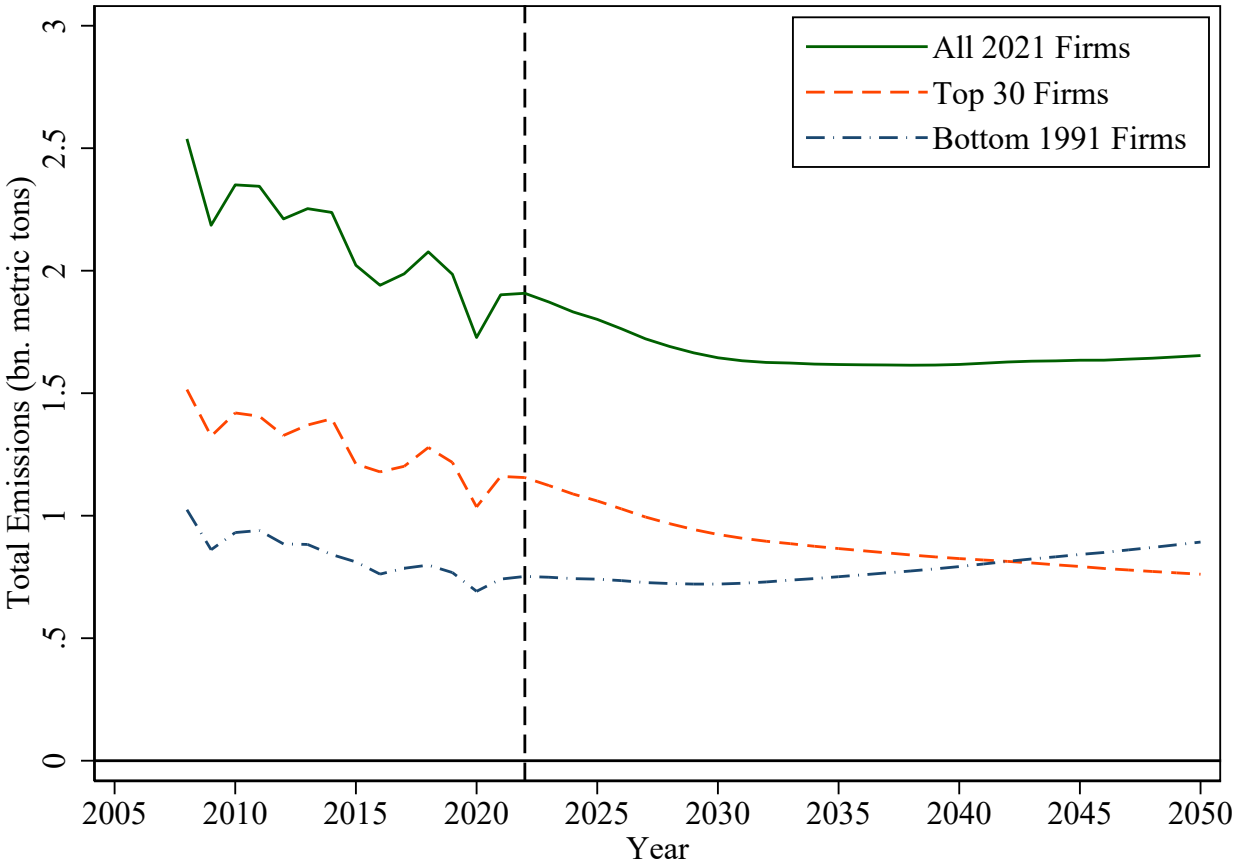


Figure A.5. Past and future emissions from VAR forecasts. This figure is the same as Figure 8, except future emissions are from VAR-based forecasts. The VAR is estimated using historical MSCI scope 1 data. The sample includes firms for which we can forecast emissions after 2022 using the VAR model.

A.5. Emission data discrepancies: MSCI vs. Trucost

In this section, we compare the historical emissions data from the two sources that we use in our firm-level analysis, MSCI and Trucost. If firms' emissions were directly observable, the data from the two sources would presumably be identical. However, emissions are not easy to measure. Some firms disclose their emissions, and both MSCI and Trucost collect such data from publicly available sources such as firms' annual reports, sustainability reports, and regulatory filings. However, many firms do not disclose their emissions, in part because such disclosure is not mandatory as of this writing.[39] Scope 3 emissions are particularly rarely disclosed. Moreover, even emissions that are disclosed are not always credible. Both MSCI and Trucost engage with firms to clarify disclosure-related information. Both also use their own proprietary models to estimate the emissions that are not disclosed as well as emissions that they do not view as credible. The various differences in the data collection processes translate into differences in the emissions data from the two sources.

Panel A of Figure A.4, discussed in Section A.4.2, reveals some differences between MSCI-based and Trucost-based VAR estimates of carbon burdens. Nontrivial differences emerge also between the MSCI- and Trucost-based VAR estimates in columns 3 and 4 of Table A.1, which use the same firm-year samples. In this section, we go deeper, focusing more directly on the differences between the emissions data from MSCI and Trucost, which we refer to as "discrepancies." We summarize the basic properties of these discrepancies, relate them to the levels of emissions and disclosure, and quantify their economic significance. Our bottom line is that these discrepancies are substantial.

We first compute simple correlations between the emission levels from MSCI and Trucost, using firm-by-year panel data from 2016 to 2022. Panel A of Table A.2 shows that these correlations are high, ranging from 81% for scope 3 emissions to 98.2% for scope 1 emissions. While these high correlations might seem reassuring, they obscure some large discrepancies given the enormous variation in emissions across firms. In a cross section in which emissions differ by several orders of magnitude, the correlation between MSCI's and Trucost's numbers can be high even if these numbers differ by a factor of, say, three.

[39] Currently, emission disclosure is mandatory only at the facility level through the EPA's Greenhouse Gas Reporting Program, and only for sufficiently large emitters.

Table A.2: Measurement discrepancies in levels and growth rates

Panel A shows the correlation between MSCI and Trucost emissions levels, using panel data from 2016 to 2022. Panel B shows the cross-sectional percentiles of firms' ratios of (i) the absolute difference between MSCI and Trucost emissions to (ii) the average of MSCI and Trucost emissions, using 2022 data only. Panel C shows the cross-sectional percentiles of the absolute difference between MSCI and Trucost emissions growth rates. Growth rates are computed as the fraction change in emissions from 2021 to 2022. We compute Trucost scope 3 emissions as the sum of scope 3 upstream and scope 3 downstream.

	Scope 1	Scope 2	Scope 3
Panel A: Correlations			
	0.982	0.916	0.809
Panel B: Percentiles of discrepancies in levels			
10th	0.000	0.000	0.070
25th	0.004	0.001	0.235
50th	0.317	0.337	0.710
75th	1.019	0.890	1.611
90th	1.544	1.488	1.963
Panel C: Percentiles of discrepancies in growth rates			
10th	0.000	0.000	0.038
25th	0.013	0.022	0.136
50th	0.110	0.206	0.369
75th	0.379	0.496	1.127
90th	0.858	0.998	6.749

Let $\mathcal{C}_{n,t,s,MSCI}$ and $\mathcal{C}_{n,t,s,Trucost}$ denote scope s carbon emissions of firm n in year t from the two data sources. We measure the MSCI-Trucost discrepancy in levels by computing

$$L_{n,t,s} = \frac{|\,\mathcal{C}_{n,t,s,MSCI} - \mathcal{C}_{n,t,s,Trucost}\,|}{(\mathcal{C}_{n,t,s,MSCI} + \mathcal{C}_{n,t,s,Trucost})\,/2} \tag{A.27}$$

for each firm, year, and scope. Panel B of Table A.2 shows the cross-sectional percentiles of $L_{n,t,s}$, for $t = 2022$ and $s \in \{1, 2, 3\}$. These percentiles show large heterogeneity across firms.

First, consider scope 1 and 2 emissions. The 25th percentiles of $L_{n,2022,1}$ and $L_{n,2022,2}$ are both smaller than 0.005, indicating that for more than a quarter of firms, the discrepancies are negligible. These are mostly firms that disclose their emissions and whose disclosures are accepted at face value by both MSCI and Trucost. The medians of $L_{n,2022,1}$ and $L_{n,2022,2}$ indicate that, for a typical firm, the difference between MSCI's and Trucost's assessments of the firm's emissions is about one third as large as the firm's average emission level. The 90th percentiles of $L_{n,2022,1}$ and $L_{n,2022,2}$ are both about 1.5, indicating that for about 10% of firms, the discrepancy is 1.5 times larger than the emission level itself. The discrepancies thus range from tiny to huge.

For scope 3 emissions, the discrepancies are larger. For example, the 90th percentile of $L_{n,2022,3}$, 1.963, implies that for 10% of firms, the MSCI-Trucost discrepancy is almost twice as large as the emission level itself. This is not surprising, as scope 3 emissions are notoriously difficult to measure. They are rarely disclosed, so both MSCI and Trucost rely on their own internal models to estimate firms' scope 3 emissions. Our results show that those models produce meaningfully different estimates.

A natural question is whether the MSCI-Trucost discrepancies have shrunk over time as a result of the growing amount of emission disclosure and its rising quality. Figure A.6 plots the time series of the cross-sectional distributions of $L_{n,t,s}$ for all three emission scopes. The plots reveal clear but modest reductions in the level of the discrepancies over time. Even at the end of our sample, the discrepancies remain substantial.

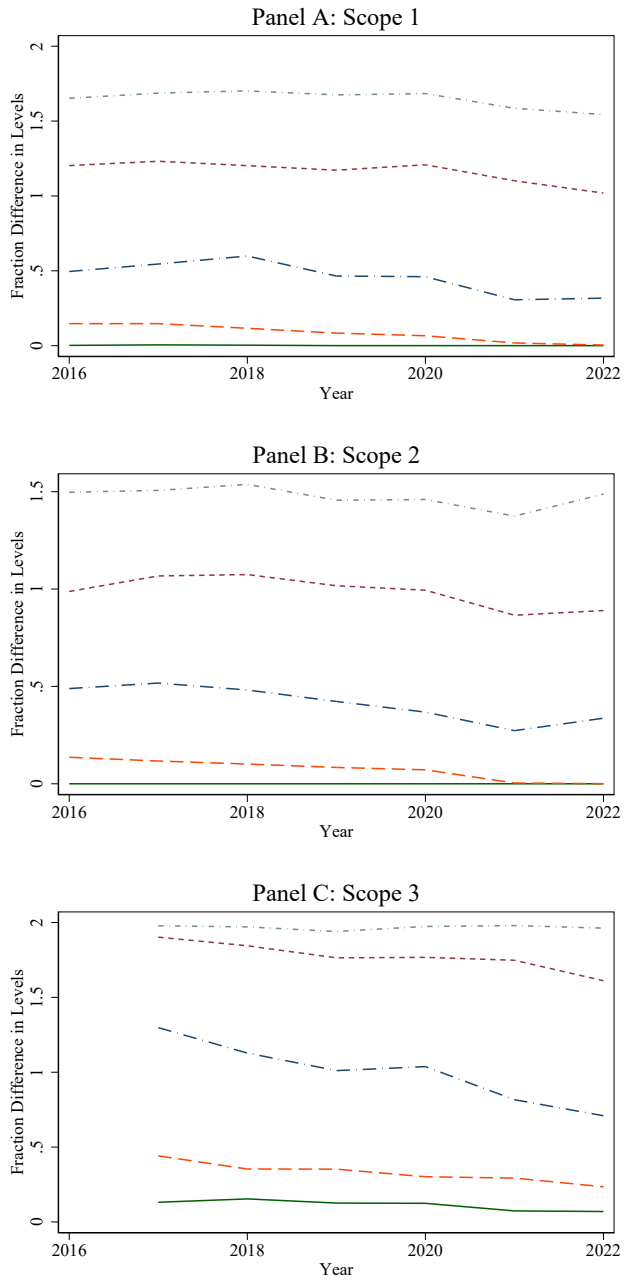


Figure A.6. MSCI-Trucost discrepancies: The time series. This figure plots cross-sectional percentiles each year of the fraction discrepancy between Trucost and MSCI for each scope. From bottom to top, the lines represent the 10th, 25th, 50th, 75th, and 90th percentiles of the fraction discrepancy. For a given firm-year observation, the fraction discrepancy equals the absolute difference between Trucost emissions and MSCI emissions, divided by the average of Trucost and MSCI emissions. Mechanically, the fraction discrepancy cannot exceed 2.

Having examined discrepancies in the levels of emissions, we turn to discrepancies in the growth rates. We measure the MSCI-Trucost discrepancy in growth rates by computing

$$G_{n,t,s} = \left| \frac{\mathcal{C}_{n,t,s,MSCI} - \mathcal{C}_{n,t-1,s,MSCI}}{\mathcal{C}_{n,t-1,s,MSCI}} - \frac{\mathcal{C}_{n,t,s,Trucost} - \mathcal{C}_{n,t-1,s,Trucost}}{\mathcal{C}_{n,t-1,s,Trucost}} \right| \tag{A.28}$$

for each firm, year, and scope. Panel C of Table A.2 shows the cross-sectional percentiles of $G_{n,t,s}$, for $t = 2022$ and $s \in \{1, 2, 3\}$. The patterns are similar to those in Panel B, but the magnitudes are mostly smaller, due to persistence in the levels of the discrepancies.

The 10th percentiles of $G_{n,2022,1}$ and $G_{n,2022,2}$ both round to 0.000, indicating no discrepancies in scope 1 or 2 emission growth rates for at least 10% of firms. The medians of $G_{n,2022,1}$ and $G_{n,2022,2}$ are 0.11 and 0.21, respectively, pointing to nontrivial discrepancies for a typical firm. The 90th percentiles are almost 1, indicating discrepancies exceeding 100% for almost 10% of firms. These are large discrepancies; for example, MSCI might be saying that a given firm's emissions grew by 50% between 2021 and 2022, whereas Trucost is saying that the same firm's emissions fell by 50%. Just as in the levels, discrepancies in the growth rates range from tiny to huge, and they are even larger for scope 3 emissions.

Which firms exhibit the largest MSCI-Trucost discrepancies? We consider two firm characteristics on a priori grounds. First, we hypothesize that the discrepancies could be larger for firms with smaller emissions. Small emitters are less likely to disclose their emissions as well as less likely to be scrutinized by activists or data providers, because whether a firm emits little or very little does not make much difference to society. Second, it would make sense for the discrepancies to be larger for firms that do not disclose emissions, regardless of the emission level. For such firms, MSCI and Trucost estimate emissions based on their own in-house models, which could differ in meaningful ways.

Figure A.7 examines the cross-sectional relations between both characteristics and $L_{n,t,s}$, our discrepancy measure from equation (A.27). Each panel shows a binscatter plot of $L_{n,t,s}$ against the log of firm n's emissions, which we take to be $(\mathcal{C}_{n,t,s,MSCI} + \mathcal{C}_{n,t,s,Trucost})/2$, at the end of our sample ($t = 2022$). There are six panels; the three rows correspond to three different scopes, $s \in \{1, 2, 3\}$, and the two columns represent different sets of firms, either all firms or the subset that disclose their own emissions. To classify a firm as disclosing or not, we follow Aswani, Raghunandan, and Rajgopal (2024). If the Trucost variable "*Scope s disclosure*" contains the string "estimate" (not case sensitive), then we assume the emissions are estimated by Trucost; otherwise we view them as disclosed by the firm.[40]

Figure A.7 shows clear relations between $L_{n,t,s}$ and both characteristics. First, the estimated slope is negative in all six panels, indicating that the discrepancies are larger for smaller emitters. This effect is strong; for example, in Panel A, the average value of $L_{n,t,s}$ for the largest 5% of emitters is less than 0.1, but for more than two-thirds of emitters,

[40] We are able to replicate summary statistics from Aswani et al. (2024) for this variable fairly closely. Our scope 1 (3) data represent Trucost estimates in 71% (93%) of firm-year observations.

the average $L_{n,t,s}$ exceeds 0.5. Second, for scope 1 and 2 emissions, the levels of $L_{n,t,s}$ are substantially larger in the first column of panels, indicating that the discrepancies are larger for firms that do not disclose their emissions. We do not observe the latter result for scope 3 emissions, perhaps because those emissions are disclosed by very few firms (only 68, compared to more than 1,000 for scope 1 and 2). For all scopes, these are still surprisingly large discrepancies even among firms that do disclose. Overall, Figure A.7 shows that the discrepancies are larger for firms that emit little and firms that do not disclose emissions.

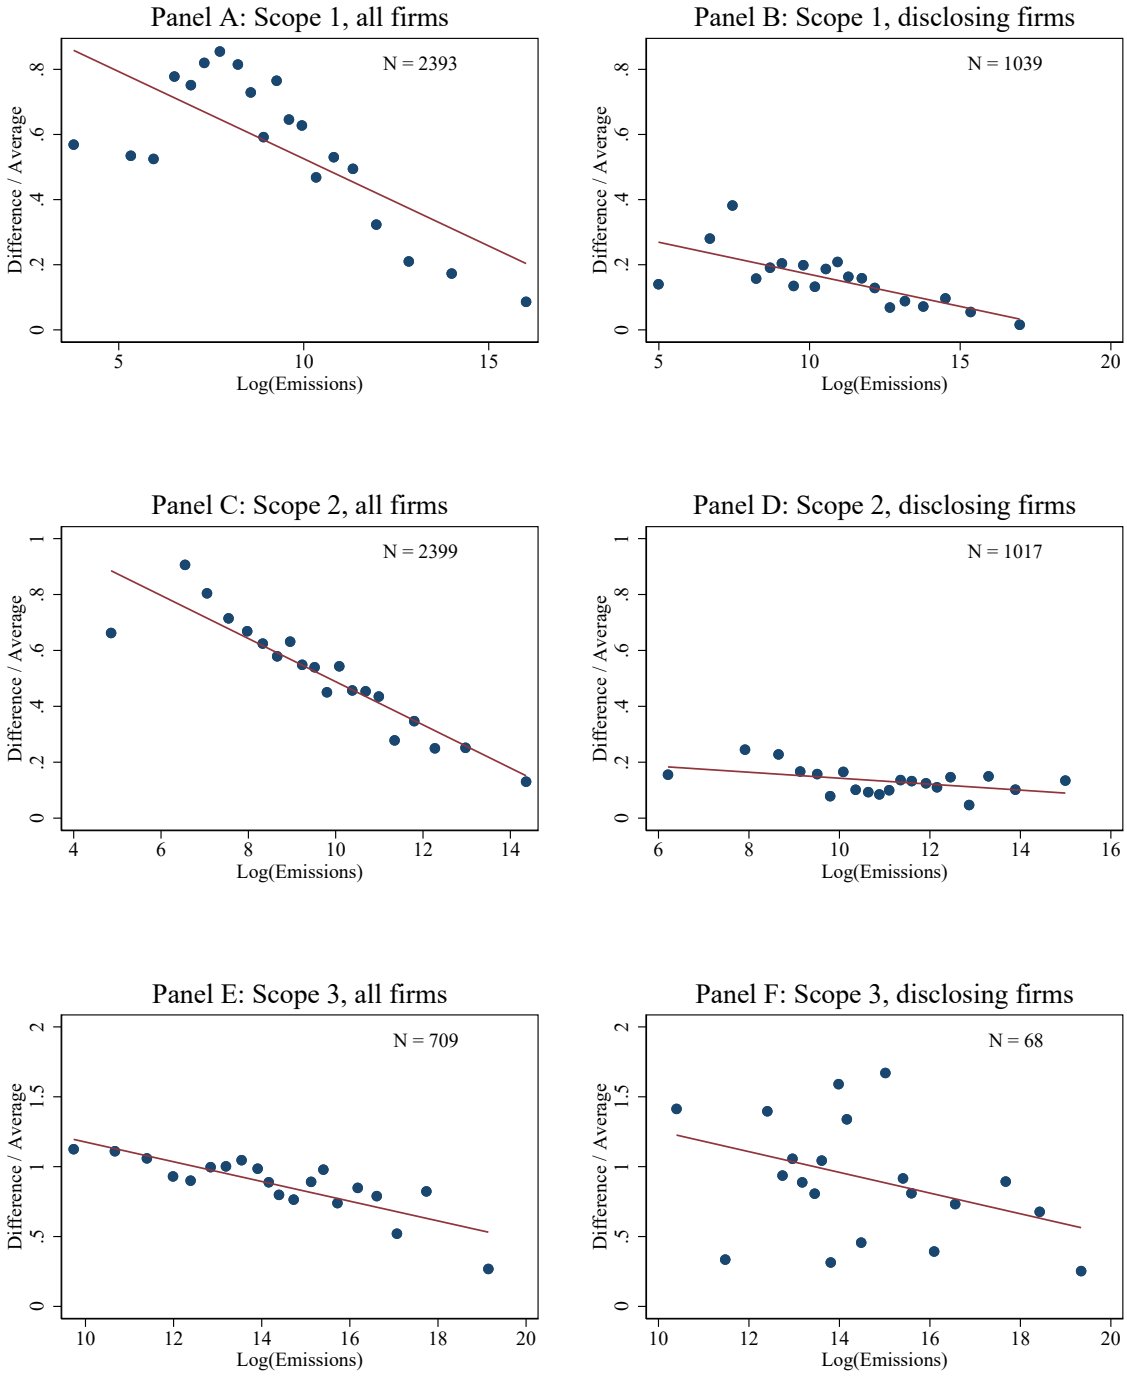


Figure A.7. Discrepancies between Trucost and MSCI. In these binscatter plots, the x-axis denotes the log of the firm's emissions (equal to the average of MSCI and Trucost emissions), and the y-axis denotes the firm's ratio of (i) the absolute difference between MSCI and Trucost emissions to (ii) the average of MSCI and Trucost emissions. Mechanically, that ratio cannot exceed 2. Data are from 2022. A firm is considered to be disclosing if the Trucost variable "Scope X disclosure," for X = 1, 2, or 3, does not contain the string "estimate." Each panel shows the number of firms with non-missing data in both MSCI and Trucost.

Finally, we analyze the economic significance of the emissions-reporting discrepancies between MSCI and Trucost. We consider a hypothetical carbon tax and translate the discrepancies into differences in carbon taxes. We use data from year 2022. We assume a carbon tax rate of \$200 per ton, which equals the EPA's SCC in 2022 with a 2% discount rate.[41] First, we calculate how much each firm would pay in carbon tax if its emissions were assessed by MSCI; we denote this dollar figure by CT_{MSCI}. Note that CT_{MSCI} is simply equal to \$200 times the firm's 2022 MSCI emissions in tons. We then calculate an analogous figure based on Trucost emissions, $\text{CT}_{Trucost}$, and report the absolute difference scaled by the firm's 2022 operating profit: $|\text{CT}_{MSCI} - \text{CT}_{Trucost}|/\text{Profit}$. Table A.3 reports selected properties of the cross-sectional distribution of this ratio within four different groups of firms, which we form by ranking firms on their MSCI emission levels.

Panel A of Table A.3 reports the ratios for scope 1 emissions. The MSCI-Trucost discrepancy is negligible for the median firm, but it is substantial for some firms. For example, for the top 5% of emitters, the 95th percentile of the ratio is 56.75%. This value indicates that 5% of the largest emitters have discrepancies larger than 56.75% of profits. The discrepancies therefore matter a lot for firms that would be paying the most in carbon tax.

Panels B and C of Table A.3 report the ratios for scope 1+2 and scope 1+2+3 emissions, respectively. The differences between Panels A and B are relatively small because scope 2 emissions are small relative to scope 1 emissions for most firms. However, Panel C reports much larger values compared to Panels A and B. For example, based on the means, the ratio of the MSCI-Trucost discrepancies to profits ranges from 52.87% to 165.68% across the four groups of firms. The ratio's 95th percentiles are all in excess of 174% of profits.

[41] To see the results under a different carbon tax rate, the reader can simply scale our results linearly. For example, for a tax rate of \$100 per ton, all numbers in Table A.3 should be multiplied by half (= 100/200).

Table A.3: Implications of measurement discrepancies for carbon taxes

This table considers a hypothetical carbon tax and shows how emissions-reporting discrepancies between MSCI and Trucost would translate to discrepancies in firms' carbon taxes. We use data from 2022 only. We consider a carbon tax rate of $200 per ton, which equals the EPA's social cost of carbon in 2022 with a 2% discount rate. We compute the tax discrepancy as the assumed $200 carbon tax rate (in dollars per ton) times the absolute value of the difference in emissions (in tons) between MSCI and Trucost. We then compute each firm's ratio of the tax discrepancy to operating profit (i.e., revenues minus the sum of COGS, SG&A, and interest expense). The table shows means and percentiles of this ratio, expressed as a percent, across firms within four different groups. The groups, noted in the column headers, are formed by ranking firms based on their MSCI emissions levels. The analysis uses data on 1836 firms for scope 1 and scope 1+2 and 635 firms for scope 1+2+3.

	Emissions Level			
	Bottom 50%	Next 25%	Next 20%	Top 5%
Panel A: Scope 1				
Mean	1.00	2.81	9.35	7.26
50th pctile	0.06	0.18	0.04	0.06
75th pctile	0.93	1.40	1.40	0.56
95th pctile	4.20	13.77	28.59	56.75
Panel B: Scope 1+2				
Mean	2.26	4.68	9.60	12.18
50th pctile	0.30	0.37	0.15	0.07
75th pctile	1.97	3.00	2.19	0.80
95th pctile	8.77	16.61	31.76	36.24
Panel C: Scope 1+2+3				
Mean	165.68	52.87	78.18	54.90
50th pctile	11.06	17.52	25.34	34.15
75th pctile	39.57	44.41	85.90	92.60
95th pctile	463.90	233.20	314.48	174.30

To summarize, we find high correlations between the data from the two providers, especially for direct emissions, similar to Busch, Johnson, and Pioch (2022). However, we show that these correlations mask large discrepancies between the data providers. The correlations are high in spite of these discrepancies because emission levels range widely across firms. The discrepancies are economically significant, as they translate into meaningful differences in hypothetical carbon taxes. We also find that the discrepancies tend to be larger for smaller emitters and for firms that do not disclose their emissions. Firms' emissions are clearly difficult to measure. The substantial divergence between the emissions data from these two leading providers is reminiscent of the divergence of ESG ratings documented by Berg, Koelbel, and Rigobon (2022). Given the growing interest in firm-level emissions data, it seems important to understand the data's limitations.

Finally, note that the measurement problem is even bigger than our results suggest. Even if MSCI and Trucost completely agree on the magnitude of a given firm's emissions, that magnitude need not perfectly match reality. Agreement between MSCI and Trucost often occurs when the firm discloses emissions and those disclosed values are simply accepted by both data providers. However, this acceptance masks the difficulties that the firm itself faces in estimating its own emissions. The fact that neither MSCI nor Trucost challenge the firm's own emission estimates does not necessarily mean that those estimates are precise.

A.6. Additional tables and figures

Table A.4: Share of current emissions and carbon burden by industry

We work with firms at the end of 2023 for which we can measure carbon burden from MSCI forecast data, and which can be assigned to a Fama-French-12 industry. In the "Present" column, we sum year-2023 emissions (measured in tons, taken from MSCI forecasts) within each Fama-French-12 industry and express that industry's sum as a fraction of the sum across all industries. In the "Future" column, we report analogous fractions after replacing current emissions with carbon burdens, computed using MSCI forecasts for all future years. The shares reported in the Scope 1+2 and Scope 1+2+3 columns (but not Scope 1) may to some extent be affected by the double-counting of emissions across firms.

	Scope 1		Scope 1+2		Scope 1+2+3	
Industry	Present	Future	Present	Future	Present	Future
1 Nondurables	0.019	0.022	0.024	0.027	0.029	0.024
2 Durables	0.002	0.002	0.006	0.006	0.033	0.031
3 Manufacturing	0.066	0.072	0.082	0.088	0.106	0.107
4 Energy	0.196	0.202	0.185	0.190	0.254	0.285
5 Chemicals	0.054	0.055	0.061	0.061	0.027	0.026
6 Business Equipment	0.006	0.005	0.019	0.017	0.038	0.028
7 Telecom	0.002	0.001	0.008	0.007	0.006	0.005
8 Utilities	0.414	0.369	0.362	0.323	0.067	0.056
9 Shops	0.032	0.043	0.045	0.057	0.073	0.079
10 Health	0.004	0.004	0.009	0.010	0.017	0.015
11 Money	0.002	0.003	0.010	0.010	0.277	0.282
12 Other	0.202	0.222	0.188	0.204	0.073	0.063

Table A.5: Fraction of firms whose carbon burden exceeds their market cap

Corresponding to Figure 3, this table shows the fraction of companies whose ratio of carbon burden to market cap is greater than 1.

	Scope 1		Scope 1+2+3	
Discount Rate	All future years (Panel A)	Through 2050 (Panel B)	All future years (Panel C)	Through 2050 (Panel D)
2.5%	0.102	0.060	0.655	0.443
2.0%	0.132	0.078	0.769	0.557
1.5%	0.183	0.105	0.873	0.664

Table A.6: Fraction of aggregate market cap belonging to firms whose carbon burden exceeds their market cap

Corresponding to Figure 4, this table shows the fraction of aggregate market cap that belongs to companies whose ratio of carbon burden to market cap is greater than 1.

	Scope 1		Scope 1+2+3	
Discount Rate	All future years (Panel A)	Through 2050 (Panel B)	All future years (Panel C)	Through 2050 (Panel D)
2.5%	0.056	0.028	0.366	0.236
2.0%	0.095	0.036	0.496	0.313
1.5%	0.109	0.065	0.624	0.390

Table A.7: Version of Table 7 dropping observations with 1% growth rate

The sample for scope 1 (1+2) includes firms for which the MSCI scope 1 (1 or 2) forecasted growth rate is not equal to 0.0100, after rounding.

	(1)	(2)	(3)	(4)
Log(Emissions)	-1.368	-1.812	-1.524	-2.072
	(-3.95)	(-4.75)	(-2.74)	(-4.11)
B/M	1.717	1.780	1.997	1.767
	(0.73)	(0.93)	(0.76)	(0.78)
Investment	11.167	10.518	11.241	10.360
	(1.47)	(1.47)	(1.55)	(1.54)
Climate Score	-16.762	-16.864	-18.258	-18.119
	(-5.08)	(-5.14)	(-5.37)	(-5.06)
Revenue Growth	5.390	8.029	4.168	6.677
	(2.37)	(1.05)	(1.62)	(0.80)
Constant	-0.010	-0.004	-0.019	-0.011
	(-2.07)	(-0.72)	(-2.82)	(-1.70)
Observations	612	597	612	597
Adjusted R^2	0.024	0.022	0.029	0.028
Scopes	1	1+2	1	1+2
Industry FE			Y	Y

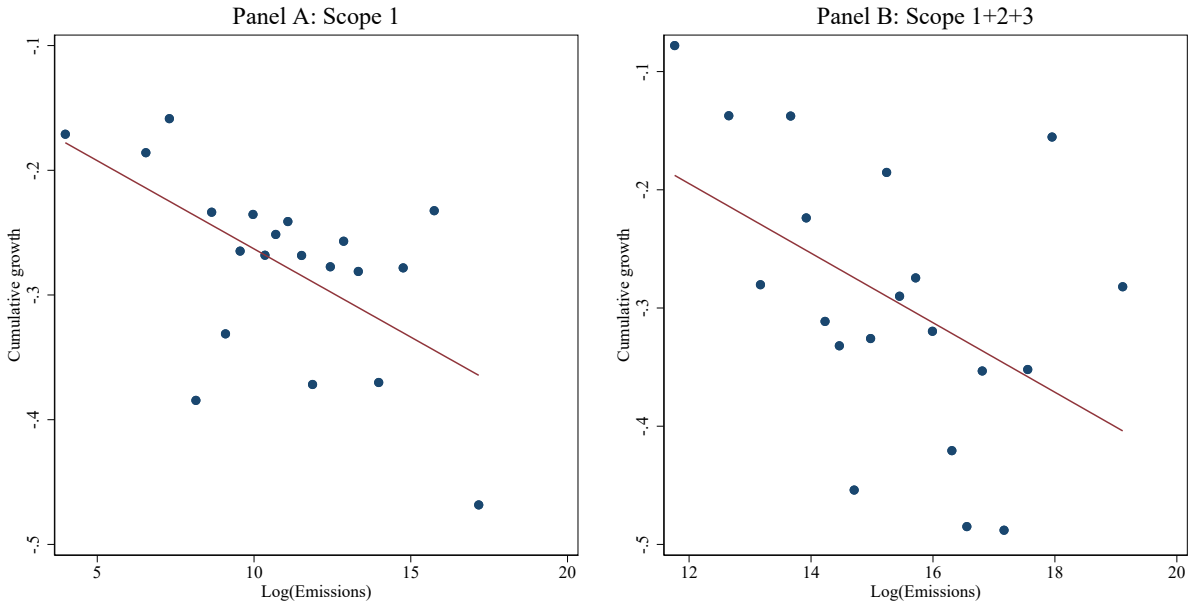


Figure A.8. Version of Figure 7 dropping observations with 1% growth rate. The sample for scope 1 (1+2+3) includes firms for which the MSCI scope 1 (1, 2, or 3) forecasted growth rate is not equal to 0.0100, after rounding. Panel A (B) includes 696 (353) firms in total.

Table A.8
Industry effects in ICCs

This table repeats the regressions from Table 9 while isolating either cross-industry variation (Panel A) or within-industry variation (Panel B). We use the Fama-French 49 industry classification. In Panel A, all variables are collapsed to the industry level using market-cap-weighted averages. Panel B adds industry fixed effects to the firm-level regressions from Table 9. All other details are the same as in Table 9.

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Variation across industries						
CB/M	5.572		20.726	3.220		13.348
	(3.98)		(3.04)	(2.48)		(2.02)
C/M		67.498	-214.556		37.580	-139.981
		(3.34)	(-2.27)		(2.10)	(-1.57)
Market beta				0.152	0.180	0.085
				(1.40)	(1.66)	(0.74)
log(B/M)				0.007	0.008	0.008
				(1.17)	(1.23)	(1.28)
log(M)				-0.001	0.000	-0.005
				(-0.14)	(0.04)	(-0.56)
Constant	0.059	0.060	0.060	-0.079	-0.123	0.030
	(12.39)	(11.97)	(13.11)	(-0.41)	(-0.63)	(0.15)
Observations	49	49	49	49	49	49
Adjusted R^2	0.236	0.174	0.298	0.467	0.448	0.484
Panel B: Variation within industries						
CB/M	0.789		0.987	0.451		0.637
	(4.18)		(1.71)	(2.55)		(1.56)
C/M		12.102	-3.503		6.802	-3.306
		(3.22)	(-0.35)		(2.32)	(-0.58)
Market beta				0.036	0.037	0.036
				(2.27)	(2.30)	(2.25)
log(B/M)				0.022	0.022	0.022
				(10.40)	(10.40)	(10.40)
log(M)				0.002	0.002	0.002
				(1.11)	(1.08)	(1.11)
Constant	0.075	0.075	0.075	0.036	0.035	0.036
	(78.90)	(77.27)	(78.26)	(1.11)	(1.11)	(1.12)
Observations	1983	1983	1983	1918	1918	1918
Adjusted R^2	0.265	0.263	0.265	0.440	0.439	0.440
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes

Table A.9
Version of Table 9 with Scope 1+2 carbon measures

This table is the same as Table 9, except carbon burden (CB) and recent emissions (C) are computed using scope 1+2 instead of scope 1 emissions.

	(1)	(2)	(3)	(4)	(5)	(6)
CB/M	0.995		0.862	0.546		0.934
	(8.07)		(1.98)	(5.23)		(2.45)
C/M		16.216	2.336		8.388	-6.852
		(7.82)	(0.32)		(4.73)	(-1.06)
Market beta				0.061	0.061	0.060
				(5.06)	(5.09)	(5.02)
log(B/M)				0.027	0.027	0.027
				(25.74)	(25.67)	(25.75)
log(M)				0.004	0.004	0.004
				(4.40)	(4.30)	(4.45)
Constant	0.074	0.074	0.074	-0.011	-0.011	-0.011
	(66.66)	(66.36)	(66.42)	(-0.52)	(-0.51)	(-0.51)
Observations	1990	1990	1990	1925	1925	1925
Adjusted R^2	0.031	0.029	0.031	0.315	0.313	0.315

Table A.10
Version of Table 9 with carbon measures winsorized at 99th percentile

This table is like Table 9, except we winsorize the two carbon measures at the 99th percentile.

	(1)	(2)	(3)	(4)	(5)	(6)
CB/M	1.708		2.053	0.898		1.867
	(8.41)		(2.78)	(5.27)		(2.96)
C/M		25.440	-5.653		12.538	-15.956
		(7.94)	(-0.49)		(4.63)	(-1.60)
Market beta				0.059	0.060	0.057
				(4.87)	(4.96)	(4.77)
log(B/M)				0.026	0.027	0.027
				(25.57)	(25.51)	(25.63)
log(M)				0.004	0.004	0.004
				(4.11)	(4.08)	(4.17)
Constant	0.073	0.073	0.073	-0.006	-0.008	-0.005
	(66.25)	(66.10)	(66.16)	(-0.30)	(-0.35)	(-0.25)
Observations	1990	1990	1990	1925	1925	1925
Adjusted R^2	0.034	0.030	0.033	0.315	0.313	0.316

Table A.11
Version of Table 9 dividing by revenue instead of market cap

This table is like Table 9, except we divide the carbon measures by revenue instead of market cap.

	(1)	(2)	(3)	(4)	(5)	(6)
CB/R	0.688		1.704	0.202		1.135
	(4.06)		(3.70)	(1.44)		(2.99)
C/R		7.186	-15.975		0.764	-14.720
		(2.90)	(-2.37)		(0.37)	(-2.64)
Market beta				0.060	0.060	0.057
				(4.91)	(4.96)	(4.70)
log(B/M)				0.027	0.027	0.027
				(25.89)	(25.98)	(26.04)
log(M)				0.004	0.004	0.004
				(4.07)	(4.14)	(4.15)
Constant	0.075	0.075	0.075	-0.007	-0.008	-0.004
	(66.86)	(67.06)	(66.98)	(-0.31)	(-0.35)	(-0.18)
Observations	1977	1977	1977	1913	1913	1913
Adjusted R^2	0.008	0.004	0.010	0.306	0.306	0.309

REFERENCES

Aswani, Jitendra, Aneesh Raghunandan, and Shiva Rajgopal, 2024, Are carbon emissions associated with stock returns? *Review of Finance* 28, 75–106.

Berg, Florian, Julian F. Koelbel, and Roberto Rigobon, 2022, Aggregate confusion: The divergence of ESG ratings, *Review of Finance* 26, 1315–1344.

Busch, Timo, Matthew Johnson, and Thomas Pioch, 2022, Corporate carbon performance data: Quo vadis? *Journal of Industrial Ecology* 26, 350–363.