NBER WORKING PAPER SERIES

CARBON BURDEN

Lubos Pastor
Robert F. Stambaugh
Lucian A. Taylor

Working Paper 33110
http://www.nber.org/papers/w33110

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2024

We are grateful for comments from Jaroslav Borovicka, Xavier Gabaix, Johannes Stroebel, Jeff Wurgler, and seminar participants at NYU and Ohio State. We thank Vitor Furtado Farias, Vanessa Hu, Aria Liu, and Chris Yaoyao Zhu for excellent research assistance. 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. Pastor serves as an independent director and trustee of Vanguard. The views expressed here do not necessarily reflect those of Vanguard or its funds, nor those of the National Bureau of Economic Research.

Carbon Burden
Lubos Pastor, Robert F. Stambaugh, and Lucian A. Taylor
NBER Working Paper No. 33110
November 2024
JEL No. D62, G30, G38, Q51, Q54

ABSTRACT

We quantify the U.S. corporate sector's carbon externality by computing the sector's “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. Among individual firms, 77% have carbon burdens exceeding their market capitalizations, as do 13% of firms even with indirect emissions omitted. 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. Firms' emissions are predictable by past emissions, investment, climate score, and book-to-market.

Lubos Pastor
University of Chicago
Booth School of Business
5807 South Woodlawn Ave
Chicago, IL 60637
and NBER
lubos.pastor@chicagobooth.edu

Robert F. Stambaugh
Finance Department
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367
and NBER
stambaugh@wharton.upenn.edu

Lucian A. Taylor
Finance Department
The Wharton School
University of Pennsylvania
2300 Steinberg Hall - Dietrich Hall
3620 Locust Walk
Philadelphia, PA 19104-6367
luket@wharton.upenn.edu

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 attempt to quantify 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 carbon externality while incorporating the impact of future emissions.

To measure the magnitude of the carbon externality, we propose a metric that we refer to as the "carbon burden." Consider an economic unit, such as a firm or a collection of firms, that emits carbon. We define its 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), which is the dollar cost of societal damages resulting from the emission of one additional ton of carbon into the atmosphere. For an additional 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 the economic unit's expected carbon emissions τ years from now. We define the unit'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 the aggregate, industry, and firm levels to compute carbon burdens at each of those levels. 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, and we also aggregate those forecasts to the industry level.

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. A firm's market value 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 dollar terms, and we compare them in our analysis.

To interpret the carbon burden as measuring an externality, future emissions cannot be priced, such as by levying a carbon tax. There is no nationwide carbon tax in the U.S. as of this writing, though there are some state and local taxes. If any carbon taxes are expected in the future, their estimated present value must be subtracted from the carbon burden to gauge the externality. We measure the externality gross of any carbon tax.

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. 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 we consider other discount rates and shorter time horizons. To help interpret the burden's magnitude, we present a simple framework that links the carbon burden to the corporate profit margin and the cost of capital.

After quantifying the aggregate U.S. carbon burden, we analyze its potential reductions stemming from the 2015 Paris Agreement, under which the U.S. aims to reduce its emissions by at least 26% by 2025 and 50% by 2030, relative to the 2005 level. Again applying the 2% discount rate to all future years, we find that achieving the Paris 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, we find that the largest emitters' reported emission targets are insufficient for the U.S. to meet its Paris goals, even if we take those targets at face value. Moreover, when we replace firms' targets by emission forecasts from MSCI, the shortfall relative to Paris widens further.

At the industry level, carbon burdens differ greatly across sectors. The ratios of carbon burden to market value are as high as 7 and 3 for the utilities and energy sectors, respectively, and as low as 0.01 for the financial and business equipment sectors. These estimates are based on direct (scope 1) emissions, which are emissions from sources owned by the firm. The numbers change little when we add scope 2 emissions, indirect emissions from the consumption of purchased energy, but adding scope 3 emissions, indirect emissions incurred in the firm's entire value chain, matters more.[1] Based on total (scope 1+2+3) emissions, the energy sector's ratio of carbon burden to market value grows to 66, and the ratio exceeds 10 for four other sectors. One of these is financials, whose ratio of 17 for total emissions stands in stark contrast to the 0.01 ratio for direct emissions.

We also divide each sector's carbon burden from all future years' emissions by the sector's burden from a single year's emissions in 2023. This ratio ranges from 49 to 80 across sectors for direct emissions, and the range for total emissions is even broader. Heterogeneity in carbon burdens reflects not only present-year emissions but also differences in forecasts of emissions growth, which are substantial. Firms' cumulative direct emission growth from 2023 to 2050 is −37% for utilities but −6% for financials, based on MSCI's forecasts. The corresponding growth rates based on firms' own reported emission targets are much more negative, ranging from −47% for nondurables to −92% for utilities. The former growth rates are less negative because MSCI views firms' own emission reduction targets as too optimistic. Looking across sectors, MSCI views the 92% and 89% reductions targeted by utilities and chemicals, respectively, as the least credible.

At the firm level, we find high cross-sectional dispersion in the ratios of carbon burden

[1] 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.

to market value. For many firms, these ratios are small; for example, they are smaller than 0.05 for 55% of firms, based on direct emissions. However, for 13% of firms, these ratios are greater than one. For these firms, which represent 9% of total market capitalization, their carbon burdens exceed their market capitalization. Carbon burdens are even larger when we add indirect emissions. Based on total emissions, 77% of firms, which represent 50% of total market capitalization, have carbon burdens larger than their market capitalization.

We also examine the ratio of a firm's carbon burden from all future years to its burden from current-year emissions only. We find the ratio's distribution is quite dispersed across firms, as a result of a large dispersion in MSCI's forecasts of future emission growth rates, which range from −100% to +33% on a cumulative basis between 2023 and 2050. Given this wedge between current and future emissions, it is not sufficient to look at firms' current emissions when judging which firms have larger carbon externalities.

Given its focus on future emissions, a firm's carbon burden could potentially be used to measure the firm's greenness. Suppose two firms have the same emissions today, but the first firm has a credible decarbonization plan whereas the second firm does not. The first firm's carbon burden is then lower, making the firm greener. If firms' carbon burdens were widely reported, they could incentivize firms to develop emission reduction strategies.

We find a negative cross-sectional relation between firms' current 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 U.S. corporate emissions by year 2050. Between 2023 and 2050, 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 is expected to come from only 30 firms.

Besides current 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.

As an alternative to firm-level emission forecasts from MSCI, we also consider forecasts from an econometric model. Our vector autoregression (VAR) model uses data on historical emissions, which we take from MSCI and Trucost, to forecast individual firms' future emissions. Similar to the results based on MSCI forecasts, emissions are expected to grow faster for firms that invest more and firms with lower climate scores. VAR-based results also support our prior conclusion that all of U.S. decarbonization is expected to come from the 30 largest emitters. VAR-based estimates of carbon burdens tend to be even larger than their counterparts based on MSCI's emission forecasts, especially for low emitters. The VAR-based carbon burdens are similar whether estimated from MSCI or Trucost data, though the Trucost-based estimates tend to be somewhat larger. In investigating those differences, we find pervasive and economically significant discrepancies between MSCI's and Trucost's historical emissions data.

The literature on corporate externalities is too large to summarize here.[2] 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). Our contribution is to quantify the carbon externality and compare it with corporate market value.

The study closest to ours is Greenstone, Leuz, and Breuer (2023), who introduce the concept of corporate carbon damages. For a given firm, they compute these damages as the product of the firm's current direct emissions and the SCC (also obtained from the EPA), divided by the firm's current profit or sales. The main difference between our studies is that they study current emissions, whereas we study future emissions. Unlike Greenstone et al., we describe patterns in forecasted future emissions, compute their present value, and compare it to firms' market values. 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 measuring the carbon externality, 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

[2] 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).

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.

Our finding of a big role for large emitters is a symptom of right skewness in the distribution of emissions across firms (e.g., Hartzmark and Shue, 2023). This result is consistent with the finding 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.

The literature that examines carbon emissions from the finance perspective also includes studies on the relations between carbon emissions and the cross section of stock returns (e.g., Bolton and Kacperczyk, 2021, 2023, Aswani, Raghunandan, and Rajgopal, 2024, Zhang, 2024) and on the carbon exposures of institutional investors' equity portfolios (e.g., Bolton and Kacperczyk, 2021, Atta-Darkua, Glossner, Krueger, and Matos, 2023, and Bolton, Eskildsen, and Kacperczyk, 2024). 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.

This paper is organized as follows. Section 2 explains how we compute the carbon burden. Sections 3, 4, and 5 then compute carbon burdens at the aggregate, industry, and firm levels, respectively. Section 6 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.[3] 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). 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 (U.S. Environmental Protection Agency, 2024). 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 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 series of SCC_τ for three initial discount rates that could prevail in τ years: 1.5%, 2.0%, and 2.5% per year. The EPA's choice of discount rates 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

[3] For a recent meta-analysis of the SCC estimates across 207 studies, see Tol (2023).

for long-term cash flows from investments in climate change abatement. When computing its SCC estimates, the EPA starts with the above three discount rates but then allows them 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.

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

When computing an entity's carbon burden, we set C_τ in equation (1) equal to a forecast of the entity's emissions τ years from now. The EPA defines SCC_τ as a marginal cost, thus technically applicable to a relatively small amount of emissions. Applying SCC_τ with C_τ equal to expected emissions for even the entire U.S. corporate sector seems reasonable, however, because U.S. emissions of GHGs in any given year are small relative to the total already 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.[4] This fraction is small because the amount of carbon emitted globally in any given year is small relative to the amount already present in the atmosphere, and also because our analysis is confined to the U.S., whose GHG emissions account for only 17% of global carbon emissions, based on CO_2 equivalents in 2022.[5] Other studies have also applied a social cost per ton to an aggregate of emissions, even at the global 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. Of course, when C_τ equals expected emissions for entities smaller than the total corporate sector, such as industries and individual firms, the argument for applying SCC_τ to those smaller values of C_τ is even stronger.

[4]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.

[5]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.

2.2. Discounting to the present

The quantities C_τ and SCC_τ apply τ periods ahead. Computing the carbon burden in equation (1) requires we discount $C_\tau \times SCC_\tau$ back to the present, using a discount rate ρ_τ. 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 the risk in $\tilde{C}_\tau - C_\tau$ be priced when discounting $C_\tau \times SCC_\tau$? The answer seems elusive. 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.[6]

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.[7] Given this rather flat yield curve at levels just below 2%, one simple specification we choose, especially since we wish to extend τ well beyond 30, is to set $\rho_\tau = 2\%$ for all τ. At that baseline value, ρ_τ includes virtually no risk premium 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: plus and minus 50 basis points. When added to the 2% baseline, those premia give alternative values of $\rho_\tau = 1.5\%$ and $\rho_\tau = 2.5\%$. Therefore, we 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

[6] 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.

[7] 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.

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 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.[8]

The second type, non-energy-related CO_2, is the smallest part of GHG emissions. Non-

[8]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).

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.[9] 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.[10] We use linear interpolation to 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 put 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.[11] Panel B of Table 1

[9] 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.

[10] 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.

[11] This amount equals total issues at market value net of holdings of foreign equities by U.S. residents, as

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.

3.3. Interpreting the burden's magnitude

While the carbon burden is large, particularly 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.

Also, 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 reasonably viewed as 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.[12] It seems difficult to say how much of the choice can be attributed to a country's corporate sector, let alone its

reported in Table L.224 of Board of Governors of the Federal Reserve System (2024).

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

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. At the same time, within 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.

As responsibility for the carbon burden is shared, it seems natural to compare the burden not only to the value of corporate equity but also to society's aggregate 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). Our baseline estimate of the U.S. carbon burden, \$87.1 trillion, thus represents 61% of total wealth. Besides exceeding the total value of equity, the carbon burden thus also constitutes a substantial fraction of U.S. national wealth.

In contrast, the social cost of U.S. emissions seems modest relative to U.S. output. In 2023, when its GDP was \$27.4 trillion, the U.S. emitted 6.28 billion tons of carbon (i.e., GHG in CO_2-equivalent tons). Multiplying this amount by the EPA's baseline SCC estimate of \$204 per ton, the social cost of 2023 U.S. emissions is \$1.28 trillion, which is 4.7% of the 2023 U.S. GDP. How can we reconcile this modest ratio with the large ratio of carbon burden to equity market value? A simple no-growth framework provides the basic intuition.

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

$$Y_t = Y + \epsilon_t\,, \tag{2}$$

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:

$$\mathcal{E}_t = fY\,. \tag{3}$$

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

$$\text{CB} = \frac{fY}{r}\,, \tag{4}$$

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:

$$D_t = hY_t\,, \tag{5}$$

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:

$$M = \frac{D}{r_S}\,, \tag{6}$$

where $D = hY$ is the expected dividend in each period. Combining equations (4) and (6), the ratio of the carbon burden to market value is

$$\frac{\text{CB}}{M} = f\,\frac{1}{h}\,\frac{r_S}{r}\,. \tag{7}$$

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 $1/h = 10$.[13] 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 (7) along with $f = 4.7\%$, which is based on the calculation in the first paragraph of this subsection, we obtain 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.

The purpose of this simple framework is to provide intuition, rather than to be fully calibrated to match a variety of empirical results. In the Appendix, we provide a richer framework that models production using the standard Cobb-Douglas production function, endogenizing the corporate profit margin. The insights we obtain there are very similar to those presented above.

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. Participation by the U.S. in the agreement was withdrawn in 2017 but reinstated in 2021. Under the agreement's Article 4, the U.S. targets 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

[13] 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%.

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 reductions stemming from 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.[14] 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 under the first Paris scenario. Panel A reports the dollar amounts, Panel B divides 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 aggregate predicted emissions to the industry level by summing firm-level forecasts from MSCI, which we describe

[14] 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).

in Section 4.1. We then analyze the carbon burdens associated with industries' predicted emissions 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 the data containing MSCI's forecasts of individual firms' future emissions, which we use not only in this section but also in Section 5.1. These forward-looking data are unique and valuable for the computation of firm-level carbon burdens, which are inherently forward-looking. From the same source, we also obtain MSCI's historical emissions data, which we use in Section 5.2.

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.[15]

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

[15] 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.

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:

$$\text{Forecast}_{n,T} \;=\; w_n \times \text{Target}_{n,T} \;+\; (1 - w_n) \times \text{Base}_{n,T}\,, \tag{8}$$

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.

4.2. Industry carbon burdens

Table 4 reports properties of carbon burdens for our 12 industries. Firm membership in each industry follows the SIC-code classifications of Fama and French, which we obtain from Kenneth French's website. 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 of Table 4 express the carbon burden of an industry's future emissions as a fraction of the industry's market capitalization. The first column considers just scope 1 emissions, the second column adds scope 2, and the third column sums all three scopes. Note that the second and third columns inherently double-count emissions across industries. For example, scope 2 emissions for non-utilities are part of scope 1 emissions for utilities, and a given emission source (e.g., gasoline combustion) can be part of scope 3 emissions for multiple industries (e.g., durables and energy). At the same time, scopes 1, 2, and 3 are mutually exclusive for a given firm. To the extent that double counting is more prevalent across industries than across firms within an industry, the quantities reported in the second and third columns are also meaningful on an industry-by-industry basis.

Table 4 shows that carbon burdens differ greatly across industries. For scope 1, utilities have a carbon burden through 2050 that is 2.71 times the industry's market cap, and energy has the second-largest ratio at 1.06. In contrast, six industries have ratios of 0.05 or less. Adding the later years more than doubles the largest values, with utilities and energy increasing to 6.95 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. Although scope 3 greatly double-counts emissions, it captures more than half of aggregate emissions not captured by firms' scopes 1 and 2.[16] The energy industry's scope 3 emissions subsume much of aggregate emissions, so its carbon burden then becomes the largest by far, 66 times that industry's market cap when including all future years. The carbon burdens of other industries also become much larger when including scope 3. For example, four other industries have ratios 10 or higher in Panel B. One of them is financials, which has a ratio of just 0.01 for scope 1.

The direct emissions of financial firms are small, given the sector's service-based nature. The sector's scope 3 emissions are large, however, because they include also the emissions of companies and projects financed by financial institutions. For example, scope 3 emissions are high for banks providing loans 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 (category 15). 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 express an industry's carbon burden over future years as its ratio to the burden from a single year's emissions in 2023. This future/present ratio generally ranges in the mid-20s when including emissions through 2050, and it is roughly

[16] 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 likely 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 (scope 2 emissions are scope 1 emissions for utilities). 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 aggregate quantity, because scope 3 inherently double counts emissions across firms, unlike scope 1.

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 the present 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 92 for energy, 51% higher. In short, computing carbon burdens of just present-year 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. Based on direct emissions, utilities account for 37% of the total, and energy accounts for 20%. Five other industries have shares below 1%: business equipment, durables, health, money, and telecom. Based on total emissions, however, the utilities' share is only 6%, and the largest shares, over 28%, belong to energy and money. Again, the carbon burden of financials changes dramatically after including scope 3 emissions.

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 credibility MSCI gives the targets (i.e., the higher is the value of w_n in equation (8)).

For scope 1, all of the ratios are well below 1, meaning that MSCI views targets as too optimistic. All industries target substantial emission reductions, but the 92% and 89% reductions targeted by utilities and chemicals are judged the least credible, with forecast-to-target ratios of just 0.11 and 0.12, respectively. The non-durable sector's targeted reduction of 47% is the most modest, but it is also judged the most credible, 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 credible, in that the forecast-to-target ratio increases for every industry. The most credible 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 carbon burdens in two ways: based on MSCI's forecasts of firms' future carbon emissions (Section 5.1) and based on emission forecasts from an econometric model (Section 5.2).

Our emissions data come from MSCI and Trucost. We obtain MSCI data on both historical and forecasted emissions from the MSCI ESG Manager. The forecast data are described in Section 4.1. The historical data start in 2008, when they become available in the ESG Manager. We obtain historical Trucost data from WRDS. We use Trucost data from years 2016 to 2022 because data coverage before 2016 is low. Both MSCI and Trucost report 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 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 we merge in Trucost by gvkey.

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 compute carbon burdens from MSCI emission forecasts as of the end of 2023, covering all future years. We consider three dependent variables, all in logs: unscaled carbon burden, carbon burden divided by the firm's market capitalization, and carbon burden divided by the burden from the firm's emissions in year 2023 only.

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 0.654. 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. Carbon burdens based on MSCI emission forecasts

We compute firms' carbon burdens at the end of 2023 by substituting MSCI's forecasts of firms' future carbon emissions into equation (1). As before, we use three different discount rates and the EPA's SCC estimates. We scale each firm's carbon burden by the firm's market capitalization, denoting the resulting ratio by CB/M.

5.1.1. 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.

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 19% 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 9% 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.

5.1.2. 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 current emissions? Put differently, do MSCI's emission forecasts contain much information that is not already contained in firms' current 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 current 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 current 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.

5.1.3. 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, measured in logs as of 2023. 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. We run these regressions for three emission scopes and both with and without industry fixed effects. In all six specifications, current emissions enter with a significantly negative slope, with t-statistics ranging from -5.14 to -12.54.

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 coefficients are only marginally significant. Investment enters with a positive slope that is significant in three specifications, pointing to larger emission increases for firms that invest more. Only the climate score enters consistently across all six specifications. Its slope estimate is always negative, with t-statistics ranging from -2.83 to -4.51, 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 that variation: adjusted R-squareds range from 6.9% to 12.2%.[17] Clearly, MSCI's approach to forecasting emissions is much 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 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

[17] 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 six specifications, and the other regressors are almost never significant. See the Appendix.

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.[18]

5.1.4. Paris redux

Emission reductions by the largest corporate emitters will be essential in achieving the goals of the Paris Agreement discussed in Section 3.4. To see this, consider the top 10% of emitting firms in each of four emission categories: scope 1, scope 2, scope 3, and their sum. Panel A of Table 8 shows that the top 10% account for a large fraction of emissions by all firms, ranging from 79% for scope 2 to 96% for scope 1. Given their dominance, the largest emitters are pivotal in the country's efforts to cut emissions.

For the U.S. to meet its Paris goals, carbon emissions in 2030 must be 41% lower than in 2023, declining from 6.3 to 3.7 billion tons. For the corporate sector to cut emissions by 41%, the bulk of this cut must come from the top 10% of emitters; the remaining 90% of firms are relatively unimportant. Panel B of Table 8 shows the required reductions for the top 10% under various scenarios for what the other 90% of firms do. If the latter firms' emissions stay constant at 2023 levels, the largest emitters need to cut their emissions by between 43% and 52%, depending on the category. If the other 90% instead cut their emissions to zero by 2030, the top 10% still need to cut their emissions between 26% and 38%.

What reductions by the largest emitters can we anticipate? To answer this question, we turn to the MSCI forecast data. Most firms in the top 10% have emission targets.[19] Panel C of Table 8 summarizes characteristics of these firms, which for each category account for over three-fourths of the 2023 emissions from the entire top 10% of emitters. Panel D reports properties of the firms' targeted emission reductions from 2023 to 2030. The aggregate targeted reductions for scopes 1 and 2, at 28% and 33% respectively, are moderately below the Paris-mandated 41% reduction. However, targeted cuts for scope 3 are only 8%, and

[18]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.

[19]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.

recall that scope 3 captures over half of U.S. aggregate emissions not captured by firms' scopes 1 and 2. To cut half of all emissions by only 8% would leave the U.S. well short of its Paris goals. Further tempering anticipated cuts by the largest emitters are MSCI's forecasts for the target-reporting firms, computed as in equation (8). The reductions from 2023 to 2030 implied by those forecasts are summarized in Panel E of Table 8. In all emission categories, MSCI is rather pessimistic about firms' meeting their targeted reductions, predicting reductions often two or three times smaller than targeted.

5.2. VAR-based emission forecasts

Our primary source of firm-level emission forecasts, used in both Sections 4 and 5.1, 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.

5.2.1. 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:

$$Y_{n,t} = c + Y_{n,t-1}A + u_{n,t} \, , \tag{9}$$

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

$$\mathrm{E}[Y_{n,t+\tau}|Y_{n,t}; c, A] = c\left(\sum_{s=0}^{\tau-1} A^s\right) + Y_{n,t}A^{\tau} \, . \tag{10}$$

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

$$\mathrm{E}[C_{n,t+\tau}|Y_{n,t}; c, A] = \bar{C}_{t+\tau}\, \mathrm{E}[\theta_{n,t+\tau}|Y_{n,t}; c, A]\,. \tag{11}$$

We substitute these forecasts into equation (1), along with the EPA's SCC forecasts, to compute firms' carbon burdens as of year-end 2022.[20]

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 (11). 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 (9) are normally distributed, then the properties of the lognormal distribution imply

$$\mathrm{E}[\theta_{n,t+\tau}|Y_{n,t}] = \exp\left(\mathrm{E}[\log(\theta_{n,t+\tau})|Y_{n,t}] + \frac{1}{2}\mathrm{Var}(\log(\theta_{n,t+\tau})|Y_{n,t})\right). \tag{12}$$

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, $S(0)$.[21] 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.

[20] 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.

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

5.2.2. VAR-based carbon burden estimates

Table 9 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.

Table 9 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 the Appendix, 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 about 14.5% total market cap—somewhat higher than the 9% 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 about 7.5% of total market cap in both cases.

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 9, 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.

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 the Appendix and summarize it here. 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. For example, based on scope 1 or scope 2 emissions, 10% of firms exhibit discrepancies 1.5 times larger than the emission level itself. For scope 3 emissions, the discrepancies are even larger—for 10% of firms, they are twice as large as the emission level. 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. Consider, for example, a tax on direct emissions equal to $200 per ton, the EPA's current baseline SCC. Among the largest emitters (top 5% of emitters based on direct emissions), 5% of them then have discrepancies in carbon taxes that exceed 57% of their annual profits. 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.

How do VAR-based carbon burdens compare to those computed based on MSCI forecasts in Section 5.1? In Panel B of Figure 9, we produce a scatterplot analogous to that in Panel A, except that on the y axis, we replace Trucost-based VAR estimates by 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 10 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.

6. Conclusion

We estimate carbon burdens, novel measures of carbon externalities, 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 695% of market value for utilities to 1% for financials, based on direct emissions. When indirect emissions are added in, the carbon burden of utilities more than doubles, but the financials' carbon burden grows more than thousandfold. For 13% of firms, which represent 9% of total market capitalization, their direct carbon burdens exceed their market values. Adding in indirect emissions, carbon burdens exceed market values for 77% of firms, which make up half of total market capitalization. For these firms, the present value of their carbon costs to society exceeds the present value of their dividends to shareholders. The large magnitudes of the estimated 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 find that adherence to the 2015 Paris Agreement would reduce the aggregate U.S. carbon burden by 21% to 32%. Key to the achievement of the Paris 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. (Besides current emissions, other firm characteristics that help explain the cross section of emission growth forecasts include investment, climate score, and the book-to-market ratio.) 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. Alas, the largest emitters' emission reduction targets are insufficient for the U.S. to fully meet its 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. 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. We could compute an alternative and potentially more precise measure of the carbon burden if there existed 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. Until such an imaginary world arrives, carbon burden estimates are likely to remain imprecise. 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 (e.g., France and Germany have very different attitudes toward nuclear energy). 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. Nonetheless, firms can surely be held responsible for some of their emissions. Designing policies that reduce the aggregate carbon burden fairly, efficiently, and significantly is an important task for scholars and policymakers alike.

Future work should also aim to improve emission measurement. We find substantial discrepancies between the emissions data from two leading providers, MSCI and Trucost. The discrepancies are larger for smaller emitters and firms that do not disclose their emissions. Emission measurement is likely to become more precise if emission disclosures eventually become mandatory in the U.S., as proposed by the SEC.

We also need more research into the risk profile of carbon emissions, to improve the way we discount future emissions. Finally, moving beyond carbon, future research should try to quantify other externalities, positive and negative, that corporations impose on society.

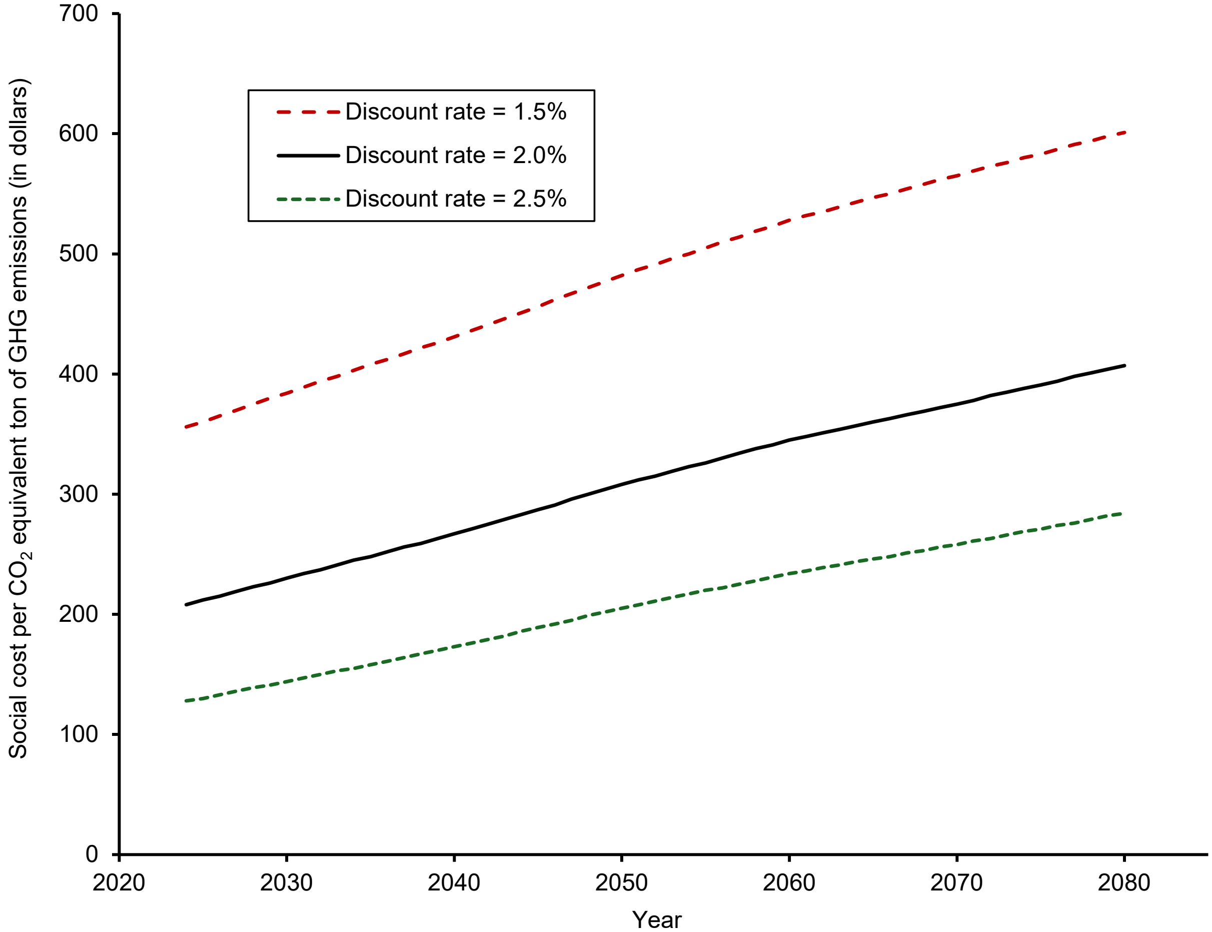

Figure 1. Social costs of GHG emissions. The figure plots EPA estimates of the social cost per CO_2-equivalent ton of GHGs emitted in a given future year. The EPA provides the costs through 2080 that are associated with each of three discount rates: 1.5% (long dashes), 2.0% (solid line), and 2.5% (short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

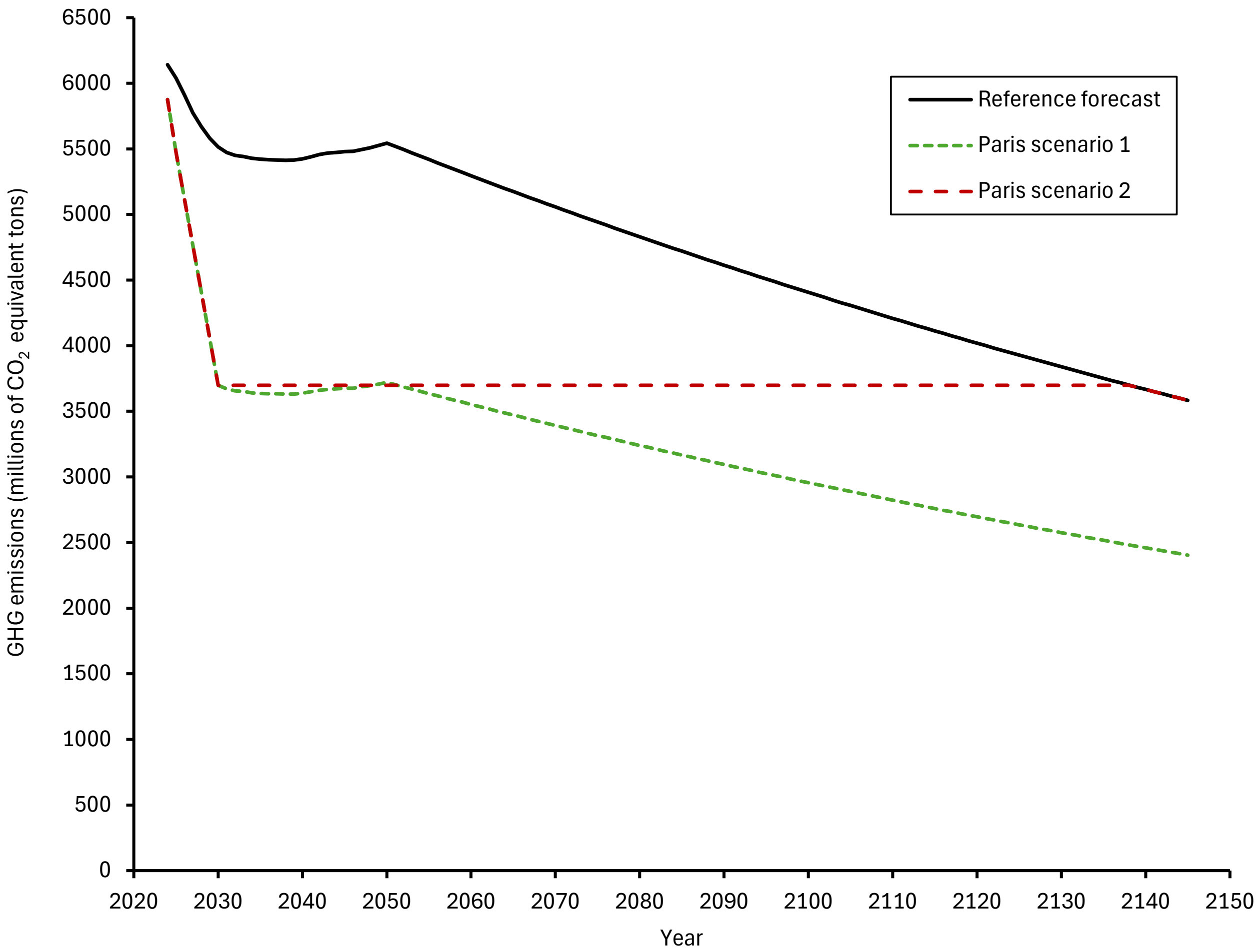

Figure 2. Forecasts of U.S. GHG emissions. The figure plots the reference forecasts (solid line) as well as forecasts under two scenarios for the Paris agreement. In the first Paris scenario (short dashes), the ratio of emissions to reference-level forecasts is maintained at the agreement's 2030 level 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.

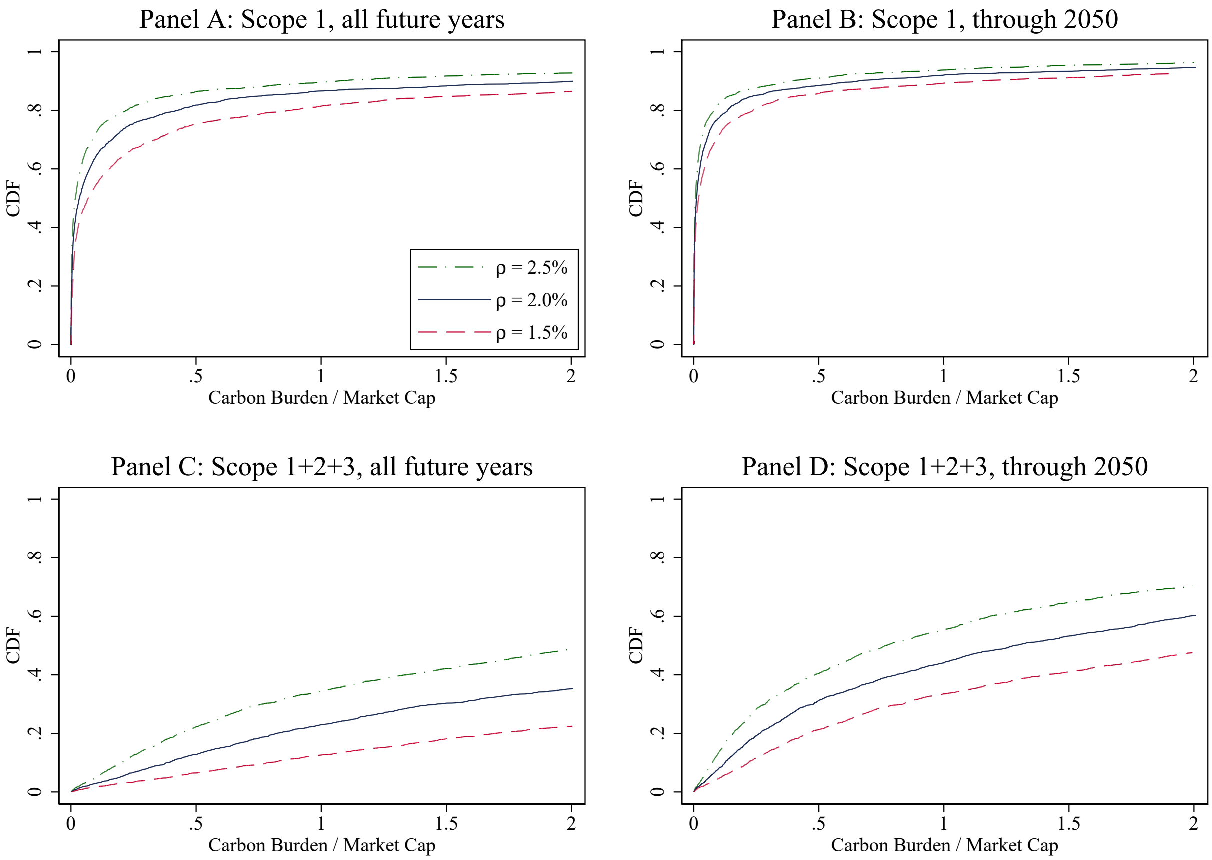

Figure 3. Distribution of firms' carbon burden as a fraction of market cap. This figure shows cumulative distribution functions (CDFs) of the ratio of carbon burden to market cap, computed in the cross section of firms in 2023. Carbon burdens are computed using MSCI's forecasts. The CDFs weight each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

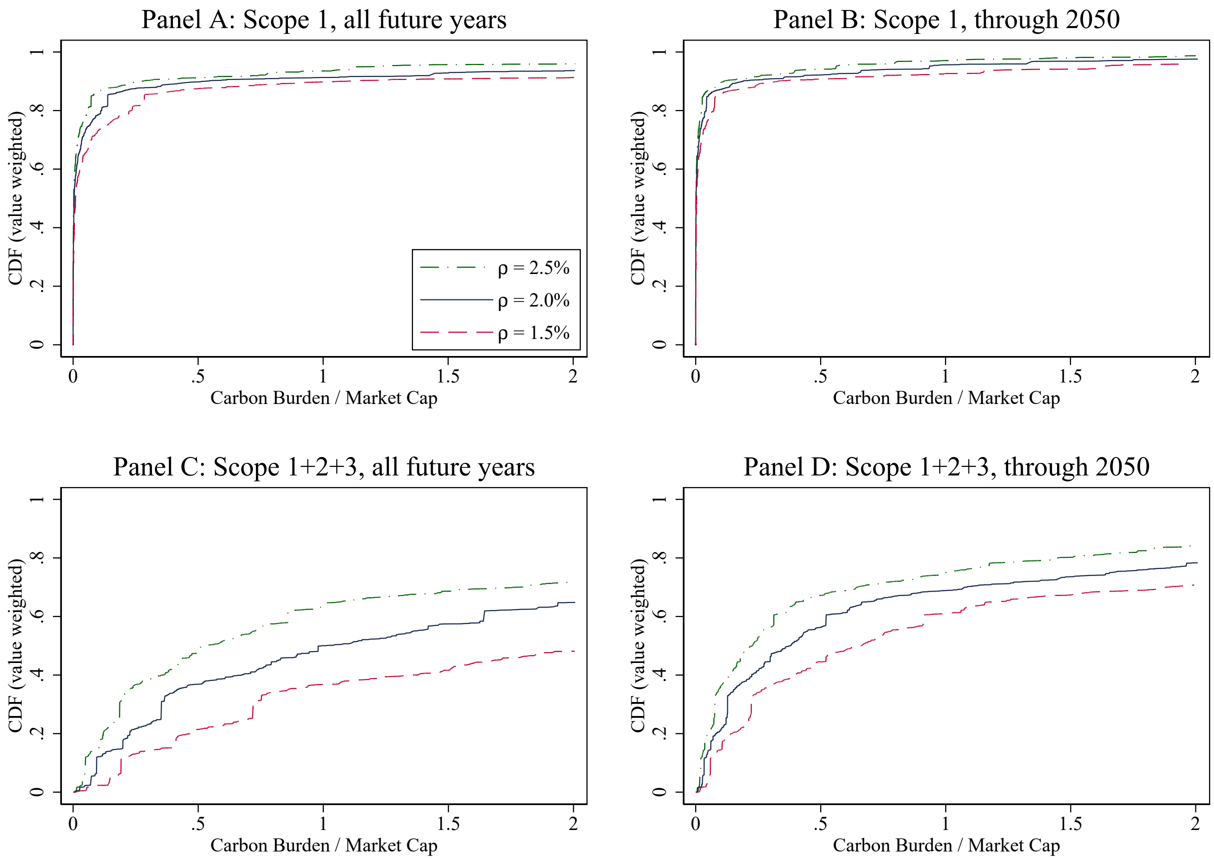

Figure 4. Value-weighted version of previous figure. Whereas the previous figure plots the fraction of firms below each x-axis value, this figure plots the fraction of aggregate market cap belonging to firms below each x-axis value.

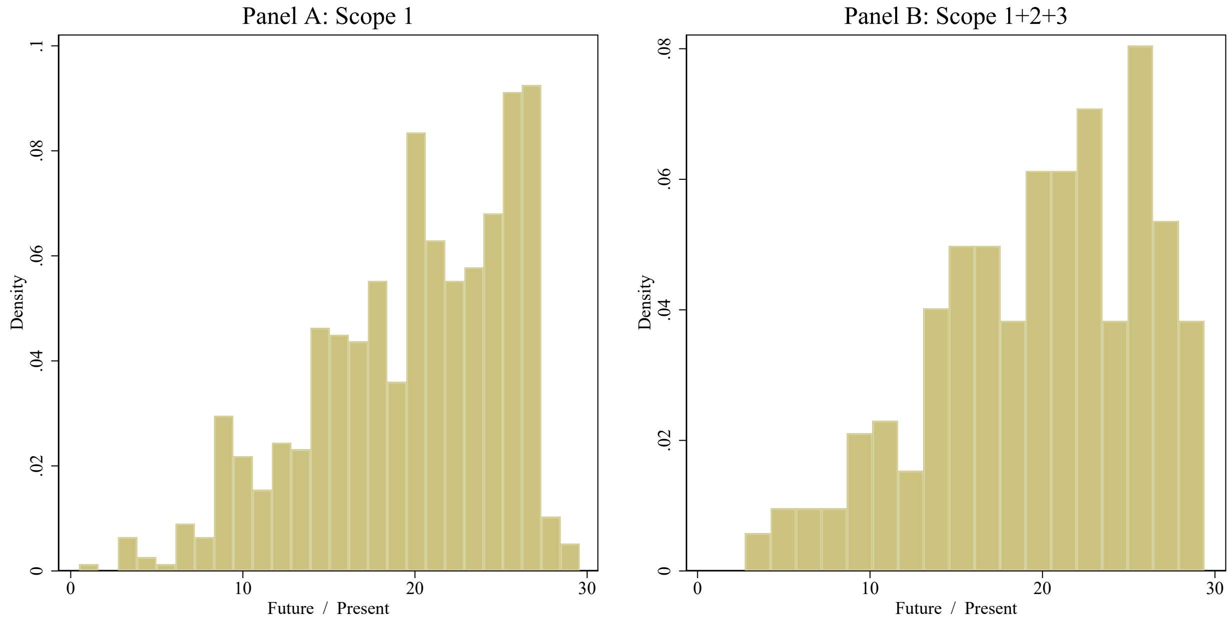

Figure 5. Firms' carbon burdens: Future years vs. current year. We compute each firm's ratio of carbon burden through 2050 to carbon burden from 2023. The figure plots this ratio's distribution across firms. Carbon burdens are computed using MSCI's emissions forecasts, with $\rho = 2\%$. Panel A (B) 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.

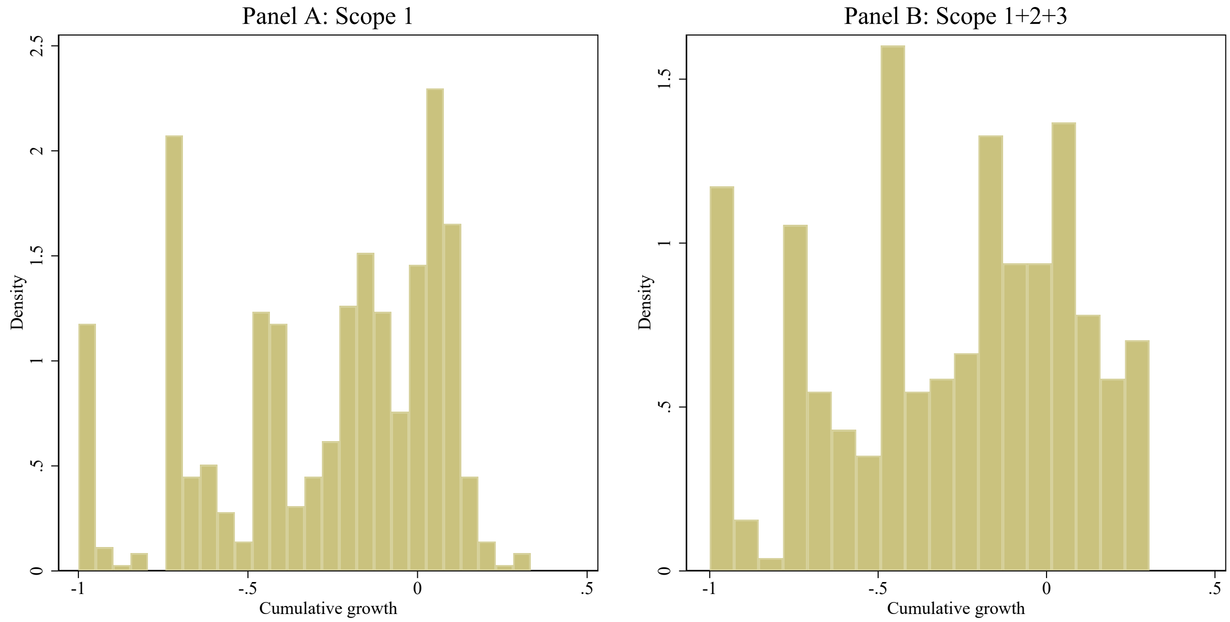

Figure 6. Firms' forecasted emissions growth. This figure plots the distribution of firms' cumulative forecasted emissions growth rates, computed as the fraction change in emissions from 2023 to 2050. Emissions forecasts are from MSCI. Panel A (B) excludes firms with scope 1 (1, 2, or 3) growth rate 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.

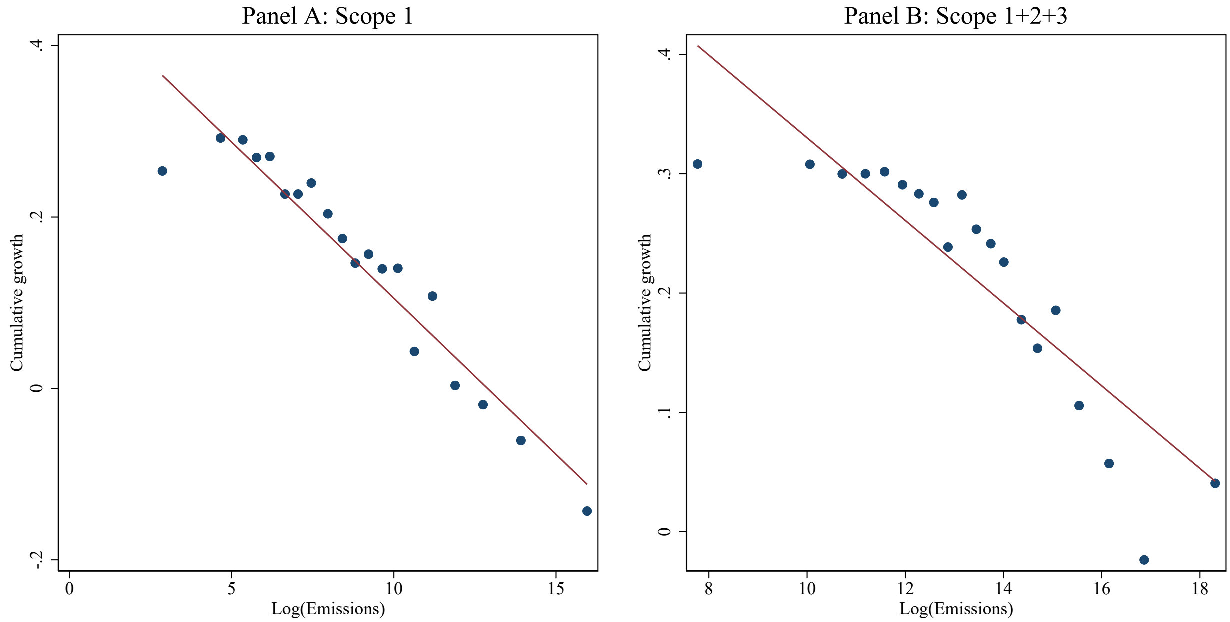

Figure 7. Current emissions and forecasted emissions growth. This figure shows the binscatter plots of firms' cumulative forecasted emissions growth rates, computed as the firm's fraction change in emissions from 2023 to 2050, against the log emissions in 2023. Emissions forecasts are from MSCI. Panel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

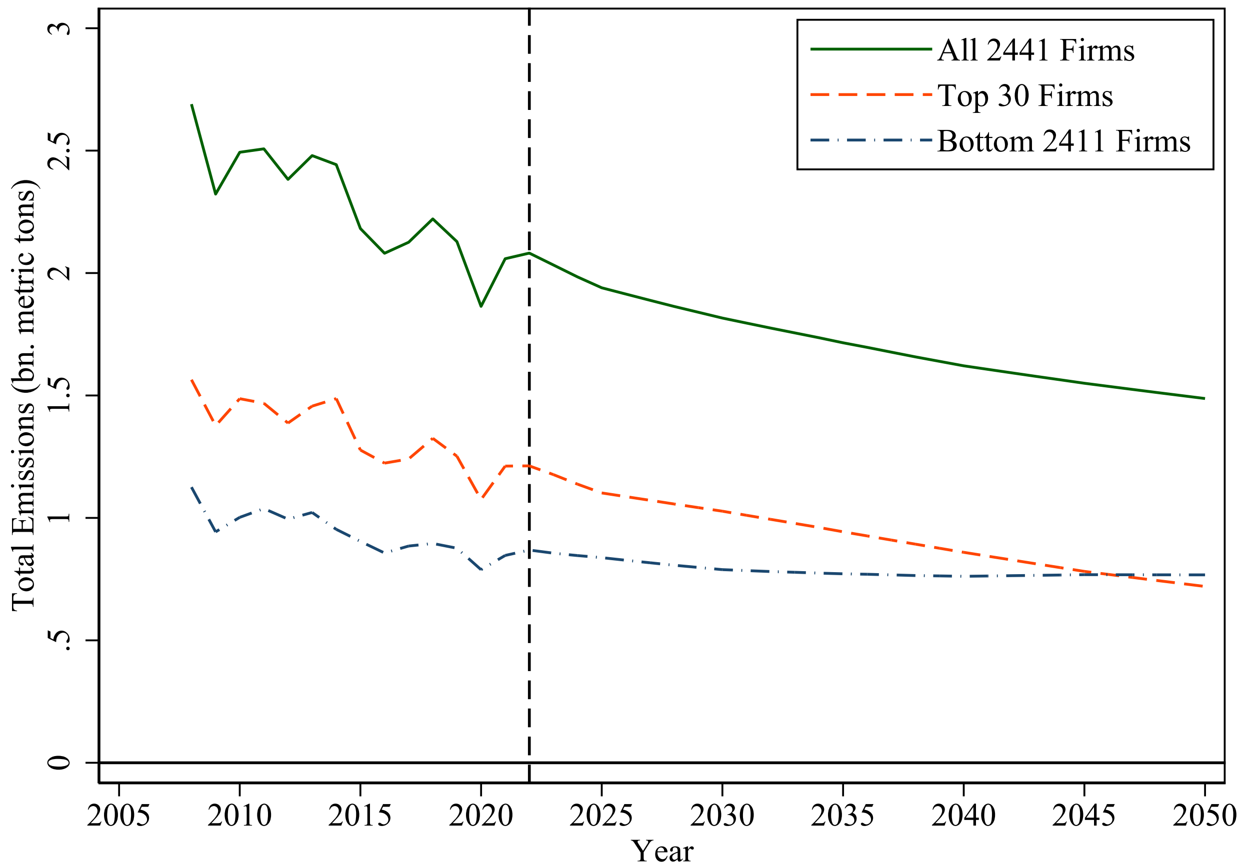

Figure 8. Past and future emissions. This figure shows past and future scope 1 emissions, using MSCI data on historical emissions and forecasts. The sample includes firms that have non-missing emission scope 1 forecasts for 2023 and historical emissions for 2022. We rank firms based on their emissions 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 $\leq$ 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.

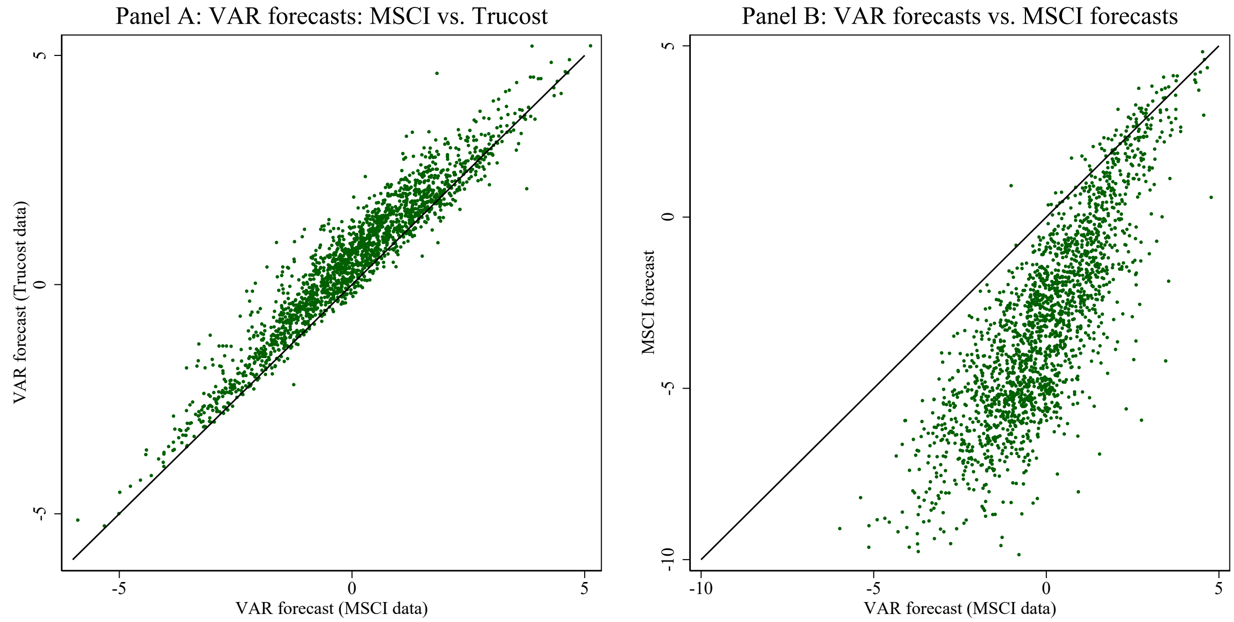

Figure 9. Comparing carbon-burden estimates. This figure shows scatter plots of firms' ratios of carbon burden to market cap, comparing estimates from one method to another. In Panel A, both dimensions use VAR-based forecasts, but the y-axis is based on historical Trucost data and the x-axis is based 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.

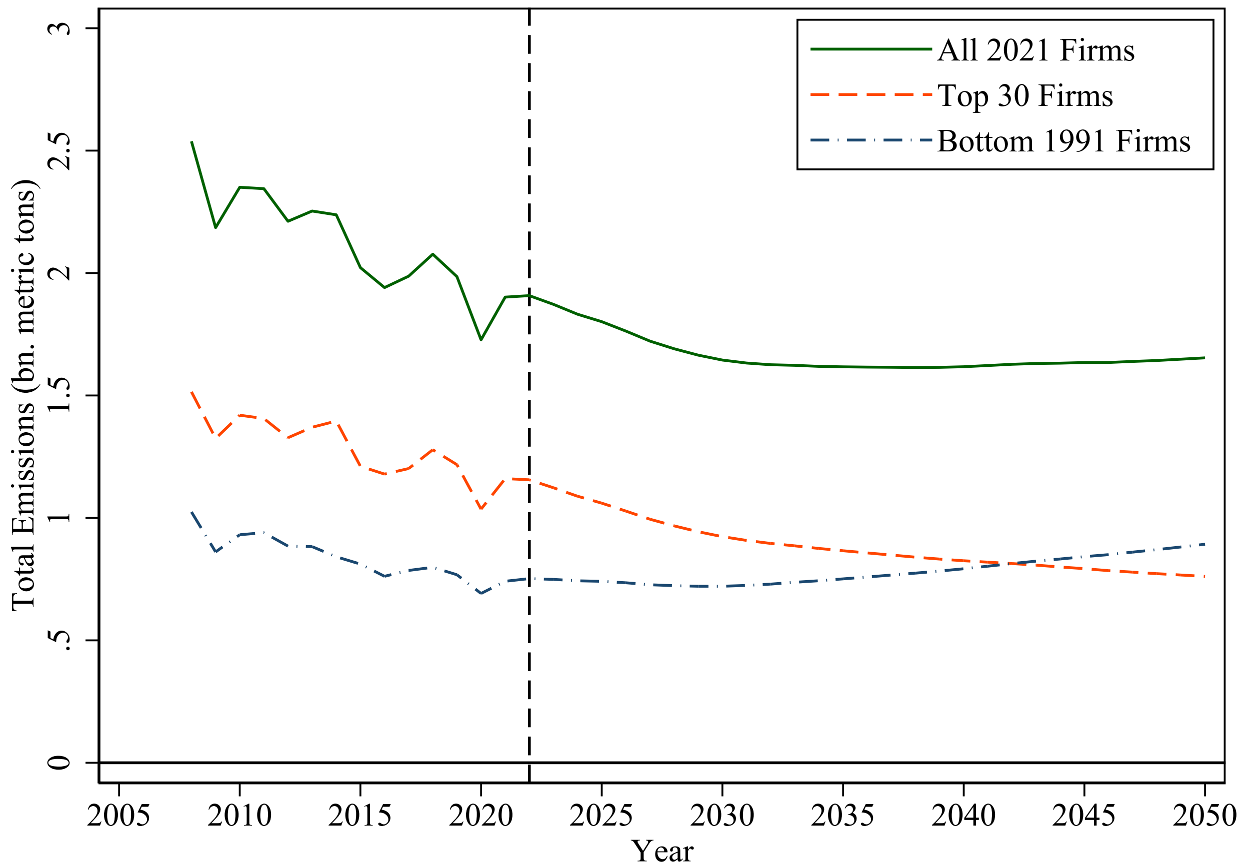

Figure 10. Past and future emissions from VAR forecasts. This figure is the same as Figure 8, except future emissions are from VAR-based forecasts. Historical scope 1 emissions are from MSCI, and the VAR is estimated using historical MSCI scope 1 data. The sample includes firms for which we can forecast emissions after 2022 using the VAR model.

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 each industry's ratio of carbon burden to market cap and ratio of future to present carbon burden. 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	2.10	22.74	22.28	22.87
Durbl	0.02	0.05	3.52	19.84	19.32	24.68
Manuf	0.26	0.38	5.60	22.33	22.41	26.22
Enrgy	1.06	1.19	20.48	22.05	22.12	28.47
Chems	0.45	0.60	3.00	22.41	22.08	25.52
BusEq	0.00	0.01	0.25	19.19	18.76	20.16
Telcm	0.01	0.07	0.67	16.54	16.26	21.40
Utils	2.71	2.81	5.54	20.35	20.46	23.61
Shops	0.05	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	0.41	0.45	1.87	22.24	22.33	24.00
Panel B: All future years						
NoDur	0.36	0.53	6.23	68.48	66.47	67.81
Durbl	0.05	0.15	10.98	57.83	55.75	77.31
Manuf	0.74	1.09	17.53	64.09	64.89	82.33
Enrgy	2.95	3.33	65.90	61.34	61.84	92.19
Chems	1.23	1.63	9.17	60.96	59.91	78.25
BusEq	0.01	0.04	0.75	52.50	53.91	60.80
Telcm	0.04	0.21	2.10	48.96	47.56	66.46
Utils	6.95	7.25	15.92	52.95	53.59	68.66
Shops	0.14	0.26	6.06	80.01	76.12	89.53
Hlth	0.03	0.07	1.41	66.66	68.79	71.15
Money	0.01	0.04	16.61	72.01	61.97	83.82
Other	1.21	1.33	5.56	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. "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/Mktcap		Future/Present	
Panel A: Using the Fama-French 49 industries						
Scope 1	0.521	0.642	0.521	0.580	0.063	0.242
Scope 1+2	0.432	0.590	0.415	0.506	0.061	0.241
Scope 1+2+3	0.367	0.654	0.392	0.461	0.051	0.154
Panel B: Using the Fama-French 12 industries						
Scope 1	0.412	0.548	0.418	0.468	0.035	0.216
Scope 1+2	0.330	0.505	0.320	0.401	0.035	0.217
Scope 1+2+3	0.272	0.586	0.294	0.352	0.028	0.131
Industry FEs	Y	Y	Y	Y	Y	Y
Log(MktCap)		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. BE/ME 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. BE/ME, 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)	(5)	(6)
Log(Emissions)	-2.344	-2.960	-2.279	-2.688	-3.212	-2.460
	(-9.53)	(-12.34)	(-5.43)	(-8.88)	(-12.54)	(-5.14)
BE/ME	1.203	1.280	0.875	1.182	1.194	0.608
	(1.74)	(1.85)	(1.51)	(1.71)	(1.83)	(1.08)
Investment	2.065	2.756	2.838	2.290	2.909	2.669
	(1.27)	(1.79)	(1.99)	(1.59)	(2.09)	(2.10)
Climate Score	-7.340	-7.063	-5.614	-9.203	-8.860	-4.200
	(-3.39)	(-3.30)	(-4.51)	(-3.40)	(-3.26)	(-2.83)
Revenue Growth	-0.552	-1.012	-1.092	-1.014	-1.377	-1.184
	(-1.30)	(-1.86)	(-1.94)	(-2.33)	(-2.49)	(-2.07)
Constant	0.020	0.029	0.034	0.017	0.025	0.028
	(7.37)	(11.09)	(6.98)	(4.90)	(8.35)	(4.11)
Observations	2191	2213	2213	2191	2213	2213
Adjusted R^2	0.100	0.107	0.069	0.118	0.122	0.085
Scopes	1	1+2	1+2+3	1	1+2	1+2+3
Industry FE				Y	Y	Y

Table 8
Firms in the top 10% of emissions

Panel A reports numbers of firms and total 2023 emissions for the firms whose emissions in a given emission category are in the top 10% of U.S. firms. For the top 10%, Panel B reports percentage emission reductions from 2023 to 2030 implied by the Paris Agreement under alternative scenarios for the remaining 90% of firms. For firms within the top 10% that also have emission targets for 2023 and 2030, as identified by MSCI, Panel C reports the number of firms and their total 2023 emissions. For the same firms, Panel D reports properties of their targeted emission reductions from 2023 to 2030, and Panel E reports properties of MSCI's forecasted emission reductions over the same period.

	Emission categories			
	Scope 1	Scope 2	Scope 3	All
Panel A. Top 10% of emitting firms				
Number of firms	285	285	285	285
Total emissions (mil. tons)	2,696	359	19,224	21,509
Percent of total emissions for all firms	96	79	85	83
Panel B. Percentage reductions for the top 10% implied by the Paris Agreement				
if other firms reduce emissions at the same rate	41	41	41	41
if other firms hold their emissions constant	43	52	48	49
if other firms reduce their emissions by 100%	38	26	31	29
if other firms increase their emissions by 100%	47	78	65	69
Panel C. Firms in the top 10% that also have emission targets				
Number of firms with emission targets	184	211	190	200
Total emissions of firms with targets (mil. tons)	2,056	284	14,635	16,539
Percent of total emissions for the top 10%	76	79	76	77
Panel D. The above firms' targeted percentage reductions				
Median targeted reduction	26	31	3	11
Average targeted reduction	31	36	16	18
Aggregate targeted reduction	28	33	8	10
Panel E. MSCI's forecasted percentage reductions for the same firms				
Median forecasted reduction	11	19	-1	3
Average forecasted reduction	16	21	8	9
Aggregate forecasted reduction	14	20	2	3

Table 9
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
	(342.10)	(849.62)	(292.86)	(426.81)
BE/ME	-0.029	-0.021	-0.037	-0.026
	(-3.98)	(-1.56)	(-3.63)	(-1.78)
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	-0.084
	(-2.10)	(-5.12)	(-2.63)	(-4.85)
Revenue Growth	-0.049	-0.102	-0.072	-0.138
	(-1.11)	(-2.54)	(-1.10)	(-2.03)
Constant	-0.099	-0.165	-0.119	-0.195
	(-2.82)	(-7.52)	(-3.66)	(-5.44)
Observations	12150	9820	7291	7291
Adjusted R^2	0.970	0.944	0.971	0.950

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.

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.

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

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.

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

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.

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

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.

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.

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.

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

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

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

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. Model

In Section 3.3, we present a simple framework that explains why the ratio of the aggregate carbon burden to total equity market value is high even though damages from emissions are only a small fraction of GDP. In that 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 in Section 3.3, total output is given by

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

but here we model expected output Y explicitly as

$$Y = K^\alpha L^{1-\alpha} \,, \tag{A.2}$$

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:

$$\mathcal{E}_t = fY \,. \tag{A.3}$$

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

$$\text{CB} = \frac{fY}{r} \,. \tag{A.4}$$

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

$$D_t = (r_K - \delta)K + \epsilon_t \,. \tag{A.5}$$

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:

$$M = \frac{D}{r_S} = \frac{r_K - \delta}{r_S} K \,, \tag{A.6}$$

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.6) pin down the marginal product of capital:

$$r_K = \delta + r_S \,. \tag{A.7}$$

Given equations (A.4), (A.6), and (A.7), the ratio of the carbon burden to market value is

$$\frac{\text{CB}}{M} = \frac{fY}{rK} = \frac{Y}{r_K K}\,\frac{r_K}{r}\,f = \frac{1}{\alpha}\,\frac{\delta + r_S}{r}\,f \,. \tag{A.8}$$

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.8) 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 in Section 3.3 and choosing a round value of $\delta = 10\%$, the value of the second term is 8, and equation (A.8) then implies CB$/M = 24f$. When $f = 4.7\%$, as in Section 3.3, we have CB$/M = 113\%$. If we increase δ slightly to 12%, we obtain CB$/M = 127\%$, which is close to our baseline estimated value of 131%.

Even though this model is richer than the model in Section 3.3, 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.8) 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.5)), 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 Section 3.3. 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. Filters applied to 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.

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.[22]

[22] 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.

A.3. 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.[23] 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 9, discussed in Section 5.2.2, 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 9, 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.1 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.

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

$$L_{n,t,s} = \frac{|\,\mathcal{C}_{n,t,s,MSCI} - \mathcal{C}_{n,t,s,Trucost}\,|}{\left(\mathcal{C}_{n,t,s,MSCI} + \mathcal{C}_{n,t,s,Trucost}\right)/2} \tag{A.9}$$

for each firm, year, and scope. Panel B of Table A.1 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.

[23]Emission disclosure may become mandatory soon. On March 6, 2024, the Securities and Exchange Commission (SEC) enacted a rule that for the first time will require company disclosures of material scope 1 and 2 GHG emissions, starting in 2026, conditional on overcoming legal challenges. Currently, emission disclosure is mandatory only at the facility level through the EPA's Greenhouse Gas Reporting Program, and only for sufficiently large emitters.

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. In the Appendix, we plot 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.

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

$$G_{n,t,s} = \left| \frac{\mathcal{C}_{n,t,s,MSCI} - \mathcal{C}_{n,t-1,s,MSCI}}{\mathcal{C}_{n,t-1,s,MSCI}} - \frac{\mathcal{C}_{n,t,s,Trucost} - \mathcal{C}_{n,t-1,s,Trucost}}{\mathcal{C}_{n,t-1,s,Trucost}} \right| \tag{A.10}$$

for each firm, year, and scope. Panel C of Table A.1 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 like 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.1 examines the cross-sectional relations between both characteristics and $L_{n,t,s}$, our discrepancy measure from equation (A.9). 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 Rajogopal (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.[24]

Figure A.1 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, 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.1 shows that the discrepancies are larger for firms that emit little and firms that do not disclose emissions.

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 the carbon tax rate of \$200 per ton, which equals the EPA's SCC in 2022 with a 2% discount rate.[25] 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, $\mathrm{CT}_{Trucost}$, and report the absolute difference scaled by the firm's 2022 operating profit: $|\mathrm{CT}_{MSCI} - \mathrm{CT}_{Trucost}|/\mathrm{Profit}$. Table A.2 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.

[24] 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.

[25] 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.2 should be multiplied by half (= 100/200).

Panel A of Table A.2 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.2 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.

To summarize, we find that the MSCI-Trucost discrepancies in measured scope 1 and scope 1+2 emissions are modest for most firms, but they are substantial at the high end, especially for large emitters. The discrepancies in scope 1+2+3 emissions are large for most firms. For all scopes, the discrepancies are smaller for firms that disclose their emissions, but they are substantial even for such firms.

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.4. Additional tables and figures

Table A.1: 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

Table A.2: 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, 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

Table A.3: 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.

	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.4: 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%	0.104	0.062	0.656	0.445
2.0%	0.133	0.079	0.770	0.559
1.5%	0.185	0.107	0.873	0.665

Table A.5: 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%	0.065	0.029	0.365	0.249
2.0%	0.087	0.044	0.500	0.311
1.5%	0.102	0.073	0.632	0.390

Table A.6: Version of Table 7 dropping observations with 1% growth rate

The sample for scope 1 (1+2) (1+2+3) includes firms for which the MSCI scope 1 (1 or 2) (1, 2, or 3) forecasted growth rate is not equal to 0.0100, after rounding.

	(1)	(2)	(3)	(4)	(5)	(6)
Log(Emissions)	-1.367	-1.810	-1.597	-1.523	-2.070	-2.497
	(-3.93)	(-4.72)	(-2.07)	(-2.73)	(-4.09)	(-2.82)
BE/ME	1.666	1.736	0.254	1.876	1.663	-1.730
	(0.76)	(0.98)	(0.09)	(0.78)	(0.82)	(-0.47)
Investment	11.171	10.520	11.568	11.237	10.355	10.266
	(1.47)	(1.47)	(0.81)	(1.55)	(1.54)	(0.71)
Climate Score	-16.770	-16.869	-25.999	-18.279	-18.131	-25.124
	(-5.11)	(-5.15)	(-5.07)	(-5.38)	(-5.06)	(-4.44)
Revenue Growth	5.396	8.050	12.759	4.168	6.687	16.054
	(2.37)	(1.05)	(1.06)	(1.62)	(0.80)	(1.09)
Constant	-0.010	-0.004	-0.003	-0.019	-0.011	0.005
	(-2.08)	(-0.72)	(-0.23)	(-2.83)	(-1.71)	(0.35)
Observations	612	597	318	612	597	318
Adjusted R^2	0.024	0.022	0.014	0.029	0.028	0.005
Scopes	1	1+2	1+2+3	1	1+2	1+2+3
Industry FE				Y	Y	Y

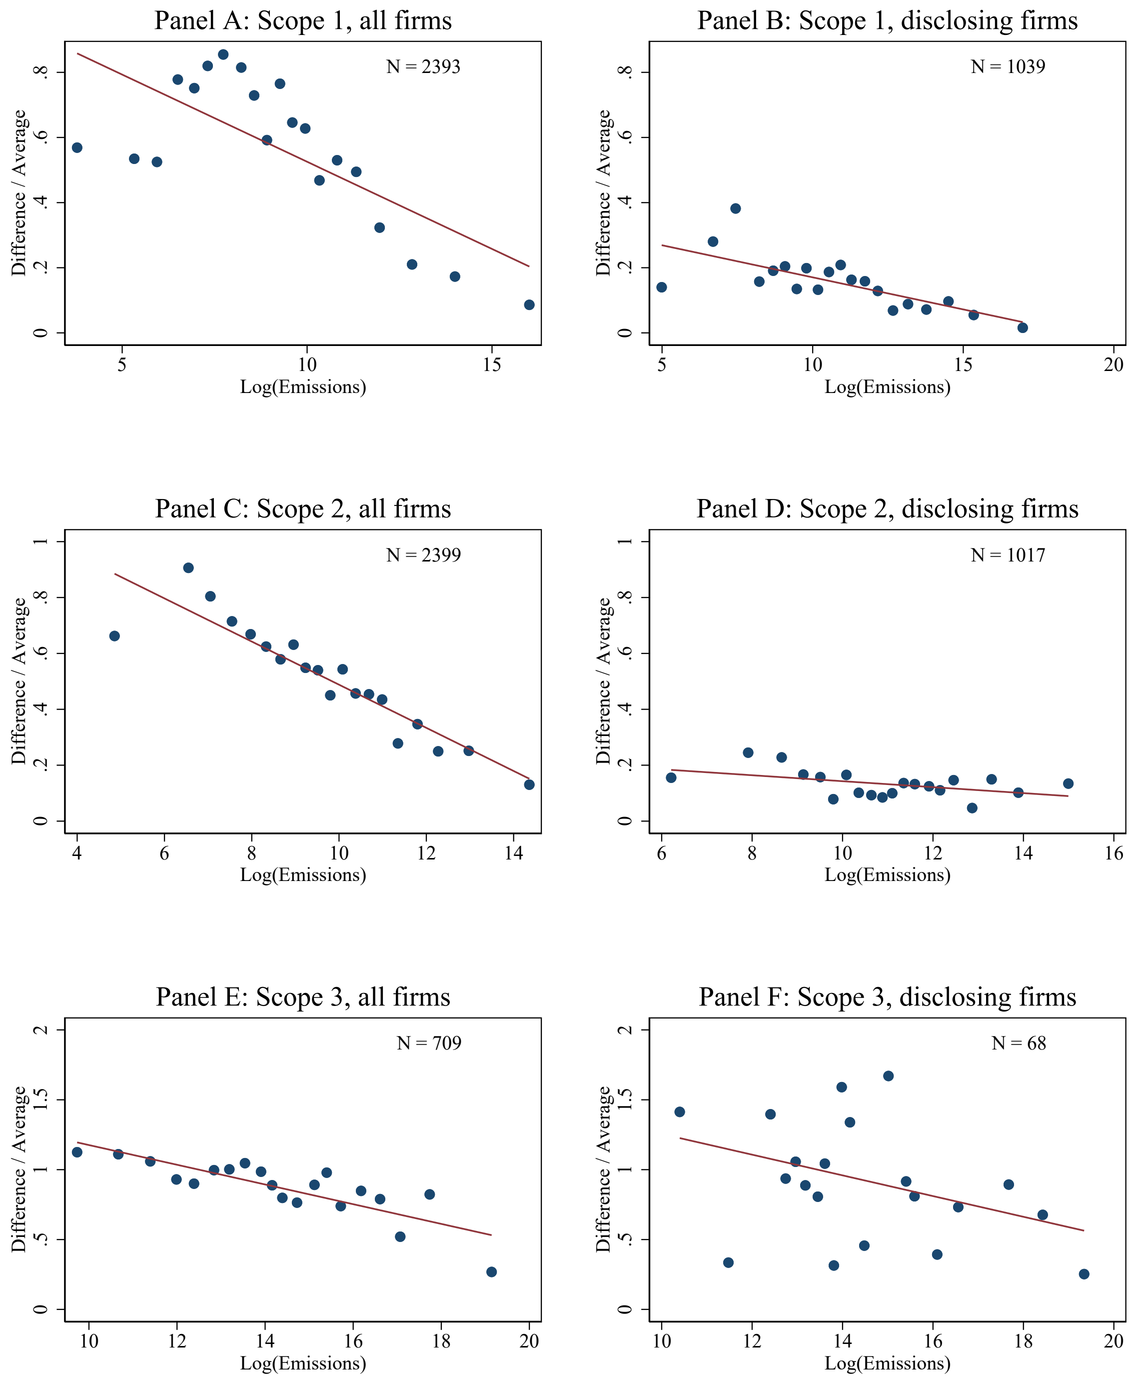

Figure A.1. Discrepancies between Trucost and MSCI. In these binscatter plots, the x-axis denotes the log of the firm's emissions (equal to the average of MSCI and Trucost emissions), and the y-axis denotes the firm's ratio of (i) the absolute difference between MSCI and Trucost emissions to (ii) the 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.

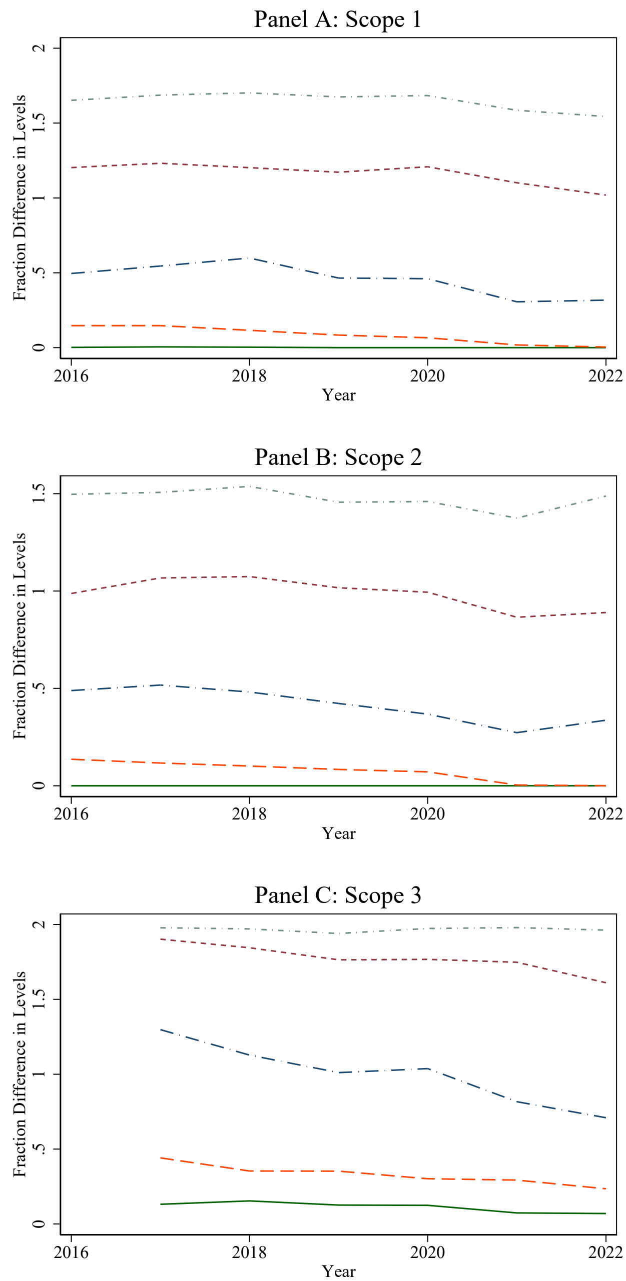

Figure A.2. MSCI-Trucost discrepancies: The time series. This figure plots cross-sectional percentiles each year of the fraction discrepancy between Trucost and MSCI for each scope. From bottom to top, the lines represent the 10th, 25th, 50th, 75th, and 90th percentiles of the fraction discrepancy. For 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.

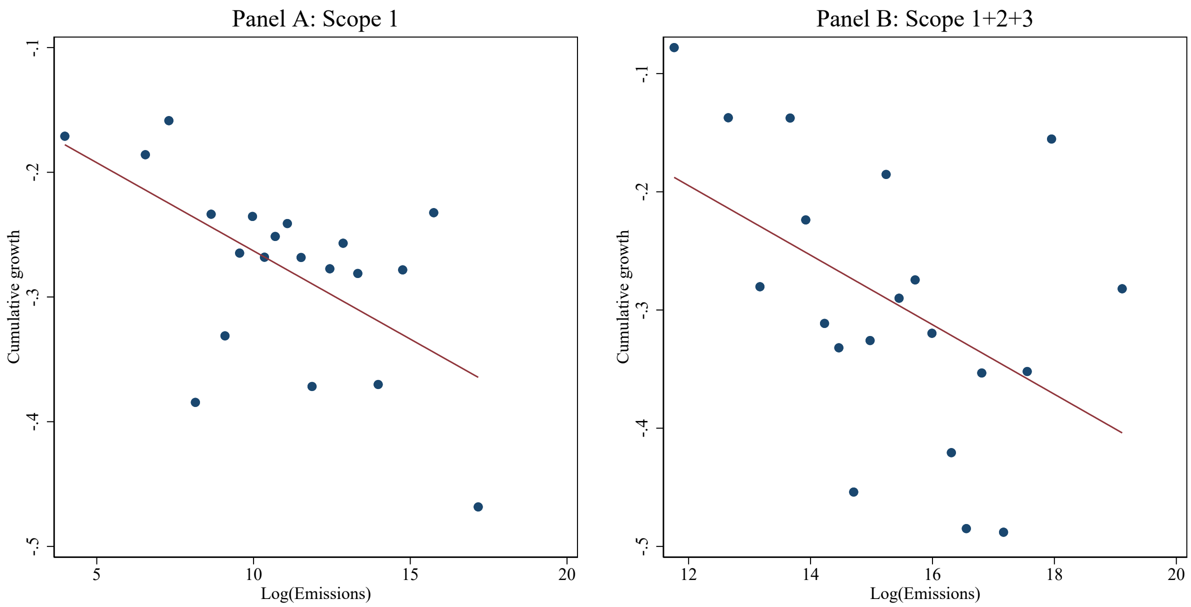

Figure A.3. Version of Figure 7 dropping observations with 1% growth rate. The sample for scope 1 (1+2+3) includes firms for which the MSCI scope 1 (1, 2, or 3) forecasted growth rate is not equal to 0.0100, after rounding. Panel A (B) includes 696 (353) firms in total.

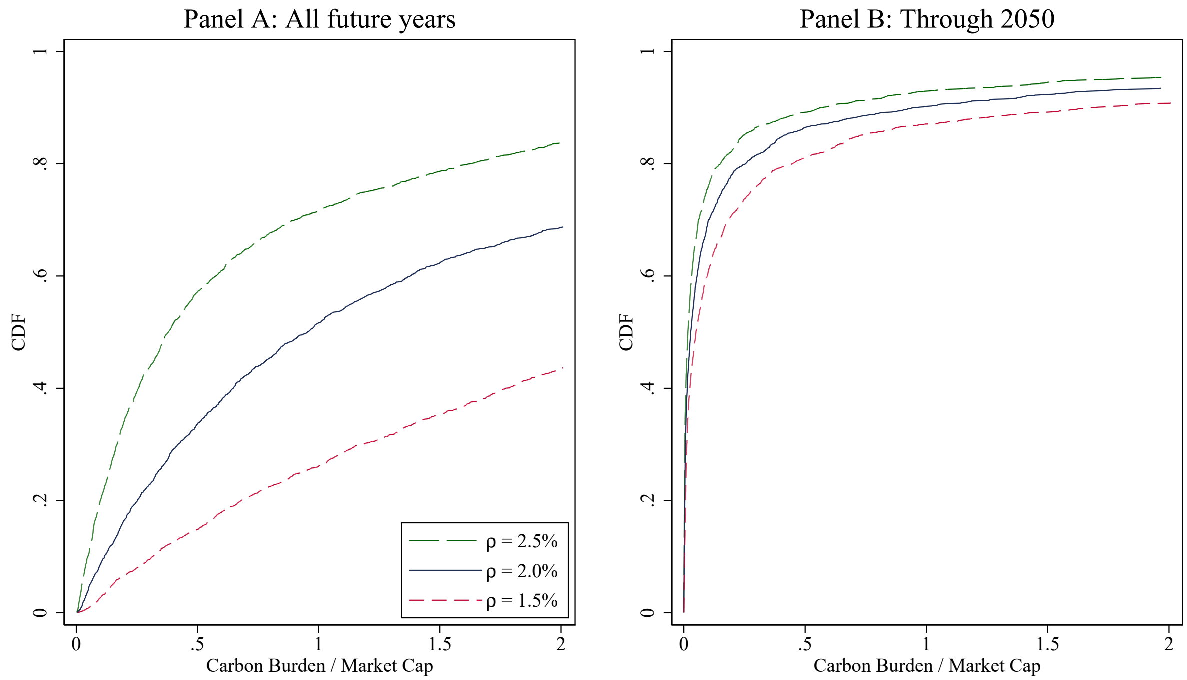

Figure A.4. Distribution of carbon burden / market cap from VAR approach. This figure shows CDFs of carbon burden / market cap, computed from the VAR approach with historical scope 1 emissions data from MSCI. Carbon burden and market cap are both measured as of the end of 2022. The CDFs weight each firm 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%	0.283	0.070
2.0%	0.483	0.097
1.5%	0.739	0.129

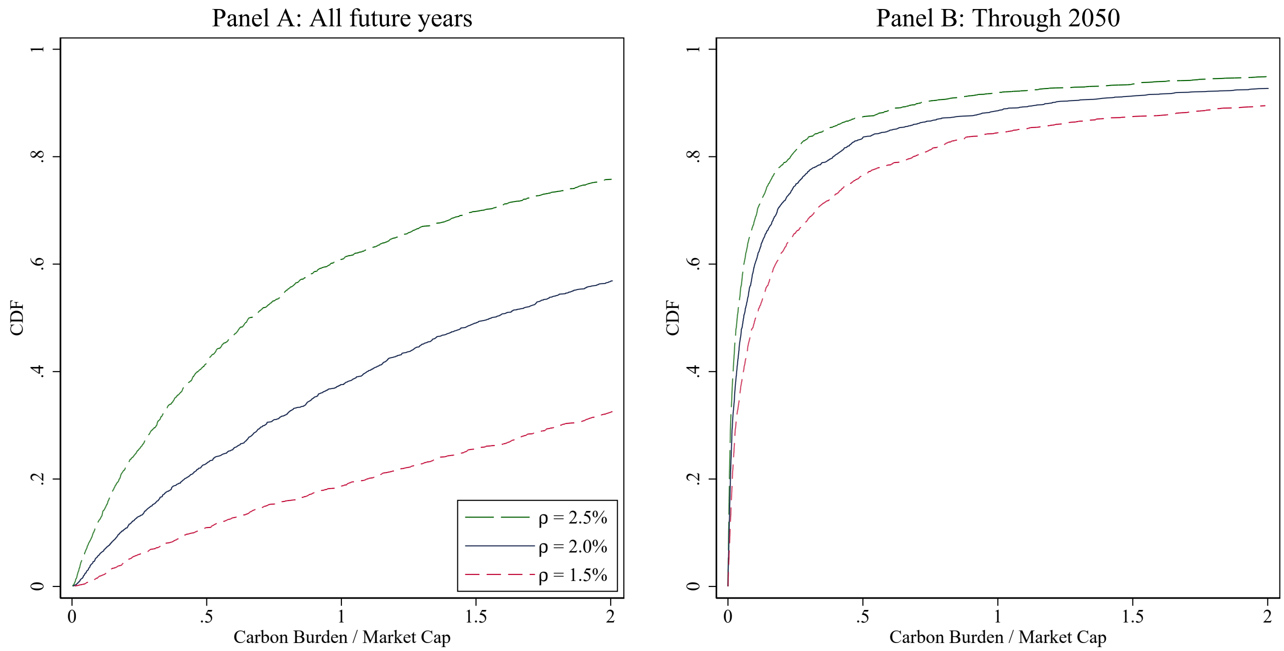

Figure A.5. Distribution of carbon burden / market cap from the VAR approach with Trucost data. This figure is the same as Figure A.4 but shows results based on historical Trucost emissions data.

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%	0.624	0.114
1.5%	0.813	0.155

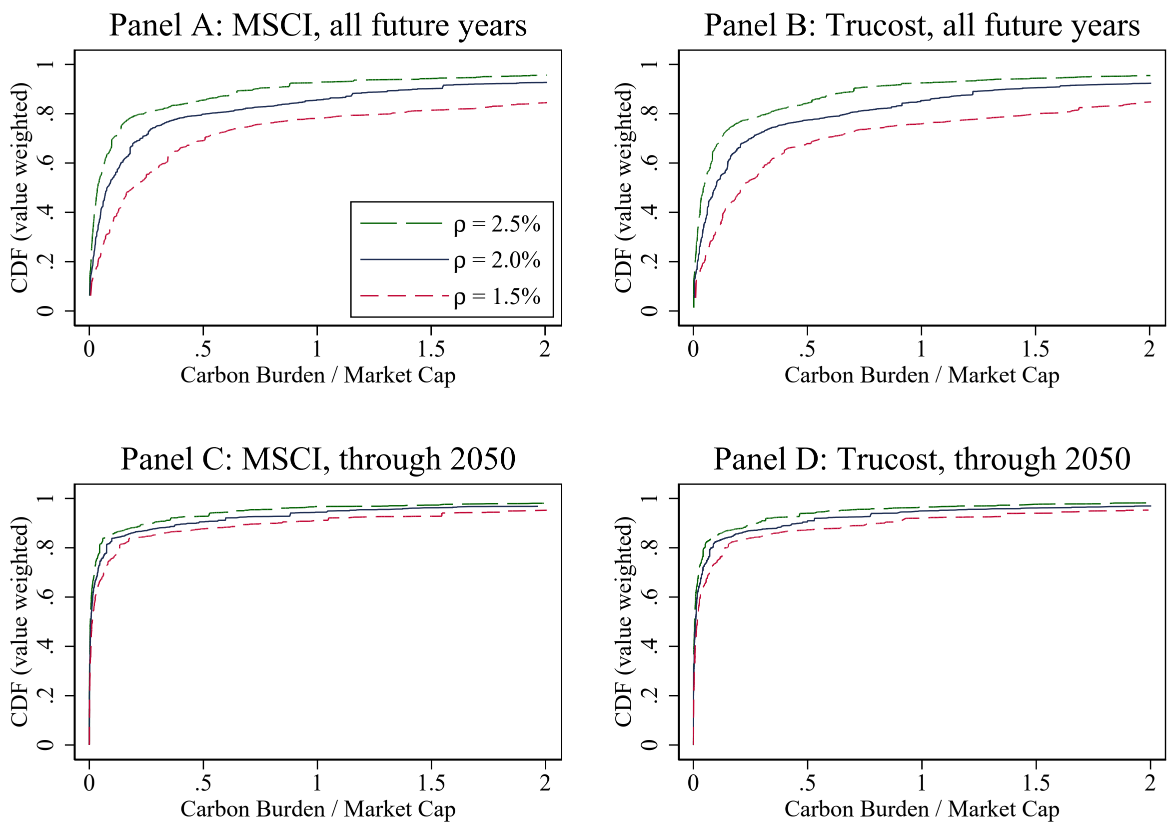

Figure A.6. Value-weighted distribution of carbon burden / market cap from the VAR approach. This figure shows cumulative distribution functions (CDFs) of H = carbon burden / market cap, computed in the cross section of firms in 2022. We use carbon burden computed using the VAR approach. The VAR model 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 H below the x-axis value.

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

	All future years		Through 2050	
Discount Rate	MSCI (Panel A)	Trucost (Panel B)	MSCI (Panel C)	Trucost (Panel D)
2.5%	0.074	0.075	0.033	0.036
2.0%	0.144	0.146	0.056	0.050
1.5%	0.219	0.241	0.089	0.080