Through the Looking Glass: Leveraging Machine Learning to Price Corporate Carbon Footprints*

First version: April 2024
This version: Nov 2024

Haoxu Wang[†], Ahmed Syalabi Seet[‡], Zhou Zhou[§], Johan Sulaeman[¶]

Abstract

We estimate the pricing of corporate Greenhouse Gas (GHG) footprints in the equity market using machine learning. The valuation model based on gradient boosted regression trees (GBRT) substantially outperforms conventional linear valuation models, yielding estimates of equity valuations that are reasonably close to the observed market data for a comprehensive global dataset constructed by Jensen, Kelly, and Pedersen (2023), which includes 250+ accounting-based features of companies listed publicly in 70+ exchanges around the world. The model allows us to perform valuation attributions of fundamental features, revealing the higher relevance of corporate carbon footprints for the valuation of energy sector firms. Employing recent advances in explainable artificial intelligence (XAI), we document heterogeneous pricing patterns across regions, industry sectors, and emission levels, and uncovers investors' distinct views on emissions generated directly by the firm (Scope 1) vs. indirectly via energy purchases (Scope 2). Our valuation framework generates equity-market-imputed carbon prices ranging from US$ 30 to 150 per tonne of GHG Scope 1 emissions in recent years, highlighting the increased importance of corporate emission in firm valuation.

Keywords: Fundamental Analysis; Carbon Footprints; Carbon Pricing; Machine Learning; Explainable AI

JEL Classification: G12; G15; Q51

*We thank seminar participants at Sungkyunkwan University and National University of Singapore for helpful comments and suggestions. The research is supported by Sustainable and Green Finance Institute (SGFIN ; WBS A-0006413-00-00) and Asian Institute of Digital Finance (AIDF) at NUS

[†]Sustainable and Green Finance Institute (SGFIN), National University of Singapore (NUS). haoxuw@nus.edu.sg

[‡]SGFIN, NUS. syalabi@nus.edu.sg

[§]Asian Institute of Digital Finance (AIDF), NUS. zhou.zhou@u.nus.edu

[¶]NUS Business School; SGFIN; AIDF, NUS. sulaeman@nus.edu.sg

Through the Looking Glass: Leveraging Machine Learning to Price Corporate Carbon Footprints

Abstract

We estimate the pricing of corporate Greenhouse Gas (GHG) footprints in the equity market using machine learning. The valuation model based on gradient boosted regression trees (GBRT) substantially outperforms conventional linear valuation models, yielding estimates of equity valuations that are reasonably close to the observed market data for a comprehensive global dataset constructed by Jensen, Kelly, and Pedersen (2023), which includes 250+ accounting-based features of companies listed publicly in 70+ exchanges around the world. The model allows us to perform valuation attributions of fundamental features, revealing the higher relevance of corporate carbon footprints for the valuation of energy sector firms. Employing recent advances in explainable artificial intelligence (XAI), we document heterogeneous pricing patterns across regions, industry sectors, and emission levels, and uncovers investors' distinct views on emissions generated directly by the firm (Scope 1) vs. indirectly via energy purchases (Scope 2). Our valuation framework generates equity-market-imputed carbon prices ranging from US$ 30 to 150 per tonne of GHG Scope 1 emissions in recent years, highlighting the increased importance of corporate emission in firm valuation.

Keywords: Fundamental Analysis; Carbon Footprints; Carbon Pricing; Machine Learning; Explainable AI

JEL Classification: G12; G15; Q51

1 Introduction

Climate change has become one of the most critical global issues. With stakeholders clamoring for more information and concrete actions regarding the impacts of commercial enterprises on climate change as well as their risk exposures, corporate leaders and investors urgently need a clearer understanding of the financial implications of corporate climate-related features. These implications include both economic impacts of climate change on corporate operations as well as potential impacts of regulatory changes – resulting from countries' commitments to reducing Greenhouse Gases (GHG) emissions – on firms' future cash flows.

One approach to quantify these impacts is through carbon pricing mechanisms, such as "carbon taxes" imposed by governments on corporate emissions and "carbon prices" prevailing in regulated carbon trading schemes such as the European Union (EU) Emissions Trading Scheme (ETS) and Chinese National ETS. Extant studies have examined these schemes, with a focus on the pricing of current emissions associated with corporate activities.[1] In this paper, we extend this literature by examining how the *equity market* prices corporate carbon emissions of publicly listed companies around the world.

Stock prices reflect many factors that are relevant to market participants, including information regarding discount rates and future cash flows. To disentangle the relevance of these factors, we develop a comprehensive fundamental-based valuation model that generates robust and accurate estimates of equity valuation. Using this model, we perform model-based attribution analyses to quantify the relevance of corporate carbon emissions (relative to other factors, e.g., profitability, investments, capital structure) in equity market valuation. As an additional feature, the model allows us to generate estimates of equity-market-imputed prices of carbon emissions for publicly listed firms in various markets and sectors over various time periods.

[1] Zhang and Wei (2010) and Weng and Xu (2018) review the studies on the two carbon markets, respectively.

Leveraging on recent developments in artificial intelligence, we develop a machine-learning-powered fundamental analysis model capable of uncovering complex and nonlinear relations between corporate fundamentals and firm valuation. This model is based on gradient-boosted regression trees (GBRT), which can (1) accommodate an extensive set of fundamental variables and (2) harness complex and nonlinear explanatory patterns across these variables for firm valuation.

We incorporate 259 accounting-based features provided in Jensen, Kelly, and Pedersen (2023) in the GBRT-powered valuation model.[2] We then add corporate emission features into the model to uncover GHG emissions' *unique* contribution to stock market valuation, beyond its indirect contribution captured by the accounting-based features. Additionally, our valuation model also accounts for fixed effects including time, country (exchange), and industry sectors.

Our machine-learning-powered valuation models deliver highly robust explanatory power for market capitalization (MCAP) and market-to-book equity ratios (M2B) of publicly listed companies in 90+ exchanges around the world. To facilitate comparisons with more conventional linear-regression-based valuation models (OLS models, henceforth), we calculate the GBRT model's R-squared (R^2) values. We find that the R^2s of the contemporaneous GBRT models are over 95% for MCAP and over 90% for M2B, much higher than corresponding OLS models incorporating the same comprehensive set of corporate features and fixed effects. These substantial improvements in explanatory power illustrate the importance of the GBRT models' ability to capture complex and nonlinear relationships between corporate features and equity valuation.[3]

The GBRT model estimates' deviations from the observed market values are typically less than 15% of the observed values, which are about half of the typical deviations of

[2]For the past 40 years, practitioners and academics have proposed numerous accounting measures that are indicative of either current stock prices (or valuation multiples) or future stock returns. Jensen, Kelly, and Pedersen (2023) confirm that many of these are also relevant in global stock markets.

[3]The GBRT-based valuation model proposed in the current study, which employs a comprehensive set of explanatory variables, also outperforms the GBRT-based valuation model developed by Geertsema and Lu (2023) that considers fewer accounting features.

OLS models' estimates using comparable feature sets. With the GBRT valuation model providing a more accurate reflection of equity market participants' evaluations of firm features, it would facilitate more precise valuation attributions to these features, including GHG emission information.[4]

Our valuation attribution analysis focuses on global firms that disclosed their GHG emissions during the period from 2003 to 2022. To capture potential changes in investors' views on climate change and sustainability following the Paris Agreement (December 2015) and potential shifts in sentiments due to the COVID-19 pandemic and the Ukrainian crisis, we split the sample into three periods: 2003–2015, 2016–2019, and 2020–2022. We train a separate GBRT model for each of the three periods. We conduct subsample analysis for each period at the regional or sectoral level. The regional classification that we consider is (1) the U.S., (2) Europe, and (3) Asia. We also perform a separate analysis on carbon-intensive firms (i.e., the "Energy+" sector), which includes fossil fuel producers and electricity generators.

In line with prior studies, we initially examine the firm's direct GHG emission (i.e., Scope 1) as well as the firm's indirect emissions (i.e., Scope 2) coming from purchase of electricity, steam, heat, and cooling.[5] Since firms' emission information is released at a lag relative to the end of the corresponding accounting year (and typically also at a lag relative to the financial statements), we conservatively assume that such information is publicly observable 12 months after the accounting year end to avoid the potential look-ahead bias highlighted in Zhang (2024). Our workhorse valuation models employ the

[4]In particular, we find that our benchmark linear regression model that incorporates the comprehensive set of corporate features performs better in explaining stock valuation than the corresponding regression models in Bolton and Kacperczyk (2023) and Pedersen, Fitzgibbons, and Pomorski (2021) that utilize smaller sets of corporate features, highlighting the importance of including as many corporate characteristics as possible for a valuation attribution analysis.

[5]Similar to previous studies (Bolton and Kacperczyk 2021a, 2023; Zhang 2024; Pedersen, Fitzgibbons, and Pomorski 2021), we source firm-level emission data from Trucost. In addition, we also incorporate emission data from Bloomberg and Refinitiv when the reported emission information is not available in Trucost. In our main analyses, we drop all firm-year observations with missing reported GHG data, including firm-year observations with "estimated" GHG values in Trucost. These firm-year observations are included in our robustness checks, in which we observe qualitatively similar results to our main analyses.

emission intensity (i.e., emissions-to-sales ratio) for Scopes 1 and 2 separately to proxy for corporate carbon footprints, similar to recent works by Pedersen, Fitzgibbons, and Pomorski (2021) and Zhang (2024).

To facilitate attribution analysis, we employ recent advances in explainable artificial intelligence (XAI). First, we use the SHapley Additive exPlanations (SHAP) measure proposed by Lundberg and Lee (2017) to examine the contribution of individual fundamental features to model estimates.[6] The SHAP analysis indicates that profitability measures (e.g., return-on-equity) and dividend payout (i.e., dividends-to-asset ratio) play dominant roles in explaining M2B ratios. Crucially for our objective, corporate GHG Scope 1 emission intensity (i.e., emissions-to-sales ratio) exhibits persistent value relevance over time, ranking in the top 10 important features (out of 261 features) in all three subsample periods. In contrast, GHG Scope 2 emission intensity ranks between 25 to 50, indicating that stock market participants view corporate direct emissions (i.e., Scope 1) differently from their indirect emissions (i.e., Scope 2) in corporate valuation analyses.

Second, we employ the Accumulative Local Effects (ALE) method proposed by Apley and Zhu (2020). While the SHAP measure is quite useful to describe feature importance – akin to analysis of "marginal" R^2 in conventional models, it is direction-silent and therefore unable to identify the direction (i.e., the sign) of each individual feature's effect on model estimates.[7] To uncover the directional effects of carbon emissions on the M2B or MCAP estimates generated by our valuation model, we employ the ALE method, which is designed to isolate the effect of individual features on the model estimates, and provides us with plots that are useful to visualize the (potentially non-linear) impacts of individual

[6]The SHAP measure is a cooperative game theory-based tool for analyzing feature importance in machine learning models. Similar to Geertsema and Lu (2023), we calculate the normalized SHAP measures to perform comparative analysis across subsamples.

[7]The SHAP measure only allow us to observe the direction of each effect for the estimate associated with each data point.

emission variables after controlling for other fundamental variables.[8]

The ALE plots generated from our valuation models exhibit pronounced nonlinear relationships between M2B and GHG Scope 1 emission intensity, with slopes varying across regions and firms with different levels of emission intensity. We also observe that the relationship patterns vary over time. In particular, the ALE curve for the effects of Scope 1 emission on M2B is generally *convex* in all regions with the *negative* slopes more pronounced after the Paris Agreement, consistent with excessive emissions from the operations directly controlled by the firm being regarded as a negative signal by the market.

The post-Paris Agreement carbon pricing patterns in the U.S. are stronger from those in Asia and Europe. Interestingly, the ALE curve for U.S. stocks exhibits positive slopes for firms with low Scope 1 emission intensity – indicating negative carbon pricing for small emitters, but we do not observe this in the sample of carbon-intensive firms in the "Energy+" sector. We therefore conjecture that the inconsistencies in recent studies examining carbon pricing in the equity market that rely on panel linear regression models (e.g., Bolton and Kacperczyk 2021b, 2023; Aswani, Raghunandan, and Rajgopal 2024; Zhang 2024) are likely manifestations of these complex and nonlinear relationship patterns, particularly for U.S. stocks.

We further examine how Scope 1 GHG emission *levels* directly affects market values of firms by generating another set of ALE plots. Consistent with the patterns for M2B and direct emission intensity discussed above, these ALE plots show that the levels of direct emissions (i.e., Scope 1) are negatively associated with MCAP, especially following the Paris Agreement.

[8]Within a certain interval for the feature of interest, the ALE method first calculates the effect of the feature at the observation level by replacing the value of the feature with the corresponding values of both the right and left interval endpoints while keeping the values of other features constant. Then, it calculates the mean of these effects across all observations within the interval to obtain the estimate of local effects in the interval, which are then accumulated over the range of the values of the feature to produce ALE plots. If ALE is to be performed on a linear regression model, the plots would be straight lines with slopes corresponding to parameter estimates from the regression.

Third, to further quantify the effects of carbon emissions on corporate valuation, we propose a numerical method – built upon the ALE method – to estimate equity-market-imputed emission prices at the regional level for each subsample period. We find that in all three regions, the imputed prices for Scope 1 emission are always positive and rising over time. The imputed emission prices for the most recent (2020–2022) period in the U.S., Europe and Asia are estimated to be US$ 148.37, 53.94 and 34.16 per tonne of CO_2 equivalent (tCO_2e) of emissions, respectively. Remarkably, the imputed Scope 1 emission prices for firms in the "Energy+" sector for the latest period (25.83, 14.38, and 7.54 US$/t$CO_2$e, respectively in the three regions) are much lower than the corresponding prices for firms in other industries (203.06, 98.34, and 42.48 US$/t$CO_2$e, respectively). This suggests that the market penalizes the excessive direct emissions of "easier-to-abate" firms, such as airline companies, much more heavily than those of energy firms whose entire business models depend entirely on economic activities that produce substantial GHG emissions.

In the last part of the study, we explore a potential channel through which direct emissions are negatively associated with corporate valuation. We examine whether direct emission intensity is associated with *future* firm profitability, after controlling for other fundamental aspects of the firm, including current profitability. We document that direct emission intensity exhibits substantial *predictive* power for 5-year-ahead return on equity (ROE): relatively carbon-intensive firms exhibit lower future profitability. Combining this with the negative pricing effects of emission intensity indicate that equity market participants had expected the negative cash flow shocks that are uniquely predictable by firms' direct carbon emissions, and these expectations were impounded into stock prices in the global equity markets.

Our study makes following contributions to the literature. First, we propose a robust equity valuation model that can efficiently harness the complex and nonlinear explanatory patterns of fundamental information for firm valuation. Second, we evaluate the importance of corporate environmental features relative to conventional fundamental information in capturing firm value. Relying on an extensive global sample of firms disclosing their GHG emissions, we uncover nonlinear and complex effects of carbon emissions on corporate valuation, highlighting similarities and differences in the pricing of corporate carbon footprints across stock markets in different regions. Third, we propose a method to impute the corporate GHG emission prices implied by the equity market, and reveal the secctoral and regional variations in carbon pricing patterns.

2 Related Literature

Our paper contributes to the literature on climate change and equity market. Early studies on the value relevance of carbon emissions published in the 2010's typically examine a limited sample of GHG reporters in a single region, and employ linear regression models with a limited set of control variables, which may not be sufficient to isolate the value effects of carbon emissions from those of conventional fundamental variables.[9] More importantly, these studies focus on the firm-value impacts of carbon emissions prior to the Paris Agreement. The current study examines the complex and potentially nonlinear effects of corporate carbon emissions on equity valuation around the world, documenting cross-industry and regional variations in the effects as well as their time-series dynamics up to the most recent period.

Our results offer novel perspectives on the ongoing debate about asset pricing implications of carbon emissions. Bolton and Kacperczyk (2021a, 2023) document that stocks of carbon-intensive firms tend to earn higher abnormal returns than less carbon-intensive counterparts not only in the U.S. but also in other countries around the world. They attribute this return spread to carbon-intensive firms' larger exposure to carbon-transition

[9] These early studies include Matsumura, Prakash, and Vera-Muñoz (2014), who examine the S&P 500 firms that voluntarily report carbon emissions during the 2006-2008 period; Chapple, Clarkson, and Gold (2013), who study around 200 European listed firms over the period 2006-2009; Chapple, Clarkson, and Gold (2013), who use a sample of 58 Australian listed firms in 2007; and Lee, Min, and Yook (2015), who examine 362 Japanese listed firms between 2003 and 2010.

risk. However, this carbon return premium is not observed in several contemporaneous studies (e.g., Garveya, Iyera, and Nashb 2018; In, Park, and Monk 2019; Görgen et al. 2020; Duan, Li, and Wen 2023; Aswani, Raghunandan, and Rajgopal 2024; Zhang 2024), some of which even document and rationalize the existence of a carbon discount. We conjecture that these inconsistent results can be driven by the pronounced nonlinear, complex, and time-varying associations between carbon emissions and firm value that we document in this study.

In addition to carbon emissions, prior studies have also examined the equity market pricing of other corporate environmental features such as sulfur dioxide emissions (Hughes 2000; Johnston, Sefcik, and Soderstrom 2008), pollution (Hsu, Li, and Tsou 2023; Cormier and Magnan 1997; Connors, Johnston, and Gao 2013), and biodiversity risk (Giglio et al. 2023). The proposed framework in this study can also be applied on these features, although measurements of these environmental features would be quite problematic for global studies such as this one.

Our paper builds on a voluminous literature that aims to develop equity valuation models that incorporate fundamental variables, primarily from accounting items in the financial statements (Bhojraj and Lee 2002; Rhodes–Kropf, Robinson, and Viswanathan 2005; Bartram and Grinblatt 2018, 2021; Geertsema and Lu 2023). Most studies on this topic rely on linear regressions of either market values (market capitalization or stock price) or valuation multiples (e.g., market-to-book equity ratio, M2B) on firm fundamentals. A notable exception is Geertsema and Lu (2023), who apply a machine learning approach to fundamental analysis. Our framework differs from the machine-learning model proposed by Geertsema and Lu (2023) in two key aspects. First, our study incorporates a comprehensive sample of global stocks, similar to Bartram and Grinblatt (2021), while their analysis is limited to U.S. stocks. Second, our framework incorporates a more extensive set of conventional fundamental variables along with carbon emission features. Beyond the methodological differences, this study focuses on valuation attributions to carbon emissions, whereas their paper aims to examine the predictive power of the stocks' valuation gap suggested by their model for future returns and the associated investment implications, in line with Bartram and Grinblatt (2018, 2021).

This study also contributes to the growing body of literature that applies machine learning techniques to finance and accounting research. This strand of literature includes expected return estimation for stocks (e.g., Kelly, Pruitt, and Su 2019; Gu, Kelly, and Xiu 2020; Freyberger, Neuhierl, and Weber 2020; Gu, Kelly, and Xiu 2021; Chen, Pelger, and Zhu 2024) and bonds (Bianchi, Büchner, and Tamoni 2021), fund performance analysis (Kaniel et al. 2023; DeMiguel et al. 2023), firm director selection (Erel et al. 2021), earnings prediction (Chen et al. 2022), firm quality measurement (Chen, Ke, and Zhao 2024), accounting fraud detection (Bao et al. 2020), and value relevance of accounting information (Geertsema and Lu 2023; Barth, Li, and McClure 2023).

An important contribution of the current study to this growing literature is the implementation of recent advances in the explainable artificial intelligence (XAI) field. The recent development of XAI in computer science has resulted in its emergence in finance (e.g., the pioneering work conducted by Cong, Liang, and Zhang 2019; Cong et al. 2021, 2023), enhancing the transparency and interpretability of outputs produced by machine learning models. To facilitate interpretability, we incorporate the feature importance analyses using the SHAP measure developed by Lundberg and Lee (2017) and Lundberg et al. (2020), and conduct valuation attribution analyses via the visualization of individual features' (potentially nonlinear) effects on the target variable via the ALE plots from Apley and Zhu (2020).

3 Data and Methodology

3.1 Fundamental information

We obtain 259 stock-level accounting characteristics—based solely on financial statement items—along with stock market information (e.g., market capitalization) for the primary security of listed firms worldwide from the global factor dataset described in Jensen, Kelly, and Pedersen (2023).[10] The dataset, which is recorded at a monthly frequency, is sourced from annual and quarterly COMPUSTAT (North America and Global) and the Center for Research in Security Prices (CRSP). For quarterly COMPUSTAT flow items, their values are aggregated over the last four quarters. Accounting characteristics are assumed to be publicly available four months after the end of an annual or quarterly accounting period and updated with the most recent data available.

Table A.3.1 provides the list of 259 accounting characteristics used in our study and their corresponding categories. We classify these accounting variables into 13 categories: Investment, Issuance, Profitability, Profit Growth, Growth*, Financial Soundness, Payout, Accruals, Efficiency, Liquidity, Capitalization, Solvency, R&D, and Miscellaneous, based on the cluster definitions of Geertsema and Lu (2023), Hou, Xue, and Zhang (2020), and Jensen, Kelly, and Pedersen (2023). Appendix A.1 provides more detail on how we classify the accounting characteristics.

The market capitalization and market-to-book ratio of firms are updated monthly using stock prices obtained from CRSP and COMPUSTAT datasets. Our empirical analysis focuses on firms that seem to be more economically relevant, that is, we exclude firms with total assets below the 20th percentile of NYSE stocks.[11]

[10] The dataset was obtained by running the code available at `https://github.com/bkelly-lab/ReplicationCrisis/tree/master/GlobalFactors`. We thank the authors for making the code publicly available.

[11] Similarly, Geertsema and Lu (2023) train their valuation models using U.S. stocks with total assets, sales and book equity each above the 10th percentile of their sample.

3.2 Corporate emission information

Annual firm-level GHG emission data, as reported by the firm and measured in tonnes of carbon dioxide equivalent (tCO_2e), is sourced from Trucost, Bloomberg, and Refinitiv (LSEG) in that order of priority. Specifically, the order of selection is the reported emission data of Trucost, Bloomberg data, Refinitiv data, and the derived emission data of Trucost. The emission data estimated by Trucost is not considered in our primary analysis. Firms with at least one non-missing value for either Scope 1 emissions or Scope 2 emissions are included in our sample. We merge the emission data with the financial dataset by using the GVKEY identifier along with ISIN. Restricted by the availability of emission data, which begins in 2003, our analysis sample period is from 2003 to 2022. Our main sample contains 6,766 unique stocks (476,556 total stock-month observations) across 77 exchanges. Fig. A.3.1 presents number of stocks by reporting status over the sample period.

To avoid potential look-ahead bias, as highlighted in Zhang (2024), we conservatively assume that GHG emission information is publicly available 12 months from the end of the accounting year. Our empirical exercises focus on Scope 1 and Scope 2 emissions. Scope 1 emissions are direct GHG emissions that occur from sources that are controlled or owned by a firm, while Scope 2 emissions are indirect GHG emissions associated with the purchase of electricity, steam, heat, or cooling (US EPA 2020). To align with the accounting characteristics (i.e., accounting ratios) and the target variable (i.e., market-to-book equity ratio or M2B), we use the emission intensity to train the GBRT model, which is defined as

$$GHGScope1_{intensity,i,t} = \frac{GHGScope1_{level,i,t}}{Sales_{i,y}}, \tag{3.1}$$

$$GHGScope2_{intensity,i,t} = \frac{GHGScope2_{level,i,t}}{Sales_{i,y}}, \tag{3.2}$$

where $GHGScope1_{intensity,i,t}$ represents the Scope 1 GHG emission intensity observed for firm i at the end of month t following the timing convention above, $GHGScope1_{level,i,t}$ denotes the observed Scope 1 GHG emission levels, and $Sales_{i,y}$ denotes revenues of firm i during the corresponding fiscal year y of the emission data. We define Scope 2 intensity similarly.

3.3 Summary of Scopes 1 and 2 GHG Emission Intensities

When conducting subsample analysis, we divide the sample into regions (Asia, Europe and the U.S.) or sectors ("Energy+" sector and "Others"). The "Energy+" sector includes all firms classified under the "Energy" industry in Fama-French-12 (FF-12) classification plus those in the electricity-related sectors of the FF-12 "Utility" industry classification. We also divide the sample period into three subperiods using the adoption of the Paris Agreement (December 2015) and the onset of the COVID-19 pandemic (December 2019) as cutoff points. Thus, our study focuses on three sample subperiods: 2003-2015, 2016-2019, and 2020-2022.

[Insert Table 1 near here]

Table 1 summarizes Scopes 1 and 2 GHG emission intensities, along with the number of firm-month observations, the number of firms (i.e., GHG reporters), and the valuation multiple (M2B) across four regions (Global, U.S., Europe, and Asia) during the three focal sample subperiods. We calculate the emission intensities and M2B at the region-period level, that is, we divided the total emissions by the total sales aggregated from all stock-month observations within a specific region (r) and period(p) combination:

$$GHGScope1_{intensity,r,p} = \frac{\sum_{i\in r,t\in p} GHGScope1_{level,i,t}}{\sum_{i\in r,y\in p} Sales_{i,y}}, \tag{3.3}$$

$$GHGScope2_{intensity,r,p} = \frac{\sum_{i\in r,t\in p} GHGScope2_{level,i,t}}{\sum_{i\in r,y\in p} Sales_{i,y}}. \tag{3.4}$$

The "Firms" column of Panel A suggests that the number of firms disclosing their GHG emissions has increased over time, both globally and within each of the three specific regions. The next column shows that U.S. firms typically have higher M2B than firms in Europe and Asia. For example, the average M2B in the U.S. (3.32) is more than twice as high as in Europe (1.60) and Asia (1.20) during the 2020-2022 period.

The "Scope 1" column indicates that global firms have GHG Scope 1 emission intensities of 0.36, 0.32 and 0.29 tCO_2e per US$1000 of sales. In all three regions, firms *directly* emit less GHGs per dollar earned (GHGs/USD) after 2015, which is likely to be driven by the adoption of the Paris Agreement in 2015. Additionally, Asian firms emit significantly more GHGs/USD than firms in the other two regions regardless of the period. As indicated by the last column, Scope 2 emission intensity of the U.S. and Europe overall exhibits a stable pattern of around 0.05 tCO_2e/$1000 over time. In contrast, Asian firms during the 2020-2022 period emit at notably high Scope 2 intensity of 0.21, which might be attributable to low revenues in USD arising from COVID or outliers caused by relatively poor reporting quality.

Panel B reports GHG emission intensities of energy and electricity producers (i.e., "Energy+" firms). Similar to the patterns observed in the full sample, the number of firms in the "Energy+" sector that report their GHG emissions has increased over time. Potentially due to the lack of growth opportunities compared to other types of firms such as those in high-tech industries, "Energy+" firms, which are often labelled as "value" firms, tend to have lower M2B. As expected, "Energy+" firms have higher Scope 1 emission intensity. Among the three regions, European firms directly emit the least GHGs/USD across all three subperiod windows. In contrast, Asian firms directly emit substantially more GHGs/USD (approximately 2 and 1.5 times more than firms in Europe and the U.S, respectively). One explanation for this is that the energy sector in Asia might be highly subsidized by the government, requiring less profit compared to its European and U.S. peers. Additionally, the last column in Panel B indicates that "Energy+"firms' Scope 2 emission intensity ranges from 0.02 to 0.07.

3.4 Valuation Model

Our workhorse stock valuation model is based on machine leaning techniques. We employ the Light Gradient Boosting Machine (LightGBM),[12] a high-performance implementation of Gradient Boosting Regression Trees (GBRT) algorithm described in Appendix A.2.1, to predict a firm's log-transformed M2B using the two GHG emission intensity variables described above along with 259 accounting characteristics (denoted as "X_{acct}"). The model can be expressed in a functional form as:

$$Log(M2B_{i,t}) = GBRT(X_{ghg,i,t}, X_{acct,i,t}, X_{other,i,t}), \tag{3.5}$$

where $X_{ghg,i,t}$ includes the two GHG emissions intenisities above, and $X_{other,i,t}$ includes 3 categorical features (year, exchange, and Fama-French 12 industry). The model-implied M2B is thus the exponential of the log of the M2B implied by the GBRT model:

$$\widehat{M2B_{i,t}} = Exp(Log(\widehat{M2B}_{i,t}). \tag{3.6}$$

The model-implied market capitalization (MCAP) is obtained by multiplying each firm's most recent book equity (BE) by their estimated $\widehat{M2B_{i,t}}$, i.e.,

$$\widehat{MCAP}_{i,t} = Exp(Log(\widehat{M2B}_{i,t}) \times BE_{i,t} \tag{3.7}$$

[12]LightGBM is a technique developed by Ke et al. (2017) from Microsoft, which is notable for its efficiency and scalability in training samples with a large number of observations and features compared to conventional implementation of GBRT (e.g., XGBoost). It creatively relies on histogram-based algorithms and a leaf-wise tree growth strategy to efficiently reveal complex predictive patterns in data. Specifically, Ke et al. (2017) propose Gradient-based One-Side Sampling (GOSS) and Exclusive Feature Bundling (EFB) to implement GBRT and document that LightGBM outperforms the conventional GBRT implementation in terms of training speed (up to over 20 times faster) without sacrificing accuracy.

The GBRT model is robust to missing values and outliers in features. Unlike linear regressions and some other machine learning models such as feed-forward neural networks, GBRT does not requires imputation of missing values in explanatory variables. As GBRT is essentially a decision-tree-based approach, it is robust to outliers in features. Nevertheless, we restrict our sample to firms with a monthly M2B ratio between 0.01 and 100 to minimize the potential effect of outliers in the target variable.[13]

For each subsample period, we train a separate GBRT model with LightGBM. Our training procedure is described in detail in Appendix A.2. In brief, firm-month observations in each subsample period (e.g., 2003-2015) are randomly divided into training and validation subsets using a 4:1 ratio. We use the validation subset to select the best-performing set of hyperparameters from 1000 candidate sets. The hyperparameter search spaces and selected set of hyperparameters for each subsample global model are provided in Table A.3.2 . Then, the model is retrained on the full subsample using the selected set of hyperparameters.[14]

3.5 Interpretability of the Valuation Model

The difficulty in interpreting machine learning models, i.e., understanding why a model produces a certain prediction, deters many researchers from switching from easy-to-interpret linear regression models to more accurate but complex machine learning models in economics and finance research. To minimize the tension between accuracy and interpretability, the machine learning literature has introduced several techniques to assist users in understanding model predictions, especially attribution of features to the target variable. We employ these techniques in this study. In particular, to analyse relative importance of GHG emission features and their value relevance, we employ SHapley Additive exPlanations (SHAP) proposed by Lundberg and Lee (2017) and Accumulative

[13]This removes around 0.14% of the stock-month observations after applying the size filter.

[14]In addition to the validation step, we use several techniques to reduce overfitting in nonlinear models as per LightGBM's official user guide (https://lightgbm.readthedocs.io/en/latest/Parameters-Tuning.html).

Local Effects (ALE) proposed by Apley and Zhu (2020).

3.5.1 SHapley Additive exPlanations (SHAP)

SHAP is a model-agnostic approach used to explain machine learning predictions at the level of individual data points, such as single stock-month observations in our study. SHAP measures individual features' contribution to a model prediction by assigning them values of Shapley (1951), a solution concept in cooperative game theory. We compute SHAP values using the tree-SHAP algorithm of Lundberg et al. (2020). For each observation, a feature's SHAP value measures how it drives the model prediction towards or away from the expected value measured by the mean of all model predictions. Furthermore, as suggested by its name, SHAP values are additive so that values of individual features can be aggregated to represent the importance of categories. In line with Geertsema and Lu (2023), we transform raw SHAP values into percentage terms by scaling absolute SHAP values across all features for each observation so that their sum equals 100%.

3.5.2 Accumulative Local Effects (ALE)

We employ the ALE plots, introduced by Apley and Zhu (2020), to examine the influence of GHG emission features on corporate valuation. The ALE function for GHG emissions can be expressed as

$$ALE(x_{ghg}) = \int_{min(x_{ghg})}^{x_{ghg}} E\big[\frac{\partial g(X_{ghg}, X_{\setminus ghg})}{\partial X_{ghg}} | X_{ghg} = z_{ghg}\big] dz_{ghg} - c_{ghg} \tag{3.8}$$

where $g(x)$ denotes a valuation model, x_{ghg} denotes the GHG emission feature of interest, $x_{\setminus ghg}$ denotes the remaining features,[15] and c_{ghg} is chosen so that the average effect over the sample is normalized to zero. The partial-derivative term $\frac{\partial f(X_{ghg}, X_{\setminus ghg})}{\partial X_{ghg}}$ measures the local (marginal) effect of x_{ghg} on $g(x)$, which is accumulated over the range of the GHG

[15] $x_{\setminus ghg}$ includes all remaining features. For example, when examining the ALE of Scope 1 emission intensity on M2B, Scope 2 emission intensity is included into $x_{\setminus ghg}$.

emission variable from the minimum value of x_{ghg} of all observations to the observed value of x_{ghg} for the specific observation.

In practice, Eq. (3.8) can be estimated by finite differences, that is,

$$\widehat{ALE}(x_{ghg}) = \sum_{b=1}^{b(x_{ghg})} \frac{1}{n(b)} \sum_{i:x^i_{ghg} \in L(b)} \left[g(z_b, x^i_{\setminus ghg}) - g(z_{b-1}, x^i_{\setminus ghg}) \right] - c_{ghg}, \tag{3.9}$$

where $z_0 ... z_b ... z_B$ is a sufficiently fine grid of x_{ghg}, $b(x_{ghg})$ denotes the index of bin in which x_{ghg} falls (i.e., $x_{ghg} \in (z_{b(x_{ghg})-1}, z_{b(x_{ghg})}]$), $L(b)$ denotes the length of the interval $(z_{b-1}, z_b]$, and $n(b)$ denotes the number of observation falling into the interval $(z_{b-1}, z_b]$. For each observation i falling into the interval $(z_{b-1}, z_b]$, we replace its value of the GHG feature (x^i_{ghg}) with the values of both right and left interval end-points—z_b and z_{b-1}—keeping the value of remaining variables unchanged, and evaluate the difference of predictions at these points. We take the average across all observations in that interval to estimate the local effect. The ALE value is therefore the sum of the estimated local effects starting from the first interval (with the lowest values of x_{ghg} in the sample) to the interval where observation i is located.

An important property of ALE is that it is model-agnostic, making it applicable to any type of predictive models, including linear and non-linear models, regardless of whether they are differentiable. This means that $g(x)$ can represent not only the LightGBM model for predicting the log of M2B, which is not differentiable like other decision-tree methods, but also the entire procedure that we used to obtain the predicted M2B and MCAP. Similarly, x_{ghg} can represent either GHG emission levels or GHG emission intensity. Fig. 1 illustrates how to apply ALE to a LightGBM-based valuation model designed to use GHG emission levels along with other firm characteristics to predict MCAP.

[Insert Figure 1 near here]

4 Results

4.1 Model Performance

Before analyzing the attribution of firms' market values to their carbon footprints, we examine the explanatory power of our machine-learning-powered valuation model by comparing it to the linear model. In order to create a level playing field, we employ a valuation model based on Ordinary Least Squares regression (OLS) that includes the same fundamental variables as the GBRT model, i.e., the 259 accounting variables plus 2 GHG variables, and the hosts of fixed effects. We compare the performance of this comprehensive linear model to that of our GBRT-based valuation model—introduced in Section 3.4—in explaining the observed M2B and MCAP values. As performance metrics, we employ both conventional R-Squared (R^2) measure and the Median Absolute Percentage Errors (MDAPE) measure frequently employed to evaluate machine learning models.

The R^2 values for M2B and MCAP are respectively defined as:

$$\mathrm{R}^2(\mathrm{M2B}) = 1 - \frac{\sum_{i,t}(\widehat{M2B}_{i,t} - M2B_{i,t})^2}{\sum_{i,t}(Mean(M2B_{i,t}) - M2B_{i,t})^2} \text{ and} \tag{4.1}$$

$$\mathrm{R}^2(\mathrm{MCAP}) = 1 - \frac{\sum_{i,t}(\widehat{MCAP}_{i,t} - MCAP_{i,t})^2}{\sum_{i,t}(Mean(MCAP_{i,t}) - MCAP_{i,t})^2}, \tag{4.2}$$

where $\widehat{M2B}_{i,t}$ and $\widehat{MCAP}_{i,t}$ denote model-implied M2B and MCAP at the stock-month level. The MDAPE value for MCAP is defined as:[16]

$$\mathrm{MDAPE} = Median\left(\left|\frac{\widehat{MCAP}_{i,t} - MCAP_{i,t}}{MCAP_{i,t}}\right|\right). \tag{4.3}$$

The pooled OLS valuation model that we use to predict M2B and MCAP is specified as:

$$Log(M2B_{i,t}) = a + b \times X_{ghg,i,t} + c \times X_{controls,i,t} + \epsilon_{i,t}, \tag{4.4}$$

[16] According to Eq. (3.6) and Eq. (3.7), the MDAPE value for M2B is the same as that for MCAP.

where $X_{ghg,i,t}$ denotes a vector of Scope 1 and Scope 2 emission intensities, and $X_{controls,i,t}$ denotes a vector of control variables including the firm characteristics identical to the GBRT model, year-month fixed effects, country-fixed effects and industry-fixed effects.[17] After obtaining the predicted log-transformed M2B ($\widehat{log(M2B)}$), we use Eq. (3.6) and Eq. (3.7) to calculate $\widehat{M2B_{i,t}}$ and $\widehat{MCAP}_{i,t}$ from the linear model.

Unlike decision-tree approaches such as GBRT, the OLS model is not robust to outliers and missing values in explanatory variables. Thus, we winsorize all explanatory variables at the 1% and 99% levels, and replace missing values of accounting characteristics with their month-exchange-industry medians. In line with the GBRT valuation model, we estimate three separate models using data from the periods: 2003-2015, 2016-2019, and 2020-2022, respectively. When assessing model performance by region or sector, we evaluate the model-implied M2B or MCAP within a specific combination of period-region or period-sector.

[Insert Table 2 near here]

Table 2 presents comparative results, with Panel A showing model performance by region and Panel B reporting results by sector. The GBRT model overall exhibits robust performance, producing estimates of M2B and MCAP that are reasonably close to the observed market values. For the global sample, it can deliver R^2 values of more than 90% for M2B and more than 95% for MCAP. The model produces accurate valuation estimates, with median valuation errors that are less than 15%. The GBRT model is stable across regions, achieving similar accuracy for firms in the U.S., Europe and Asia. In comparing across sectors, although R^2 for M2B suggests the GBRT model is more accurate for the non-"Energy+" sector, the other two metrics indicate the stability of the model for the two sectors.

[17] Bolton and Kacperczyk (2023) use the pooled linear regression with the same fixed effects when estimating the effects of GHG emissions on the log of the book-to-maker ratio for global stocks (see their Table VII).

The GBRT model consistently outperforms its equally comprehensive OLS counterpart in explaining M2B and MCAP across regions and sectors in all three periods. In addition to the GBRT model's higher R^2s in the first four columns, the last two columns of both panels report that the accuracy of the GBRT model for the two more recent periods is about twice that of the OLS model in terms of MDAPE—an outlier-robust measure—regardless of the region or sector.

4.2 Feature Importance

After documenting the relative accuracy of the GBRT model, we now turn the attention to employing the model to evaluate the importance of corporate GHG emissions in equity valuation. We measure the importance of each feature by calculating its percentage contribution (in terms of absolute SHAP values, as detailed in Section 3.5.1) to the estimated (log-transformed) M2B produced by the GBRT valuation model. We further scale each absolute SHAP value by dividing it by the largest SHAP value, so that the largest value is equal to 1.

[Insert Figure 2 near here]

Panel A of Fig. 2 presents the 25 most important firm fundamental variables in determining M2B during each of the three subsample periods. The SHAP plots indicate that the profitability measures – e.g., variants of the return on equity or assets measures (i.e., earnings-related items scaled by book equity or assets) – play substantial roles in the equity market valuations of publicly listed firms in all three periods. Additionally, dividend payout (scaled by their book assets, i.e., "div_at") persistently exhibits strong value relevance over time, ranking as the fifth, third and third variable in the three sample periods, respectively. Remarkably, Scope 1 GHG emission intensity consistently ranks among the top 10 variables.

Panel B further reports the time-series ranks of Scope 1 emission intensity derived from two metrics specific to LightGBM: "split" and "gain". The "split" metric is defined as the frequency with which a feature is used to make a split in the decision tree, while the "gain" metric reflects how much the tree benefits from using the feature to create a new split point in terms of reduction in training loss. Both metrics tend to reach a consensus with SHAP that Scope 1 carbon intensity is more relevant to equity valuation than the Scope 2 counterpart.

Additionally, Panel B depicts the ranks of Scope 2 GHG emission intensity in terms of the three importance metrics. All three metrics indicate that the Scope 2 emission intensity is less value relevant than Scope 1 emission intensity. In sum, investors have distinct views on the effects of direct and indirect corporate GHG emissions on equity valuation.

[Insert Table 3 near here]

Table 3 presents the feature importance implied by the percentage contribution of absolute SHAP values to predictions at the feature category level for the full universe of firms, as well as for "Energy+" firms and Non-"Energy+" firms, separately.

Overall, Scope 1 GHG emission remains persistently important with prediction contributions of more than 2% under all three valuation models. It is more important than Scope 2 emission during all three periods for the entire universe as well as for the two types of firms considered. Moreover, Scope 1 emission is more value-relevant to firms in the "Energy+" sector, and its importance is increasing over time. Scope 2 GHG emission is also more value-relevant to "Energy+" firms. Different from the Scope 1 counterpart, Scope 2 emission first becomes more important after the Paris agreement, but the importance attenuates during the COVID period though still higher than during the pre-Paris-agreement period.

Regarding accounting characteristics, features in the Profitability category are consistently the most important irrespective of period and firm type, but their aggregate importance significantly drops during the recent period. Notably, the value relevance of R&D features increases over time, which is consistent with the finding of Barth, Li, and McClure (2023) that accounting items relating to intangible assets exhibit increased value relevance.

4.3 Effects of GHG Emission Intensity on M2B

In this section, we examine how GHG emission intensity affects firms' valuation multiple, M2B. We employ the ALE plots introduced in Section 3.5.2 to visualize the local effects of corporate GHG emission intensity on the M2B implied by our GBRT-based valuation model, after controlling for the effects of firms' accounting characteristics.[18]

[Insert Figure 3 near here]

Fig. 3 provides the ALE plots for the effects of Scope 1 GHG (the top three sub-figures) and Scope 2 GHG (the bottom three) emission intensity on M2B. Each sub-figure includes firms in the U.S., Europe and Asia during one of the three subsample periods. For better visualization of the effects, we apply a base-10 log scale to the x-axis.

With the exception of the U.S. ALE curve for the 2016-2019 period—which has near-flat fluctuating slopes for low emitters but negative slopes for relatively high emitters—all ALE curves of Scope 1 tend to exhibit downward sloping trends, indicating a negative association of equity valuation with direct GHG emission intensity. Additionally, M2B responds to a (10-fold) increase in Scope 1 emission intensity more negatively in all three regions after the Paris agreement in 2015. That is, the equity market penalizes additional

[18] The ALE method can be employed for various estimation models. In particular, the ALE plot for a specific independent variable of interest in a linear regression model is a straight line whose slope is the estimated regression coefficient on that variable. The ALE method is arguably more suitable than partial dependence plots (PDP), another popular tool in the Explainable AI literature, for attributing model predictions to a feature that is correlated with other features (Apley and Zhu 2020).

direct emission more heavily after the Paris agreement. Among the three regions, the ALE curves for Asia and Europe almost overlap during each period, suggesting investors' similar views on the value impacts of direct emissions for firms in these two regions. In contrast, the market pricing of carbon footprints in the U.S. seems to behave differently, particularly for high emitters.

In contrast, the ALE curves of Scope 2 are much flatter than those of Scope 1. That is, given the same level of emission intensity, the financial consequences of a marginal increase in indirect emissions are generally less severe than those of direct emissions. Also, the indirect emission information is incorporated into market valuation in a similar manner across all three regions. These patterns indicate that while the Paris Agreement presumably significantly affects the value relevance of corporate emission features, investors continue to put more emphasis on emissions directly controlled by entities, which is consistent with the results from feature importance analyses reported in Table 3.

[Insert Figure 4 near here]

Given that Scope 1 emission intensity is identified as a more important feature in determining firm values compared to Scope 2 emission intensity, Fig. 4 further provides Scope 1's ALE curves for all firms (the top three sub-figures) and "Energy+" firms (the bottom three), which are plotted on the raw scale for both y-axis and x-axis. Observations with emission intensity values above the 95th percentile are omitted. Note that the filter is applied separately to the "Energy+" firms and other firms.

The top three plots show that the slopes of the ALE curves vary across different levels of emission intensity for all periods and regions considered, highlighting non-linear and complex relationships between the intensity of corporate direct GHG emissions and M2B. Overall, the direction of effects is negative; however, the highest emitters always have much flatter ALE curves than other firms, irrespective of the region and period.

The bottom three plots confirm that the weak relevance of direct emissions in equity valuation is driven by markets' tolerance to "harder-to-abate" firms, i.e., the Scope 1 emission-intensive "Energy+" firms. Specifically, after the Paris agreement in 2015, direct emissions of above-median emitters in the "Energy+" sector are weakly correlated with the valuation ratio across three regions. As for the pre-agreement period, except for the top-emitting bin of Asia, the associations of M2B with emission intensity for all relatively large emitters (i.e., bins 7 through 10) in the U.S. and in Asia are slight. Regarding European firms in the "Energy+" sector, there are negative pricing effects arising from increases in Scope 1 emission intensity across all bins during the 2003-2015 period

4.4 "Dollar" Effects of Direct GHG Emissions

We have so far focused on valuation ratio (M2B) and emission intensity (GHG emission/revenues) and found that direct emission intensity is more relevant to equity valuation than intensity of indirect emissions. In this section, we investigate how the *level*[19] of *direct* GHG emissions is associated with *market values* of the emitting firms. We report the ALE plots for this configuration in Fig. 5 for all firms (the top three sub-figures) and "Energy+" firms (the bottom three). The level of emissions has a very wide range and a highly right-skewed distribution. As such, we use the base-10 log scale for the x-axis of the plots.

[Insert Figure 5 near here]

The top three subplots indicate that direct emissions tend to have a negative effect on market cap in all three regions regardless of the period. In the U.S., the negative association is becoming more pronounced over time. In particular, after the Paris Agreement, the negative "Dollar" effects for large U.S. emitters from the top three bins tend to be stronger than their counterparts in Europe and Asia. In other words, the emission costs

[19]The level of carbon emissions is also a stock characteristic advocated by Bolton and Kacperczyk (2021a, 2023).

implied by the reduction in stock prices are the highest in the U.S. particularly for large emitters.

We perform a similar analysis for firms in the "Energy+" sector. As shown in the bottom three sub-figures, the pricing patterns for the "Energy+" firms are generally consistent with but weaker than those for all firms with an exception that direct emissions are positively associated with market values for the top two bins of European "Energy+" firms in the recent period. In addition, during the 2016-2019 period, the stock market tends to price "Energy+" firms' direct emissions in a *globally* consistent manner.

[Insert Figure 6 near here]

4.4.1 Model-implied Scope 1 GHG Emission Prices

In this section, we further quantify the effects of direct carbon emissions on corporate valuation by estimating the market-imputed Scope 1 GHG emission prices implied by our valuation model. Our estimation is performed using a quantitative model that is built on the ALE method. Specifically, we first estimate the emission price for each interval $(z_{b-1}, z_b]$ of GHG emission levels using the average marginal change in firm values with respect to GHG emission levels for all observations within the interval, that is,

$$\widehat{P}_b(x_{ghg}) = \frac{1}{n(b)} \sum_{i:x^i_{ghg} \in L(b)} \Big[\frac{g(z_b, x^i_{\setminus ghg}) - g(z_{b-1}, x^i_{\setminus ghg})}{z_b - z_{b-1}}\Big]. \tag{4.5}$$

This calculation also represents the slope of the tangent to the ALE curve at bin b. When estimating the emission price for a particular sample group, we initially divide it into 10 bins based on emission levels. We calculate the emission price for each bin and the take the average of these bin prices using the average emission levels in each bin as weights. As such, we place more importance to the estimated prices for firms with higher emission contributions. The price estimates are reported in Table 4.

[Insert Table 4 near here]

We perform this estimation of market-imputed emission prices at the regional level for each subsample period. Panel A of Table 4 indicates that the market-imputed prices of Scope 1 emissions (in US$ per tonne) are positive irrespective of the region-period combination and that these prices are rising over time within each region. Prior to the Paris agreement, the costs of direct emission are similar across the three regions, while after that, the U.S. has the highest emission price followed by Europe and Asia. The imputed emission prices from the valuation models for the most recent period are 148.37, 53.94 and 34.16 US$ per tonne of CO_2 equivalent (US$/t$CO_2$e) in the U.S., Europe and Asia, respectively.

It is important to note that, as shown in Panels B and C of Table 4, firms in the "Energy+" sector tend to have lower imputed carbon prices for their direct emissions relative to other firms. The imputed Scope 1 emission prices for firms in the "Energy+" sector for the latest period are 25.83 (U.S.), 14.38 (Europe), and 7.54 (Asia) US$ per tCO_2e, which are much lower than the corresponding prices for non-"Energy+" firms (203.06, 98.34, and 42.48 US$ per tCO_2e in the three regions). This suggests that the market penalizes the excessive direct emissions of "easier-to-abate" firms, such as airline companies, much more heavily than those of energy generators whose entire business models depend entirely on economic activities that produce substantial GHG emissions.

4.5 Do Corporate Carbon Footprints Contain Profitability Information?

In the last part of the study, we explore a potential channel driving the negative association between direct emissions and corporate valuation. In particular, we examine whether direct emission intensity is negatively associated with future firm profitability, after controlling for other fundamental aspects of the firm. In particular, we examine how firm's

GHG emission intensity, observed in month t, affects its 5-year-ahead ($t + 60$) profitability. We employ the GBRT model to predict the return on book equity (ROE), which is calculated as net income divided by book equity (ni_be). Note that "ni_be" is the most important profitability variable in our valuation model, based on its SHAP values. In addition to Scopes 1 and 2 emission intensities, we include accounting characteristics that are not in the Profitability and Profit Growth categories as explanatory variables, along with the three categorical features included in our valuation model. We also include the lagged target variable (ROE) to take into account potential persistence in profitability. The model specification can be expressed as follows:

$$ROE_{i,t+60} = GBRT(X_{ghg,i,t}, X_{\setminus profit,i,t}, ROE_{i,t}). \tag{4.6}$$

4.5.1 GHG Emissions as Predictors of Future Profitability

[Insert Figure 7 near here]

The feature importance figure (Fig. 7) based on SHAP values shows that Scope 1 emission intensity is more important than Scope 2 emission intensity in predicting firm profitability. Specifically, Scope 1 carbon intensity is the fourth important stock characteristic, while Scope 2 carbon intensity is the nineteenth important feature. Also, the current ROE is unsurprisingly the most important predictor of future ROE followed by free cash flow scaled by operating cash flow and the dividend payout ratio (i.e., "div_at").

Next, we perform a directional analysis of direct emissions' relevance in predicting future profitability by plotting the ALE of Scope 1 emission intensity in Fig. 8 for all firms (the top sub-figure) and the firms in the "Energy+" industry (the bottom sub-figure). The x-axis of the figure is on the base-10 log scale. The top plot suggests that in general, Scope 1 emission intensity tends to be positively correlated with future profitability when Scope 1 emission is less intensive (i.e., firms belonging in the first few bins) and negatively associated with future profitability when the emission intensity is high except for the top bin of Asia. This can explain why emission-weighted (market-imputed) Scope 1 emission prices are positive. Moreover, the negative association between Scope 1 emission intensity and 5-year-ahead ROE tends to be stronger in the U.S.. These results indicate that at least some market participants had expected these negative cash flow shocks predictable by direct GHG emissions of carbon-intensive firms and incorporated these expectations into ex-ante market valuations of corporate equity.

The bottom graph demonstrates that for above-median emitters in the "Energy+" industries of the three regions, their Scope 1 emission intensity overall exhibits weaker negative associations with future profitability. This can account for why the market-imputed prices of the "Energy+" firms are lower than others.

[Insert Figure 8 near here]

5 Robustness Checks

In this section, we conduct robustness checks for our study. In Section 5.1, we restrict the sample to the GHG reporters listed in the U.S. and build GBRT-based valuation models for the three periods. Next, in Section 5.2, we extend our analysis to a more comprehensive stock universe by adding stocks with GHG emissions estimated by Trucost (including half a million stock-month more observations) to our primary sample. We first directly use the estimated GHG values along with the reported values to train three valuation models. Then, we treat the estimated GHG values as missing and train three valuation models. Since the GBRT model is robust to missing values in input variables, firms with missing emission data can be included.

5.1 U.S.-only Model

A potential concern about our global valuation models is that even after controlling for country fixed effects, they may not sufficiently account for the differences in regulations, accounting standards, and investor sentiment among countries. In this section, we restrict our sample to U.S. stocks. The U.S. stock market is the world's most economically important market by total market value. We first examine the accuracy of the U.S.-only valuation model. Our U.S.-focused machine learning model outperforms the global counterpart in valuing U.S. stocks. For "Energy+ firms, it can achieve R^2 (M2B) values of 94.6%, 92.8%, and 93.8% for three periods, with median (absolute) valuation error of 5.9%, 5.0%, and 6.9%. For other firms, it can achieve R^2 (M2B) values of 95.8%, 98.0%, and 95.8% for three periods, with median (absolute) valuation error of 6.1%, 4.9%, and 6.4%.

Fig. A.3.2 reports importance of GHG emission intensity in determining M2B relative to accounting characteristics of firms. Consistent with the global model, Scope 1 emission intensity is more value relevant than Scope 2 emission intensity.

[Insert Figure 9 near here]

Fig. 9 illustrates the directions of emission intensity's value impacts using ALE curves on a log-scaled x-axis. The Scope 1 pattern of the 2016-2019 period shown in the top three sub-figures is generally consistent with that shown in the global models (Fig. 3). For the 2003-2015 period, the top three bins behave differently in the two valuation models. For the recent period, the two valuation models imply similar pricing patterns for large emitters in bins seven through ten and different patterns for the rest.

Compared to the patterns suggested by the global model (Fig. 3), the U.S.-only model implies a notable sensitivity of M2B to Scope 2 emission intensity for the recent period. The direction of the effects transitions from positive for less carbon-intensive firms to negative for large emitters. For the remaining two periods, the two models tend to suggest similar responses of M2B to indirect emission intensity.

5.2 Models with Trucost Estimated GHG Sample

5.2.1 Treating Trucost Estimates as Reported Values

We expand our baseline GHG reporter sample by including stocks with GHG emission values estimated by Trucost and train three new valuation models. In other words, we treat Trucost estimates as if they were values reported by these firms.[20] We assume that investors can either observe the estimated GHG values 12 months after the end of the accounting year or estimate the GHG values in a similar manner to Trucost.

[Insert Figure 10 near here]

Fig. 10 demonstrates the directional effects of emission intensity on M2B. Consistent with the pattern shown in the reporter-only models (Fig. 3), direct emission intensity is typically negatively correlated with M2B. M2B shows comparable sensitivities to direct emission intensity during the first two periods but still exhibits highest sensitivity to Scope 1 emission intensity during the recent period. In addition, indirect emission intensity overall shows a pretty weak association with M2B across regions and over time, consistent with the corresponding ALE curves of the reporter-only model. The exception is medium emitters in the U.S. whose M2B responds to Scope 2 emission intensity in a non-negligible negative direction.

5.2.2 Treating Trucost Estimates as Missing Values

Using the same sample as in Section 5.2.1, i.e., firms with reported or estimated GHG values from Trucost, we train three additional valuation models but treat the Trucost-

[20] Fig. A.3.3 illustrates the time-series importance of emission intensity in capturing M2B based on three metrics, absolute SHAP value, 'split,' and 'gain.' It shows that Scope 1 emission intensity is always one of the most important features and more important than Scope 2 emission intensity, which is consistent with the result based on the reporter-only model.

estimated GHG values as missing values. As the GBRT model can decide to use the "missing value" status as a criterion for splitting decision trees, we essentially enlarge our baseline GHG reporter sample using Trucost's coverage decision of corporate GHG estimation as the parameter of sample selection.

[Insert Figure 11 near here]

The results from this alternative treatment are reported in Fig. 11, which depicts the effects of emission intensity on M2B implied by the three new models. The figures show similar pricing patterns for Scope 1 as those in Fig. 3 which are based on the models ignoring non-reporting firms. For Scope 2, except for the top-emitting bin of the U.S, the pricing patterns do not deviate from those implied by the reporter-only model much.

It is important to note that these analyses are relegated to robustness checks as they are subject to the prevailing concerns that Trucost's sample selection for estimation coverage as well as its proprietary estimation methodology may be directly correlated with stock valuations.

6 Conclusion

This study examines how *equity markets* incorporate corporate carbon emissions into stock prices around the world. We propose a comprehensive set of machine-learning-powered fundamental analysis models capable of not only uncovering complex and nonlinear relations between corporate fundamentals and firm valuation, but also generating imputed prices of corporate carbon emissions.

The first model is based on gradient boosted regression trees (GBRT) which can accommodate an extensive set of fundamental variables. Our GBRT valuation model has a very robust explanatory power for equity valuation of publicly listed firms around the world. The R^2 of the GBRT model is over 95% in explaining variations in market capitalizations (MCAP) and over 90% in explaining variations in market-to-book ratios (M2B), much higher than corresponding linear (OLS) models incorporating the same comprehensive set of corporate features and fixed effects. The GBRT model estimates' deviations from the observed market values are less than 15% of the observed values, which are typically half of the typical deviations of OLS models' estimates.

With the GBRT valuation model providing a more accurate reflection of equity market participants' evaluations of corporate features, we perform model-based attribution analyses to quantify the relative importance of corporate carbon emissions. Scope 1 GHG emission intensity (i.e., emissions-to-sales) ranks within the top ten important features out of 261 corporate features included in the valuation model in each of the three subperiods of our sample. In contrast, market participants view corporate indirect (Scope 2) emissions differently, with the rank of Scope 2 GHG emission intensity falling within the 25-50 range.

To uncover the directional effects of carbon emissions on the M2B and MCAP estimates generated by our valuation model, we employ the Accumulative Local Effects (ALE) method proposed by Apley and Zhu (2020), which produces plots to illustrate the complex and non-linear patterns unveiled by the GBRT model. Indeed, the ALE plots generated from our valuation models illustrating the impact of carbon emission intensity on M2B are generally *convex* in all regions with the *negative* slopes more pronounced after the Paris Agreement, particularly for Scope 1 GHG intensity. These patterns indicate that excessive emissions from the operations directly controlled by the firm are more likely to be regarded as a negative signal by equity market participants.

Our estimation framework produces estimates of equity-market-imputed prices of carbon emissions. We find that in all three regions, the imputed equity-based prices for Scope 1 GHG emissions are positive and rising over time. The imputed emission prices from the valuation models for the most recent (2020-2022) period are about US$ 30-150 per tonne of CO_2 equivalent (tCO_2e) of corporate Scope 1 GHG emissions. The imputed

Scope 1 emission prices for firms in the fossil-dependent “Energy+” sector are much lower than the corresponding prices for firms in other industries, indicating that equity market participants penalize excessive direct emissions of “easier-to-abate” firms, such as airline companies, more heavily than those of energy firms whose entire business models depend on activities that produce substantial GHG emissions.

We identify an important channel driving the negative association between direct emissions and corporate valuation: emission intensity is negatively associated with *future* firm profitability. Firms with relatively high direct carbon emission intensity exhibit lower future profitability. This pattern is consistent with equity market participants correctly anticipating the negative cash flow shocks that are uniquely predictable by firms’ direct carbon emissions and these expectations impounded into stock prices in the global equity markets.

Table 1 GHG Emission Intensity across Regions and over Time

This table presents Scopes 1 and 2 GHG Emission Intensities, along with the number of firm-month observations, the number of firms, and the market-to-book ratio (M2B) for three sample periods across four regions (Global, U.S., Europe, and Asia). The emission intensities and M2B are calculated at the region-period level. Scope 1 (Scope 2) emission intensity (tCO_2e per thousand U.S. Dollar) is defined as the total emissions divided by the total sales aggregated from all firm-month observations within a specific region-period. Similarly, M2B is calculated as the total market capitalization divided by the total book equity for all firm-month observations within the same given region-period. The data is restricted to firms that report GHG emissions. Panel A includes all reporting firms, while Panel B covers the reporting firms in the "Energy+" sector. The "Energy+" sector includes all firms classified under the "Energy" industry in Fama-French-12 (FF-12) classification plus those in the electricity-related sectors of the FF-12 "Utility" industry classification.

Panel A: All Industries

Region	Period	Observations	Firms	M2B	Scope 1 $\left(\frac{tCO_2e}{\$1,000}\right)$	Scope 2 $\left(\frac{tCO_2e}{\$1,000}\right)$
Global	2003-2015	163,872	2,429	1.72	0.36	0.04
	2016-2019	138,915	4,243	1.78	0.32	0.05
	2020-2022	173,769	6,517	1.98	0.29	0.10
U.S.	2003-2015	37,669	516	2.12	0.31	0.04
	2016-2019	26,973	796	2.58	0.22	0.04
	2020-2022	34,243	1,303	3.32	0.17	0.03
Europe	2003-2015	56,642	770	1.51	0.30	0.03
	2016-2019	41,218	1,149	1.54	0.27	0.04
	2020-2022	43,774	1,529	1.60	0.22	0.03
Asia	2003-2015	52,914	841	1.43	0.51	0.04
	2016-2019	53,660	1,784	1.24	0.46	0.05
	2020-2022	75,988	2,965	1.20	0.42	0.21

Panel B: "Energy+"

Region	Period	Observations	Firms	M2B	Scope 1 $\left(\frac{tCO_2e}{\$1,000}\right)$	Scope 2 $\left(\frac{tCO_2e}{\$1,000}\right)$
Global	2003-2015	20,460	266	1.35	1.02	0.03
	2016-2019	13,351	366	1.09	1.10	0.06
	2020-2022	13,017	465	1.51	0.93	0.07
U.S.	2003-2015	6,439	76	1.58	1.12	0.04
	2016-2019	3,631	98	1.41	1.11	0.07
	2020-2022	3,733	135	1.51	0.95	0.07
Europe	2003-2015	6,215	82	1.16	0.64	0.02
	2016-2019	3,729	97	0.91	0.70	0.06
	2020-2022	3,398	121	0.97	0.55	0.05
Asia	2003-2015	5,715	78	1.21	1.94	0.03
	2016-2019	4,146	116	0.84	1.65	0.02
	2020-2022	4,246	145	0.76	1.42	0.06

Table 2 Model Performance: GBRT vs OLS

This table reports the performance of Gradient Boosting regression Trees (GBRT) and Ordinary Least Squares regression (OLS) in predicting Market Capitalization (MCAP) and the Market-to-Book ratio (M2B). For each sample period (e.g., 2003-2015), we obtain MCAP and M2B predictions for all global firms in the sample from GBRT and OLS, respectively. We employ Median Absolute Percentage Error (MDAPE) and R-squared (R^2) as performance metrics. For each period, we evaluate the performance of the global models not only for the full sample (Global) but also for three regions (U.S., Europe, and Asia) in Panel A or 2 sectors ("Energy+" and Others) in Panel B where the "Energy+" sector includes all firms classified under the "Energy" industry in Fama-French-12 (FF-12) classification plus those in the electricity-related sectors of the FF-12 "Utility" industry classification. The LGBM-based valuation model is introduced in Section 3.4, while the OLS model is the pooled linear regression with year-month, exchange and industry fixed effects. Additionally, when estimating OLS models, we winsorize all explanatory variables at the 1% and 99% levels, and replace missing values of accounting characteristics with their month-exchange-industry medians.

Panel A: Model Performance by Region

Region	M2B		MCAP			
	R^2		R^2		MDAPE	
	GBRT	OLS	GBRT	OLS	GBRT	OLS
2003-2015						
Global	0.93	0.69	0.97	0.84	0.12	0.26
U.S.	0.92	0.66	0.97	0.84	0.11	0.24
Europe	0.93	0.69	0.96	0.79	0.12	0.26
Asia	0.93	0.74	0.97	0.84	0.12	0.28
2016-2019						
Global	0.94	0.63	0.98	0.84	0.11	0.29
U.S.	0.94	0.64	0.98	0.85	0.11	0.28
Europe	0.93	0.48	0.97	0.75	0.11	0.29
Asia	0.94	0.72	0.96	0.74	0.11	0.30
2020-2022						
Global	0.90	0.43	0.98	0.79	0.15	0.38
U.S.	0.91	0.40	0.98	0.81	0.15	0.37
Europe	0.89	0.38	0.95	0.65	0.15	0.36
Asia	0.87	0.41	0.95	0.63	0.15	0.40

Panel B: Model Performance by Sector

Region	M2B		MCAP			
	R^2		R^2		MDAPE	
	GBRT	OLS	GBRT	OLS	GBRT	OLS
2003-2015						
Energy+	0.86	0.46	0.97	0.86	0.11	0.26
Others	0.93	0.69	0.97	0.83	0.12	0.26
2016-2019						
Energy+	0.86	0.32	0.96	0.85	0.11	0.28
Others	0.94	0.63	0.98	0.84	0.11	0.30
2020-2022						
Energy+	0.83	0.28	0.98	0.78	0.15	0.37
Others	0.90	0.43	0.98	0.80	0.15	0.38

Table 3 Category Importance

This table presents feature category importance as measured by the percentage contribution of absolute SHAP values to predictions from three global models (2003–2015, 2016–2019, and 2020–2022) for all firms (All), firms in the "Energy+" sector and Non-"Energy+" firms (Others). The "Energy+" sector includes all firms classified under the "Energy" category in Fama-French-12 (FF-12) classification plus those in the electricity-related sectors of the FF-12 "Utility" industry category. The raw SHAP values are computed at the stock-month observation level for all features. The percentage contribution of each feature for each observation is calculated by dividing the absolute SHAP value of the feature by the sum of absolute SHAP values across all features. Category SHAP, as reported in the table, is the sum of the percentage SHAP values for all features within a certain category. The "Control" category contains categorical features of year, country, and FF-12 industry. Table A.3.1 provides details about category classification.

	All			Energy+			Others		
	03-15	16-19	20-22	03-15	16-19	20-22	03-15	16-19	20-22
					GHG				
Scope 1	2.17	2.12	2.44	3.36	3.51	4.05	2.00	1.97	2.30
Scope 2	0.39	0.63	0.44	0.60	1.00	0.65	0.36	0.59	0.42
					Accounting Characteristics				
Profitability	45.95	46.46	31.92	40.68	38.04	27.11	46.71	47.40	32.32
Investment	4.92	4.61	7.50	5.24	4.79	7.19	4.88	4.59	7.52
Payout	5.55	6.85	6.25	5.29	7.70	6.77	5.59	6.76	6.21
Issuance	4.13	2.90	4.90	4.58	3.12	4.87	4.06	2.87	4.90
Profit Growth	3.59	4.05	5.25	3.72	5.15	5.62	3.58	3.93	5.21
Growth*	2.13	2.05	3.24	2.29	2.29	3.22	2.11	2.02	3.24
Accruals	0.35	0.34	0.42	0.37	0.35	0.43	0.35	0.33	0.42
R&D	0.69	1.11	1.44	0.52	0.67	1.07	0.71	1.16	1.47
Capitalization	0.77	1.13	1.91	0.66	1.11	1.62	0.78	1.13	1.93
Efficiency	3.02	2.60	2.05	3.39	3.14	2.19	2.97	2.54	2.04
Financial Soundness	2.62	2.83	3.83	2.81	3.31	4.02	2.59	2.77	3.81
Solvency	1.24	1.97	2.29	1.07	1.70	2.21	1.26	2.00	2.30
Liquidity	1.02	1.01	1.22	1.11	1.17	1.18	1.01	0.99	1.22
Miscellaneous	5.11	4.96	5.06	5.17	4.71	4.63	5.10	4.99	5.10
Control	16.35	14.38	19.86	19.13	18.23	23.17	15.95	13.95	19.58

Table 4 Model-implied Scope 1 GHG Emission Prices

This table presents the estimated Scope 1 GHG emission prices at the period-region level, covering three regions (U.S., Europe and Asia) during three periods: 2003–2015, 2016–2019, and 2020–2022. We first divide a specific subsample into 10 bins and calculate the price for each bin, which is measured by the average marginal effects of Scope 1 emission levels on model-implied market capitalization. The bin prices are then averaged using the average emission levels of bins as weights. The detailed estimation procedure can be found in Section 4.4.1. The estimated emission prices for all firms in a certain period-region combination are reported in Panel A ("All"), while those for firms in the "Energy+" sector and in Non-"Energy+" sectors ("Others") are reported in Panel B and Panel C, respectively. "Energy+" sector includes firms in the "Energy" industry category as per Fama-French-12 (FF-12) classification as well as those in the electricity-related sectors of the FF-12 "Utility" industry category. All results are based on one of three global models trained using all firms in a specific period.

	USA	Europe	Asia
Panel A: All Firms			
2003-2015	12.68	11.67	14.42
2016-2019	41.34	23.29	16.88
2020-2022	148.37	53.94	34.16
Panel B: Energy+ Firms			
2003-2015	5.13	29.86	7.88
2016-2019	7.32	8.18	5.68
2020-2022	25.83	14.38	7.54
Panel C: Non-Energy+ Firms			
2003-2015	22.69	10.40	24.95
2016-2019	49.09	34.08	27.38
2020-2022	203.06	98.34	42.48

Figure 1 Flow Chart: Calculating ALE of GHG Emission Levels on Market Capitalization

This flow chart illustrates how we calculate accumulative local effects (ALE) of an emission level variable ($x_{ghg_{level}}$) on market capitalization (mcap) predicted by our GBRT-based valuation model introduced in Section 3.4.

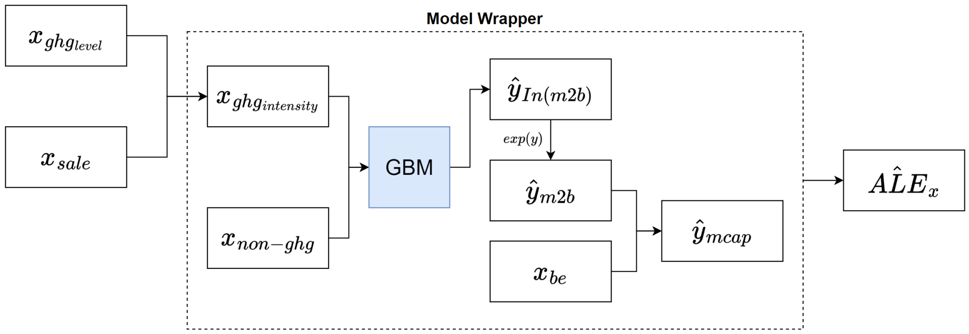

Figure 2 Feature Importance

Panel A shows the top-25 important features according to absolute SHAP values derived from three GBRT-based valuation models (each covering 2003–2015, 2016–2019 and 2020–2022 subperiods, respectively) that fit the natural logarithm of M2B. Panel B depicts the time-varying importance of Scope 1 emission intensity (left) and Scope 2 emission intensity (right) in determining log(M2B).The importance is measured by the rank of Scope 1 or 2 emission intensity in each subsample model according to absolute SHAP values, "split" and "gain," where "split" refers to the number of times a variable is used to make a split in decision trees, and "gain" refers to the reduction in training loss resulting from using the feature to create a new split point.

Panel A: Top-25 Important Variables in the Valuation Model for Predicting log(M2B)

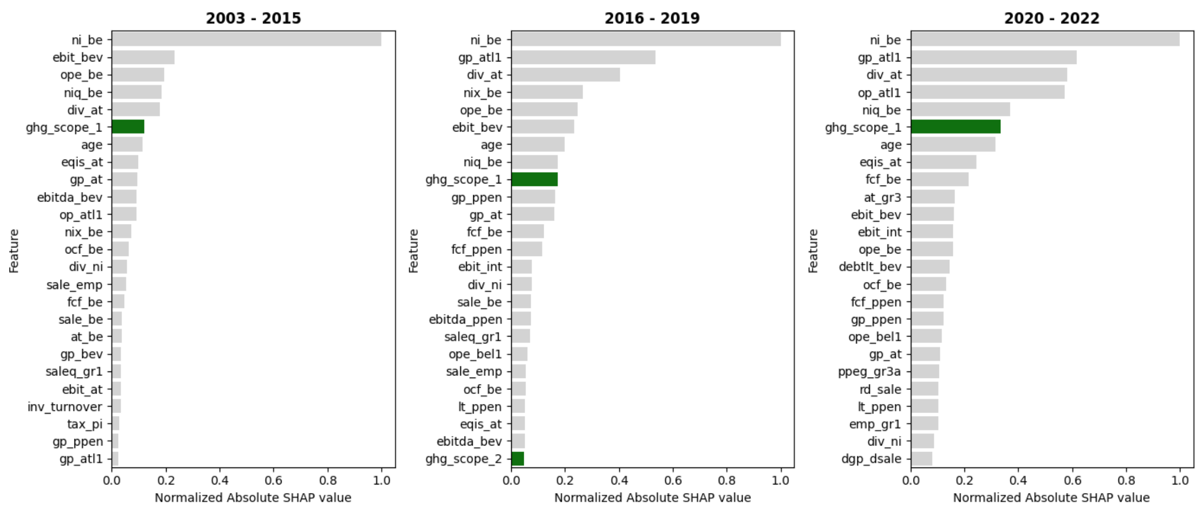

Panel B: Importance of GHG Emission Variables Over Time

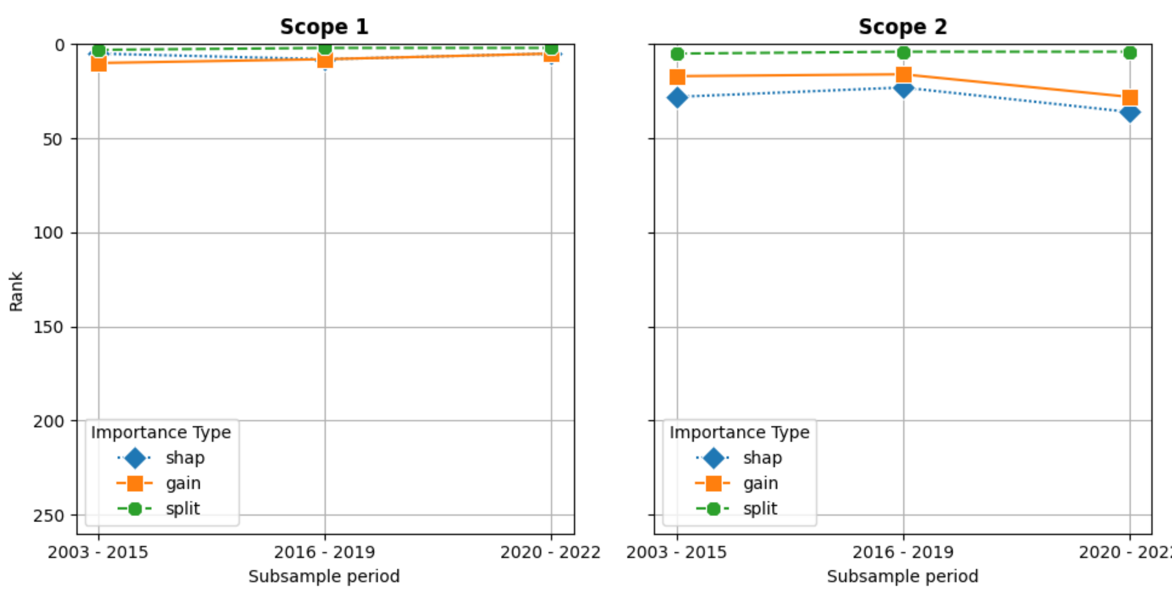

Figure 3 ALE Plots: Market-to-Book Ratio vs GHG Emission Intensity (x-axis in the log scale)

This figure includes Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission intensity (Top) and Scope 2 emission intensity (Bottom) on the market-to-book ratio for firms in the U.S., Europe or Asia during three periods: 2003–2015, 2016–2019, and 2020–2022. The x-axis is in the base-10 log scale. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission intensity variable. All plots are based on one of three subsample global models.

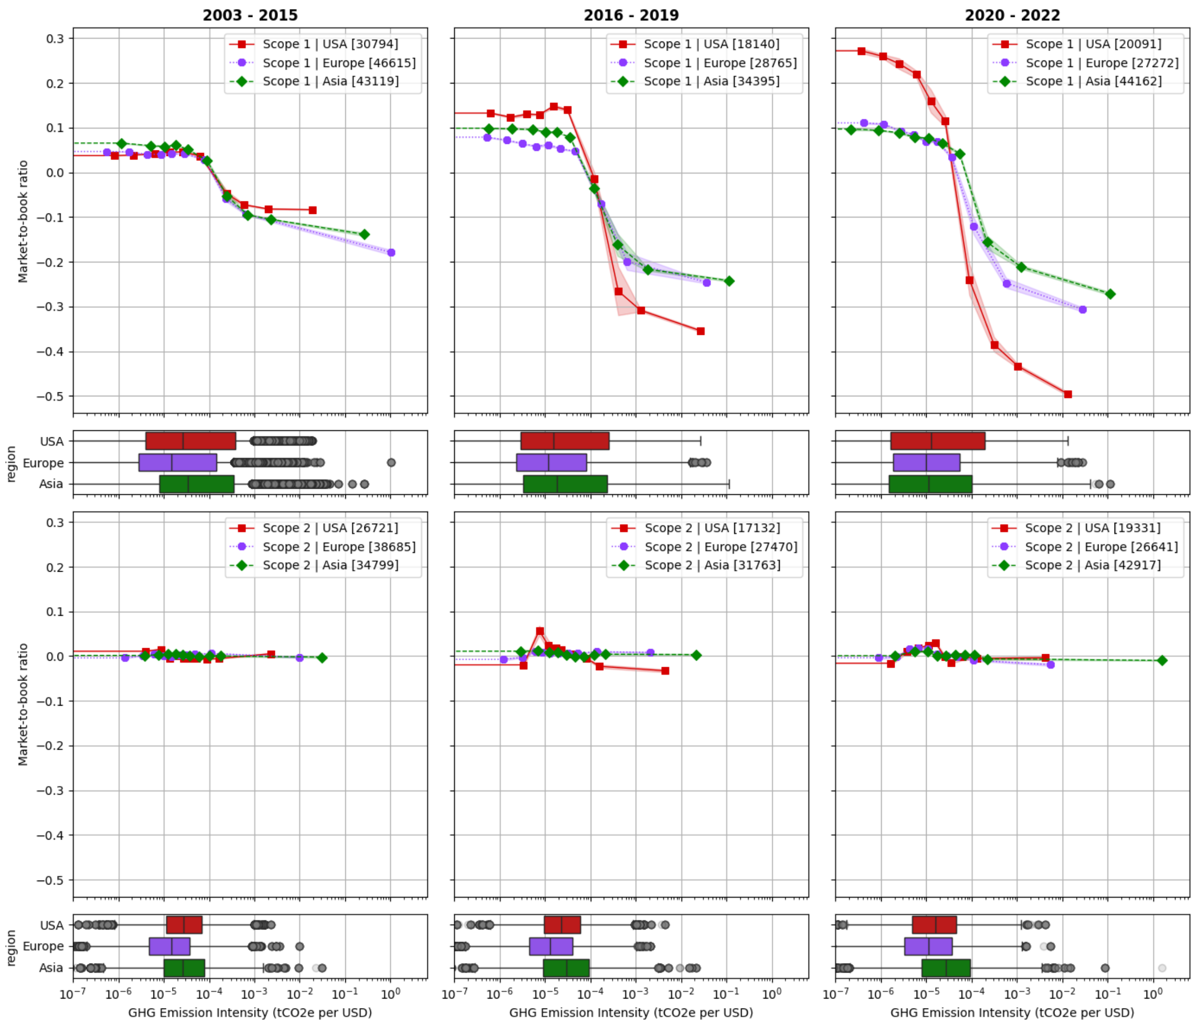

Figure 4 ALE Plots: Market-to-Book Ratio vs Scope 1 Emission Intensity

This figure shows Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission intensity on the market-to-book ratio for all firms (**Top**) and "Energy+" firms (**Bottom**) in the U.S., Europe or Asia during three periods: 2003–2015, 2016–2019, and 2020–2022. Observations with emission intensity values above the 95th percentile are omitted for better visualization. The filter is applied separately to "Energy+" firms and non-"Energy+" firms. "Energy+" firms refer to all firms in the "Energy" industry category as per Fama-French-12 (FF-12) classification as well as those in the electricity-related sectors of the FF-12 "Utility" industry category. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission intensity variable. All plots are based on one of three subsample global models.

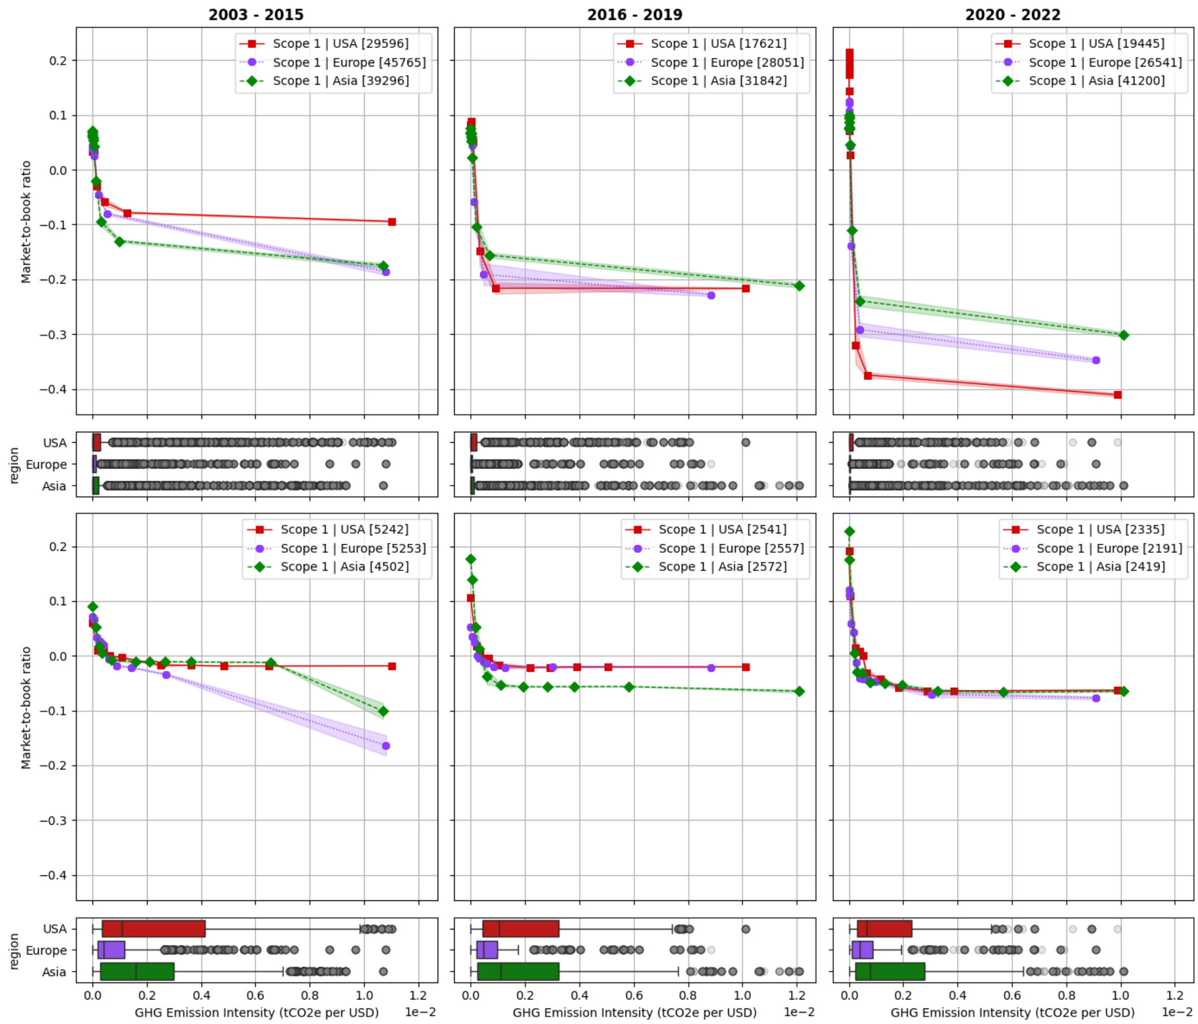

Figure 5 ALE Plots: Market Cap vs GHG Emission Levels (x-axis in the log scale)

This figure includes Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission levels (Top) and Scope 2 emission levels (Bottom) on market capitalization for firms in the U.S., Europe or Asia during three periods: 2003–2015, 2016–2019, and 2020–2022. The x-axis is in the base-10 log scale. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission level variable. All plots are based on one of three subsample global models.

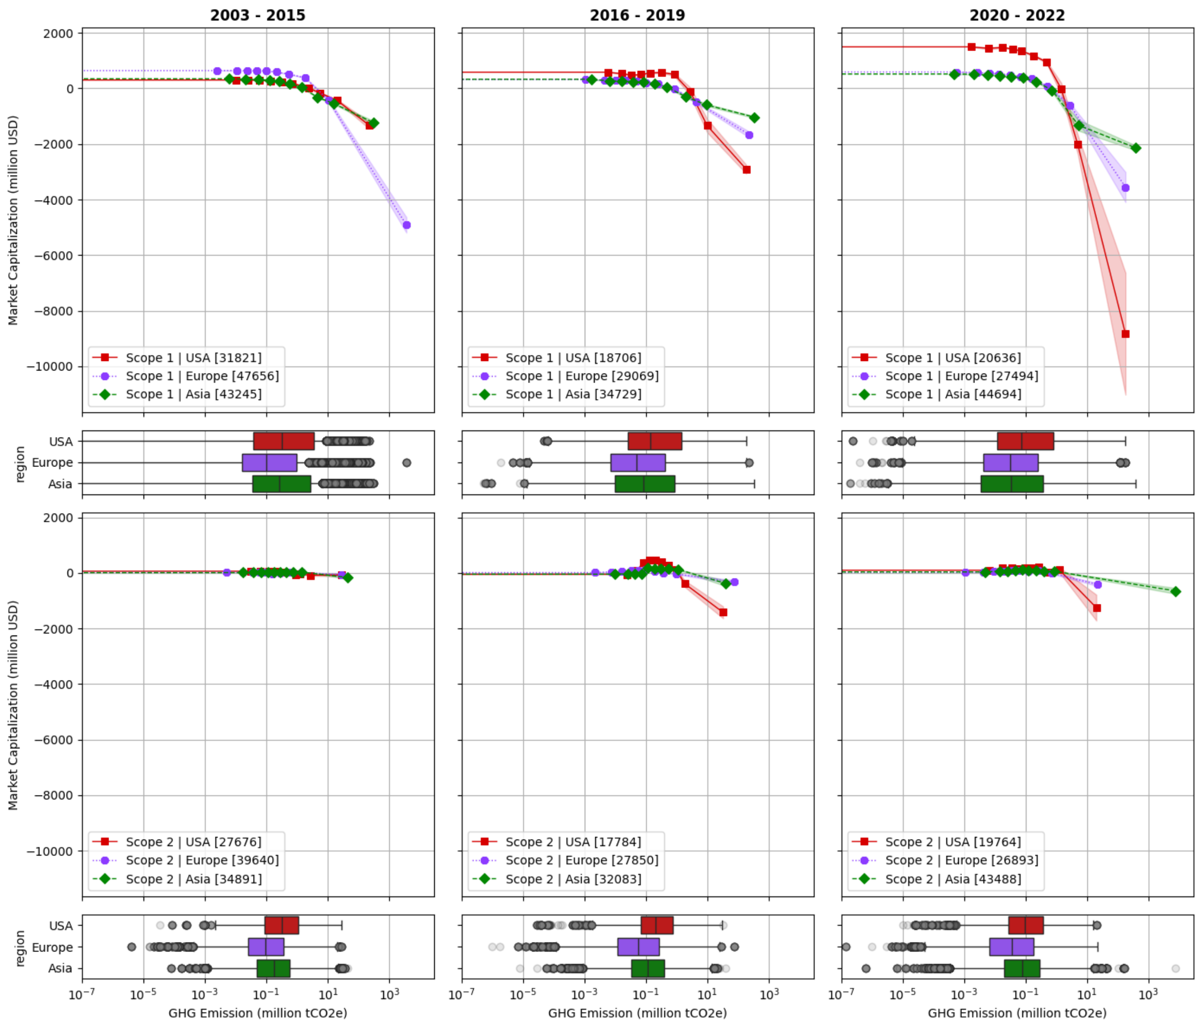

Figure 6 Market Cap vs Scope 1 GHG Emission Levels (x-axis in the log scale)

This figure shows Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission levels on market capitalization for all firms (**Top**) and "Energy+" firms (**Bottom**) in the U.S., Europe or Asia during three periods: 2003–2015, 2016–2019, and 2020–2022. "Energy+" firms refer to all firms in the "Energy" industry category as per Fama-French-12 (FF-12) classification as well as those in the electricity-related sectors of the FF-12 "Utility" industry category. The x-axis is in the base-10 log scale. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission level variable. All plots are based on one of three subsample global models.

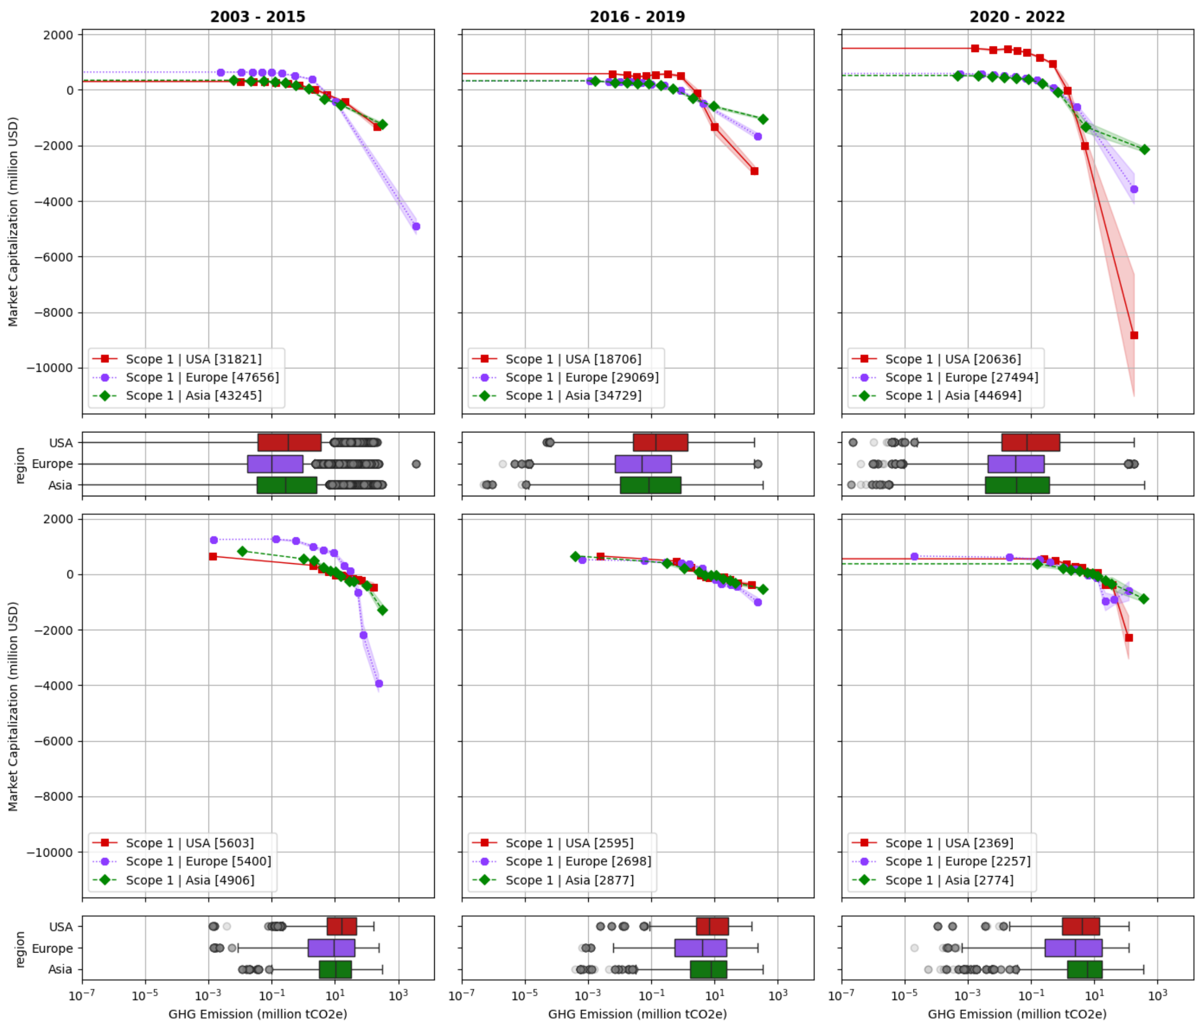

Figure 7 Feature Importance in GBRT Model for Predicting Future ROE

This figure presents the top-25 important features according to absolute SHAP values derived from a GBRT models (2003–2022) that predict firm's 5-year-ahead ROE. The features are measured during the period from 2003 to 2017, whereas the 5-year-ahead ROE is measured from 2008 to 2022

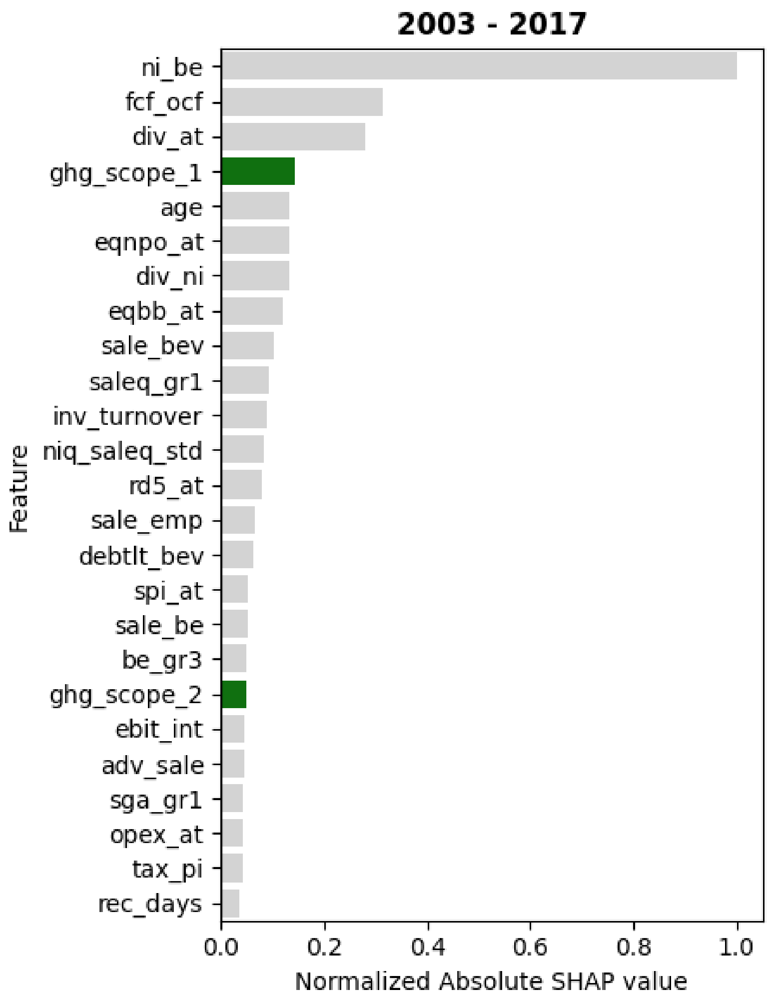

Figure 8 ALE Plots: 5-year-ahead ROE vs GHG Emission Intensity (x-axis in the log scale)

This figure includes Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission intensity (Top) and Scope 2 emission intensity (Bottom) on 5-year-ahead ROE for firms in the U.S., Europe or Asia during the 2003–2022 period. The features are measured during the period from 2003 to 2017, whereas the 5-year-ahead ROE is measured from 2008 to 2022. The x-axis is in the base-10 log scale. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission intensity variable. All plots are based on the full-sample global model.

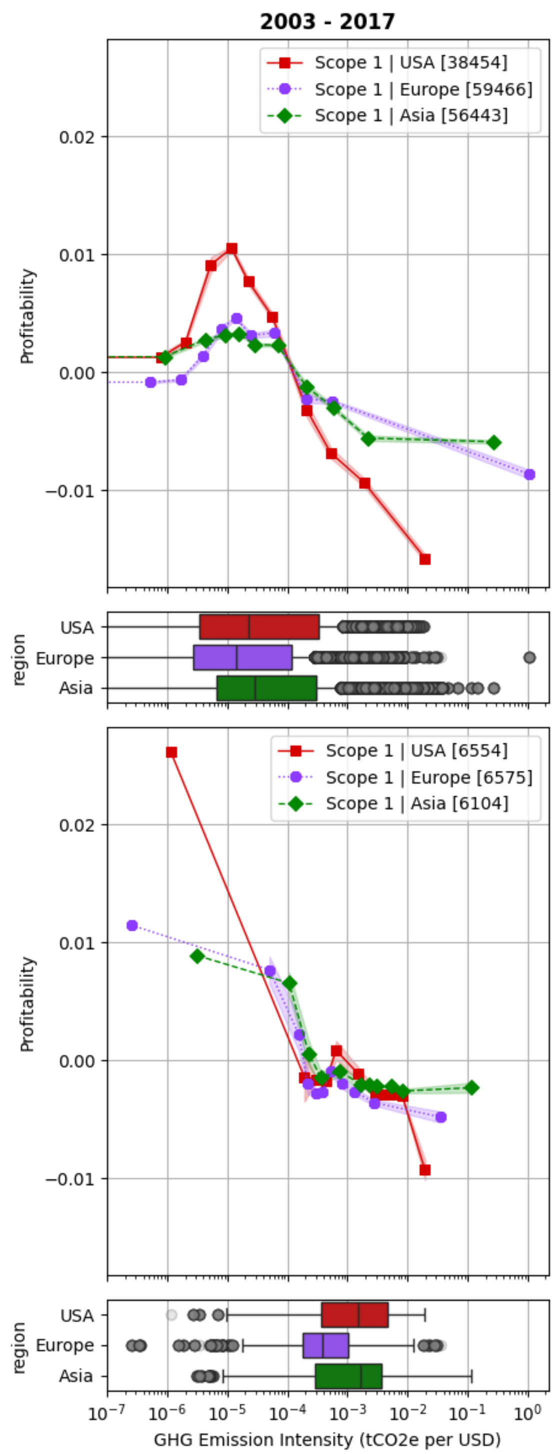

Figure 9 ALE Plots: Market-to-Book Ratio vs GHG Emission Intensity (x-axis in the log scale) – U.S. Sample

This figure includes Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission intensity (Top) and Scope 2 emission intensity (Bottom) on the market-to-book ratio for U.S. firms during three periods: 2003–2015, 2016–2019, and 2020–2022. The x-axis is in the base-10 log scale. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission intensity variable. All plots are based on one of three subsample models using U.S. stocks.

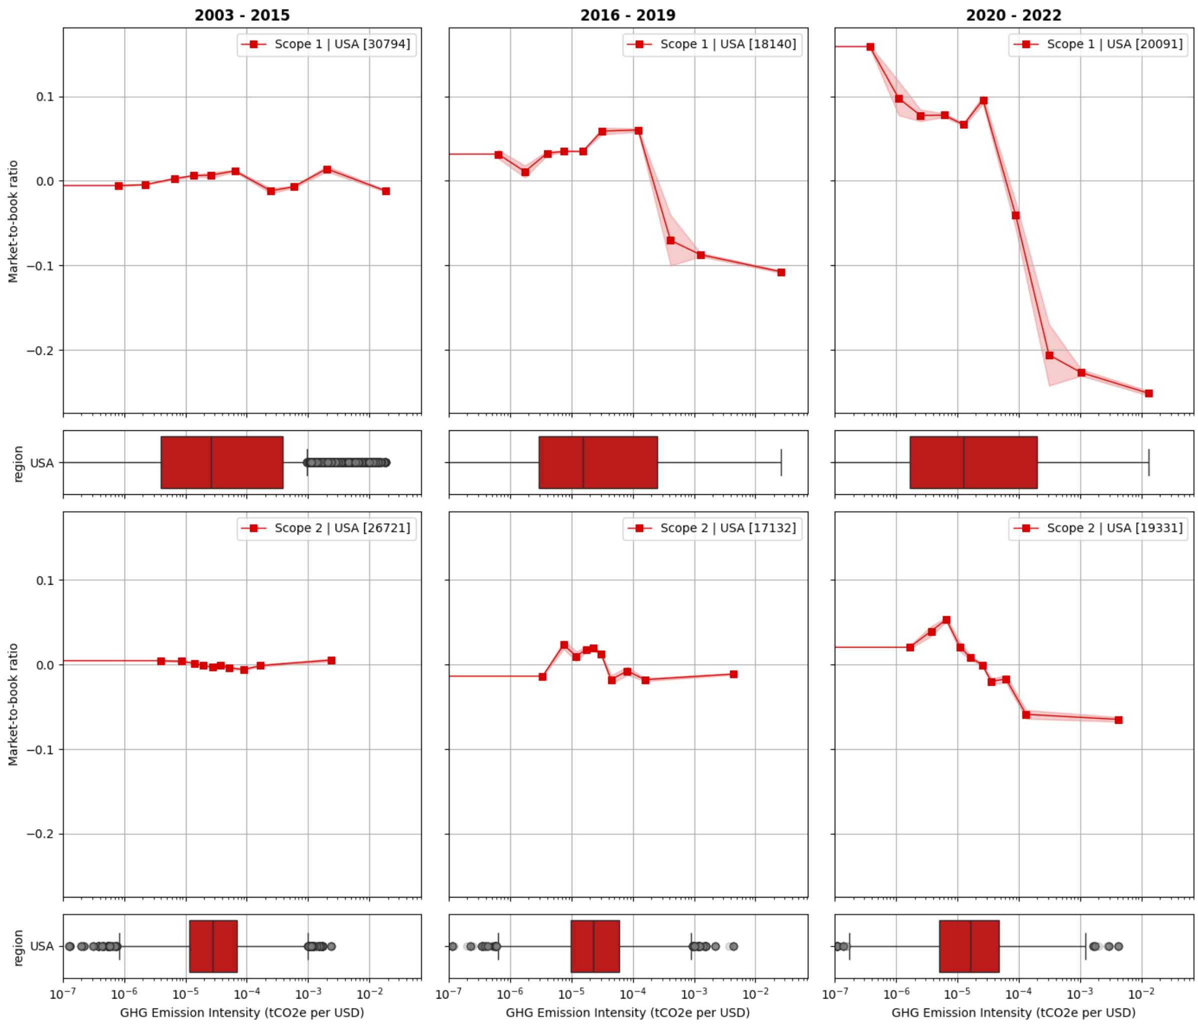

Figure 10 ALE Plots: Market-to-Book Ratio vs GHG Emission Intensity (x-axis in the log scale) – Reported GHG Values & Trucost Estimates

This figure includes Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission intensity (Top) and Scope 2 emission intensity (Bottom) on the market-to-book ratio for firms in the U.S., Europe, or Asia during three periods: 2003–2015, 2016–2019, and 2020–2022. The x-axis is in the base-10 log scale. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission intensity variable. We expand the GHG reporter sample by including stocks with GHG emission values estimated by Trucost. All plots are based on one of three subsample global models.

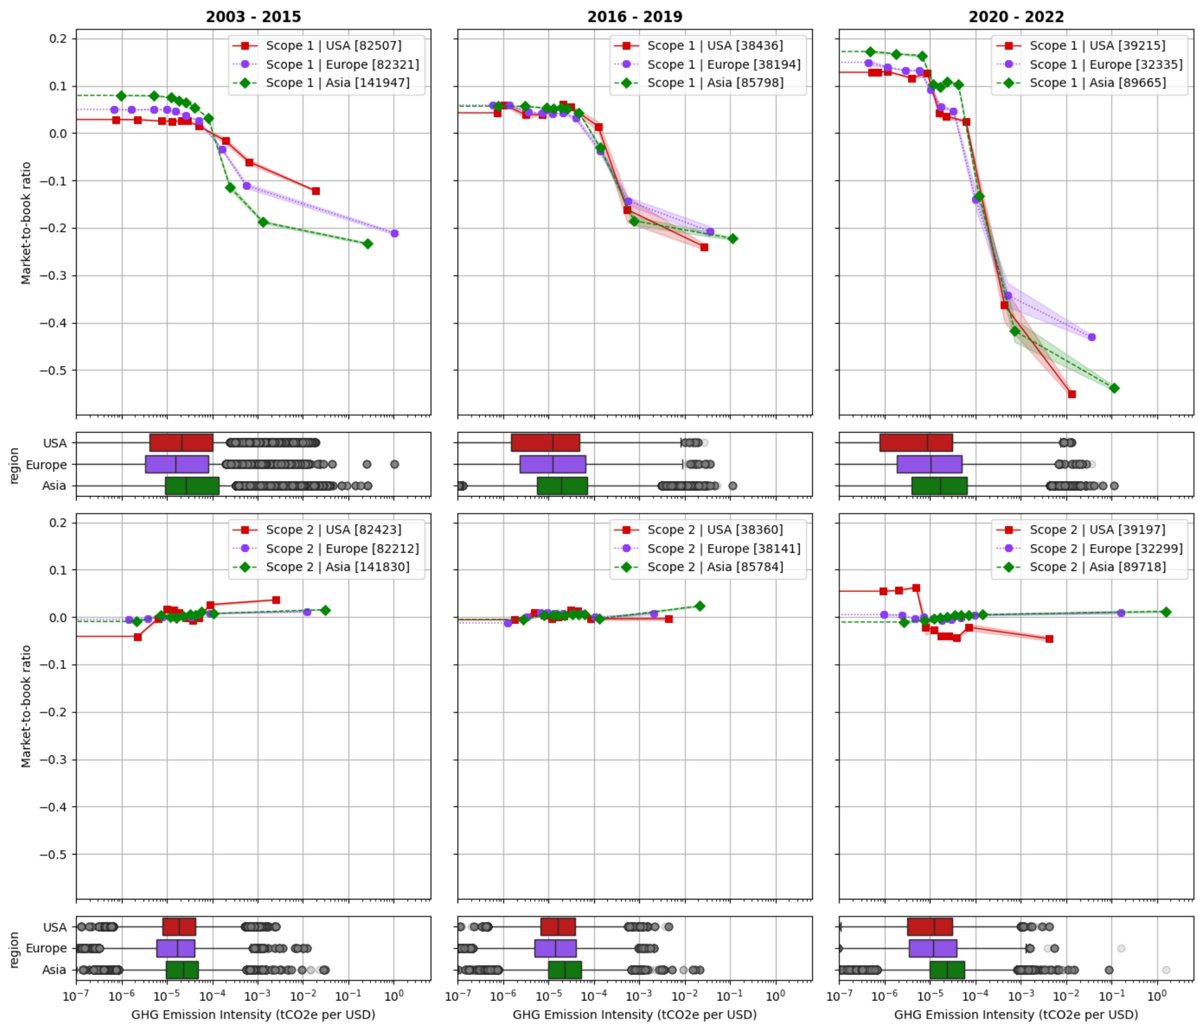

Figure 11 ALE Plots: Market-to-Book Ratio vs GHG Emission Intensity (x-axis in the log scale) – Reported GHG Values & Trucost Estimates Treated as Missing

This figure includes Accumulative Local Effect (ALE) plots for the effects of Scope 1 emission intensity (Top) and Scope 2 emission intensity (Bottom) on the market-to-book ratio for firms in the U.S., Europe, or Asia during three periods: 2003–2015, 2016–2019, and 2020–2022. The x-axis is in the base-10 log scale. Each ALE plot is supplemented with a chart below showing the distribution of the GHG emission intensity variable. We expand the GHG reporter sample by including stocks with GHG emission values estimated by Trucost but treat the estimated values as missing. All plots are based on one of three subsample global models.

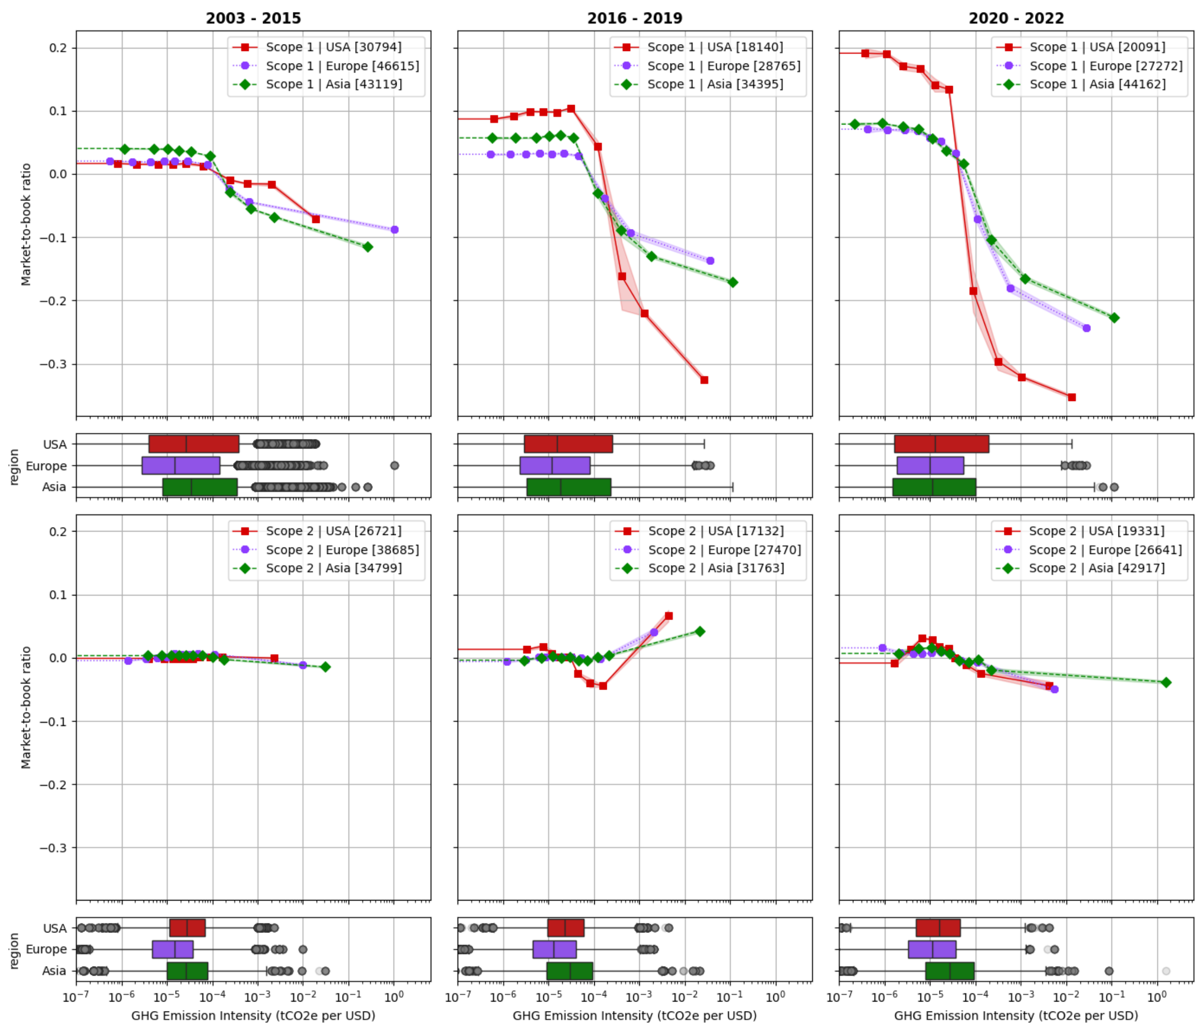

Appendix

A.1 Accounting Characteristics and Cluster Classification

Table A.3.1 presents names and clusters of 258 accounting characteristic along with firm age used in our study. Geertsema and Lu (2023) mainly rely on financial ratios constructed by Wharton Research Data Services (WRDS), which classify their ratios into 6 categories: Profitability, Capitalization, Financial Soundness, Solvency, Liquidity, and Efficiency. Note that WRDS' financial ratios are included in the dataset of Jensen, Kelly, and Pedersen (2023). Geertsema and Lu (2023) create another category, Growth, for variables measuring changes in accounting items. We break down their Growth category into Investment (e.g., asset growth), Profit Growth (e.g., change in Return on Equity), Accruals (e.g., change in net working capital), Payout (i.e., growth in equity or dividend payout), Issuance (i.e., growth in debt or equity issuance) and Growth* (for growth in other accounting items), following Hou, Xue, and Zhang (2020) and Jensen, Kelly, and Pedersen (2023).[21] The categories of Accruals, Payout, Issuance, and R&D are created by us.

[Insert Table A.3.1 near here]

[21]The categories of Investment, Accruals, Payout and Issuance also include variables not measuring changes in accounting items.

A.2 Model Training Details

A.2.1 Gradient Boosting Regression Trees (GBRT) Algorithm

Gradient Boosting Regression Trees (GBRT)

Inputs: $\{(x_i, y_i)\}_{i=1}^{N}$ and a differentiable loss function, $L(\cdot)$, such as mean square error/quadratic loss.
Step 1: Model Initialized with a constant value such as the mean value of the target variable:

$$F_0(x) = \arg\min_{\gamma} \sum_{i=1}^{N} L(y_i, \gamma)$$

Step 2:
for m = 1 to M **do**

Calculate residuals (i.e., $r_{im} = y_i - F_{m-1}(x_i)$) for each observation i, whcih is equivalent to:

$$r_{im} = -\left[\frac{\partial L(y_i, F(x_i))}{\partial F(x_i)}\right]_{F(x)=F_{m-1}(x)}$$

Fit a decision tree to the pseudo residuals ($r_{i,m}$) using x_i:

$$G_m(x) = \text{FittedTree}(\{(x_i, r_{im})\}_{i=1}^{N})$$

Update the model prediction with the learning rate of ν:

$$F_m(x) = F_{m-1}(x) + \nu G_m(x)$$

end for
Final model prediction:

$$F_M(x) = F_0(x) + \nu \sum_{m=1}^{M} G_m(x)$$

A.2.2 Model Pipeline

The pipeline of our valuation model is built using he 'Pipeline' class from the scikit-learn Python package. The pipeline also acts as a wrapper for model-agnostic algorithms including SHAP and ALE. Additionally, this enables the estimation of not only the effects of the features used in the LightGBM (e.g., emission intensity variables) on the model target (i.e, the market-to-book ratio) but also the effects of emission levels $[tCO_2e]$ on Market Capitalization $[USD]$.

The model pipeline comprised of 3 modules: 1) the GHG transformer, 2) feature transformer, and (3) the target transformer. The GHG transformer houses operations applied to the GHG-emission-related features only including transformation of emission levels to emission intensity by dividing levels by revenues from the corresponding fiscal year. The feature transformer is implemented using the 'ColumnTransformer' class and designed to apply specific transformations to a subset of model inputs. Within the feature transformer, redundant features such as revenues and book equity are dropped. These features are intentionally fed into the pipeline together with the training features to eliminate potential errors due to incorrect data indexing. Lastly, the target transformer is implemented using the 'TransformedTargetRegressor' class wraps an estimator and a mathematical operation (the log transformation) and its corresponding inverse operation (the exponential transformation) which was applied to the target variable, the market-to-book ratio, before and after model training and inference.

A.2.3 Model Selection

The search for a close-to-optimal set of model hyperparameters is conducted using Bayesian-optimization-based Optuna Python package. Bayesian optimization is arguably favored over other common selection algorithms such as grid search in terms of convergence speed and implementation difficulty. To make the optimization more stable and robust, the learning rate is fixed to the default value of 0.1. Some hyperparameters such as 'min_split_gain' are discretized to reduce the span of the search space. A total of 1000 trial sets of hyperparameters are tested in the search. In each trial, the model is optimized and evaluated using root-mean-square-error (RMSE). Early stopping is set to 10 rounds with a minimum delta of 0.01 in improvement between optimization steps. The search spaces of hyperparameters and selected set for each subsample valuation is summarized in Table A.3.2.

[Insert Table A.3.2 near here]

A.3 Additional Figures and Tables

Table A.3.1 List of Accounting Characteristics and their Clusters

Var	Name	Cluster
aliq_at	Liquidity scaled by lagged Assets	Investment
at_gr1	Asset Growth 1yr	Investment
at_gr3	Asset Growth 3yr	Investment
be_gr1	Book Equity Growth 1yr	Investment
be_gr1a	Book Equity Change 1 yr scaled by Assets	Investment
be_gr3	Book Equity Growth 3yr	Investment
ca_gr1	Current Asset Growth 1yr	Investment
ca_gr3	Current Asset Growth 3yr	Investment
capx_at	Capital Expenditures scaled by Assets	Investment
capx_gr1	CAPX 1 year growth	Investment
capx_gr1a	Capital Expenditures Change 1yr	Investment
capx_gr2	CAPX 2 year growth	Investment
capx_gr3	Capital Expenditures Growth 3yr	Investment
capx_gr3a	Capital Expenditures Change 3yr	Investment
coa_gr1a	Current Operating Assets Change 1yr	Investment
coa_gr3a	Current Operating Assets Change 3yr	Investment
emp_gr1	Employee Growth 1 yr	Investment
fna_gr1a	Financial Assets Change 1yr	Investment
fna_gr3a	Financial Assets Change 3yr	Investment
intan_gr1a	Intangible Assets Change 1yr	Investment
intan_gr3a	Intangible Assets Change 3yr	Investment
inv_gr1	Inventory Change 1 yr	Investment
inv_gr1a	Inventory Change 1yr	Investment
inv_gr3a	Inventory Change 3yr	Investment
lnoa_gr1a	Change in Long-Term NOA scaled by average Assets	Investment
lti_gr1a	Investment and Advances Change 1yr	Investment
lti_gr3a	Investment and Advances Change 3yr	Investment
nca_gr1	Non-Current Asset Growth 1yr	Investment
nca_gr3	Non-Current Asset Growth 3yr	Investment
ncoa_gr1a	Non-Current Operating Assets Change 1yr	Investment
ncoa_gr3a	Non-Current Operating Assets Change 3yr	Investment
nfna_gr1a	Net Financial Assets Change 1yr	Investment
nfna_gr3a	Net Financial Assets Change 3yr	Investment
nncoa_gr1a	Net Non-Current Operating Assets Change 1yr	Investment
nncoa_gr3a	Net Non-Current Operating Assets Change 3yr	Investment
noa_gr1a	Change in net operating assets	Investment
oa_gr1a	Operating Assets Change 1yr	Investment
oa_gr3a	Operating Assets Change 3yr	Investment
ppeg_gr1a	Property, Plans and Equiptment Gross Change 1yr	Investment
ppeg_gr3a	Property, Plans and Equipment Gross Change 3yr	Investment
ppeinv_gr1a	Change in Property, Plant and Equipment Less Inventories scaled by lagged Assets	Investment
sale_gr1	Sales Growth 1yr	Investment
sale_gr3	Sales Growth 3yr	Investment
saleq_gr1	Quarterly Sales Growth	Investment
sti_gr1a	Change in short-term investments	Investment
capex_abn	Abnormal Corporate Investment	Issuance
dbnetis_at	Net Debt Issuance scaled by Assets	Issuance
dbnetis_gr1a	Net Debt Issuance Change 1yr	Issuance
dbnetis_gr3a	Net Debt Issuance Change 3yr	Issuance
debt_gr1	Total Debt Growth 1yr	Issuance
debt_gr3	Total Debt Growth 3yr	Issuance
debtlt_gr1a	Long-Term Debt Change 1yr	Issuance
debtlt_gr3a	Long-Term Debt Change 3yr	Issuance
debtst_gr1a	Short-Term Debt Change 1yr	Issuance
debtst_gr3a	Short-Term Debt Change 3yr	Issuance
dltnetis_at	Net Long-Term Debt Issuance scaled by Assets	Issuance
dltnetis_gr1a	Net Long-Term Debt Issuance Change 1yr	Issuance
dltnetis_gr3a	Net Long-Term Debt Issuance Change 3yr	Issuance
dstnetis_at	Net Short-Term Debt Issuance scaled by Assets	Issuance
dstnetis_gr1a	Net Short-Term Debt Issuance Change 1yr	Issuance
dstnetis_gr3a	Net Short-Term Debt Issuance Change 3yr	Issuance

Var	Name	Cluster
eqis_at	Equity Issuance scaled by Assets	Issuance
eqis_gr3a	Equity Issuance Change 3yr	Issuance
eqnetis_at	Equity Net Issuance scaled by Assets	Issuance
eqnetis_gr1a	Equity Net Issuance Change 1yr	Issuance
eqnetis_gr3a	Equity Net Issuance Change 3yr	Issuance
fincf_at	Financial Cash Flow scaled by Assets	Issuance
fincf_gr1a	Financial Cash Flow Change 1yr	Issuance
fincf_gr3a	Financial Cash Flow Change 3yr	Issuance
fnl_gr1a	Financial Liabilities Change 1yr	Issuance
lt_gr1	Total Liabilities Growth 1yr	Issuance
lt_gr3	Total Liabilities Growth 3yr	Issuance
ncl_gr1	Non-Current Liabilities Growth 1yr	Issuance
ncl_gr3	Non-Current Liabilities Growth 3yr	Issuance
ncol_gr1a	Non-Current Operating Liabilities Change 1yr	Issuance
ncol_gr3a	Non-Current Operating Liabilities Change 3yr	Issuance
netis_at	Net Issuance scaled by Assets	Issuance
netis_gr1a	Net Issuance Change 1yr	Issuance
netis_gr3a	Net Issuance Change 3yr	Issuance
ni_ar1	Earnings persistence	Issuance
noa_at	Net Operating Asset to Total Assets	Issuance
pstk_gr1	Preferred Stock Growth 1yr	Issuance
pstk_gr3	Preferred Stock Growth 3yr	Issuance
cop_at	Cash Based Operating Profitability scaled by Assets	Profitability
cop_atl1	Cash Based Operating Profitability scaled by lagged Assets	Profitability
cop_bev	Cash Based Operating Profitability scaled by BEV	Profitability
ebit_at	Operating Profit after Depreciation scaled by Assets	Profitability
ebit_bev	Operating Profit after Depreciation scaled by BEV	Profitability
ebit_sale	Operating Profit Margin after Depreciation	Profitability
ebitda_at	Operating Profit before Depreciation scaled by Assets	Profitability
ebitda_bev	Operating Profit before Depreciation scaled by BEV	Profitability
ebitda_ppen	Operating Profit before Depreciation scaled by PPEN	Profitability
ebitda_sale	Operating Profit Margin before Depreciation	Profitability
fcf_be	Free Cash Flow scaled by BE	Profitability
fcf_ppen	Free Cash Flow scaled by PPEN	Profitability
fcf_sale	Operating Cash Flow Margin	Profitability
fi_at	Firm Income scaled by Assets	Profitability
fi_bev	Firm Income scaled by BEV	Profitability
gp_at	Gross Profit scaled by Assets	Profitability
gp_atl1	Gross Profit scaled by lagged Assets	Profitability
gp_bev	Gross Profit scaled by BEV	Profitability
gp_ppen	Gross Profit scaled by PPEN	Profitability
gp_sale	Gross Profit Margin	Profitability
ni_at	Net Income scaled by Assets	Profitability
ni_be	Net Income scaled by BE	Profitability
ni_emp	Net Income scaled by Employees	Profitability
ni_sale	Net Profit Margin before XI	Profitability
niq_at	Quarterly Income scaled by AT	Profitability
niq_be	Quarterly Income scaled by BE	Profitability
nix_be	Net Income Including Extraordinary Items scaled by BE	Profitability
nix_sale	Net Profit Margin	Profitability
ocf_at	Operating Cash Flow scaled by Assets	Profitability
ocf_be	Operating Cash Flow scaled by BE	Profitability
ocf_sale	Free Cash Flow Margin	Profitability
op_at	Operating profits-to-book assets	Profitability
op_atl1	Ball Operating Profit scaled by lagged Assets	Profitability
ope_be	Operating Profit to Equity scaled by BE	Profitability
ope_bel1	Operating Profit scaled by lagged Book Equity	Profitability
pi_sale	Pretax Profit Margin	Profitability
cfoa_ch5	Operating Cash Flow to Assets 5 yr Change	Profit Growth
dgp_dsale	Change Gross Profit minus Change Sales	Profit Growth
dsale_dinv	Change Sales minus Change Inventory	Profit Growth
dsale_drec	Change Sales minus Change Receivables	Profit Growth
dsale_dsga	Change Sales minus Change SG&A	Profit Growth
ebit_gr1a	Operating Profit after Depreciation Change 1yr	Profit Growth
ebit_gr3a	Operating Profit after Depreciation Change 3yr	Profit Growth
ebitda_gr1a	Operating Profit before Depreciation Change 1yr	Profit Growth
ebitda_gr3a	Operating Profit before Depreciation Change 3yr	Profit Growth

Var	Name	Cluster
fcf_gr1a	Free Cash Flow Change 1yr	Profit Growth
fcf_gr3a	Free Cash Flow Change 3yr	Profit Growth
gmar_ch5	Gross Profit to Sales 5 yr Change	Profit Growth
gp_gr1a	Gross Profit Change 1yr	Profit Growth
gp_gr3a	Gross Profit Change 3yr	Profit Growth
gpoa_ch5	Gross Profit to Assets 5 yr Change	Profit Growth
ni_gr1a	Net Income Change 1yr	Profit Growth
ni_gr3a	Net Income Change 3yr	Profit Growth
ni_inc8q	Number of Consecutive Earnings Increases	Profit Growth
niq_at_chg1	Change in Quarterly Income scaled by AT	Profit Growth
niq_be_chg1	Change in Quarterly Income scaled by BE	Profit Growth
nix_gr1a	Net Income Including Extraordinary Items Change 1yr	Profit Growth
nix_gr3a	Net Income Including Extraordinary Items Change 3yr	Profit Growth
ocf_at_chg1	Operating Cash Flow to Assets 1 yr Change	Profit Growth
ocf_gr1a	Operating Cash Flow Change 1yr	Profit Growth
ocf_gr3a	Operating Cash Flow Change 3yr	Profit Growth
ope_gr1a	Operating Earnings to Equity Change 1yr	Profit Growth
ope_gr3a	Operating Earnings to Equity Change 3yr	Profit Growth
roa_ch5	ROA 5 yr Change	Profit Growth
roe_ch5	ROE 5 yr Change	Profit Growth
sale_emp_gr1	Sales scaled by Employees Growth 1 yr	Profit Growth
saleq_su	Revenue Surprise	Profit Growth
tax_gr1a	Effective Tax Rate Change 1yr	Profit Growth
tax_gr3a	Effective Tax Rate Change 3yr	Profit Growth
ap_gr1a	Accounts Payable Change 1yr	Growth*
ap_gr3a	Accounts Payable Change 3yr	Growth*
cash_gr1a	Cash and Short-Term Investments Change 1yr	Growth*
cash_gr3a	Cash and Short-Term Investments Change 3yr	Growth*
cl_gr1	Current Liabilities Growth 1yr	Growth*
cl_gr3	Current Liabilities Growth 3yr	Growth*
cogs_gr1	Cost of Goods Sold Growth 1yr	Growth*
cogs_gr3	Cost of Goods Sold Growth 3yr	Growth*
col_gr1a	Current Operating Liabilities Change 1yr	Growth*
col_gr3a	Current Operating Liabilities Change 3yr	Growth*
dp_gr1a	Depreciation and Amortization Change 1yr	Growth*
dp_gr3a	Depreciation and Amortization Change 3yr	Growth*
ol_gr1a	Operating Liabilities Change 1yr	Growth*
ol_gr3a	Operating Liabilities Change 3yr	Growth*
opex_gr1	Operating Expenses Growth 1yr	Growth*
opex_gr3	Operating Expenses Growth 3yr	Growth*
rec_gr1a	Receivables Change 1yr	Growth*
rec_gr3a	Receivables Change 3yr	Growth*
sga_gr1	Selling, General, and Administrative Expenses Growth 1yr	Growth*
sga_gr3	Selling, General, and Administrative Expenses Growth 3yr	Growth*
txditc_gr1a	Deferred Taxes and Investment Credit Change 1yr	Growth*
txditc_gr3a	Deferred Taxes and Investment Credit Change	Growth*
txp_gr1a	Income Tax Payable Change 1yr	Growth*
txp_gr3a	Income Tax Payable Change 3yr	Growth*
cash_lt	Cash Balance scaled by Total Liabilities	Financial Soundness
cl_lt	Current Liabilities scaled by Total Liabilities	Financial Soundness
debtlt_be	Long-Term Debt to Book Equity	Financial Soundness
debtlt_debt	Long-Term Debt scaled by Total Debt	Financial Soundness
debtst_debt	Short-Term Debt scaled by Total Debt	Financial Soundness
ebitda_debt	Operating Profit before Depreciation scaled by Total Debt	Financial Soundness
fcf_ocf	Free Cash Flow scaled by Operating Cash Flow	Financial Soundness
int_debt	Interest scaled by Total Debt	Financial Soundness
int_debtlt	Interest scaled by Long-Term Debt	Financial Soundness
inv_act	Inventory scaled by Current Assets	Financial Soundness
lt_ppen	Total Liabilities scaled by Total Tangible Assets	Financial Soundness
nwc_at	Working Capital scaled by Assets	Financial Soundness
ocf_cl	Operating Cash Flow scaled by Current Liabilities	Financial Soundness
ocf_debt	Operating Cash Flow scaled by Total Debt	Financial Soundness
opex_at	Operating Leverage	Financial Soundness
profit_cl	Profit before D&A scaled by Current Liabilities	Financial Soundness
rec_act	Receivables scaled by Current Assets	Financial Soundness

Var	Name	Cluster
adv_sale	Advertising scaled by Sales	Miscellaneous
age	Age	Miscellaneous
earnings_variability	Earnings Variability	Miscellaneous
ni_ivol	Net Income Idiosyncratic Volatility	Miscellaneous
niq_saleq_std	Net Income to Sales Quarterly Volatility	Miscellaneous
nri_at	Non-Recurring Items scaled by Assets	Miscellaneous
ocfq_saleq_std	Operating Cash Flow to Sales Quarterly Volatility	Miscellaneous
pi_nix	Taxable income-to-book income	Miscellaneous
roe_be_std	ROE Volatility	Miscellaneous
roeq_be_std	Quarterly ROE Volatility	Miscellaneous
spi_at	Special Items scaled by Assets	Miscellaneous
staff_sale	Labor Expense scaled by Sales	Miscellaneous
tax_pi	Effective Tax Rate	Miscellaneous
xido_at	Extraordinary Items and Discontinued Operations scaled by Assets	Miscellaneous
div_at	Total Dividends scaled by Assets	Payout
div_gr3a	Dividend Payout Ratio Change 3yr	Payout
div_ni	Dividend Payout Ratio	Payout
eqbb_at	Net Equity Payout scaled by Assets	Payout
eqbb_gr1a	Equity Buyback Change 1yr	Payout
eqbb_gr3a	Equity Buyback Change 3yr	Payout
eqis_gr1a	Equity Issuance Change 1yr	Payout
eqnpo_at	Equity Net Payout scaled by Assets	Payout
eqnpo_gr1a	Equity Net Payout Change 1yr	Payout
eqnpo_gr3a	Equity Net Payout Change 3yr	Payout
eqpo_gr1a	Net Equity Payout Change 1yr	Payout
eqpo_gr3a	Net Equity Payout Change 3yr	Payout
cowc_gr1a	Current Operating Working Capital Change	Accruals
cowc_gr3a	Current Operating Working Capital Change 3yr	Accruals
nwc_gr1a	Net Working Capital Change 1yr	Accruals
nwc_gr3a	Net Working Capital Change 3yr	Accruals
oaccruals_at	Operating Accruals	Accruals
oaccruals_ni	Percent Operating Accruals	Accruals
taccruals_at	Total Accruals	Accruals
taccruals_ni	Percent Total Accruals	Accruals
ap_turnover	Account Payables Turnover	Efficiency
at_turnover	Asset Turnover	Efficiency
inv_turnover	Inventory Turnover	Efficiency
rec_turnover	Receivables Turnover	Efficiency
sale_be	Sales scaled by Total Stockholders' Equity	Efficiency
sale_bev	Sales scaled by BEV	Efficiency
sale_emp	Sales scaled by Employees	Efficiency
sale_nwc	Sales scaled by Working Capital	Efficiency
ap_days	Days Accounts Payable Outstanding	Liquidity
ca_cl	Current Ratio	Liquidity
caliq_cl	Quick Ratio	Liquidity
cash_at	Cash and Short Term Investments scaled by Assets	Liquidity
cash_cl	Cash Ratio	Liquidity
cash_conversion	Cash Conversion Cycle	Liquidity
inv_days	Days Inventory Outstanding	Liquidity
rec_days	Days Sales Outstanding	Liquidity
be_bev	Common Equity scaled by BEV	Capitalization
cash_bev	Cash and Short-Term Investments scaled by BEV	Capitalization
debt_bev	Total Debt scaled by BEV	Capitalization
debtlt_bev	Long-Term Debt scaled by BEV	Capitalization
debtst_bev	Short-Term Debt scaled by BEV	Capitalization
pstk_bev	Preferred Stock scaled by BEV	Capitalization
at_be	Book Leverage	Solvency
debt_at	Debt-to-Assets	Solvency
debt_be	Debt to Shareholders' Equity Ratio	Solvency
ebit_int	Interest Coverage Ratio	Solvency
rd5_at	R&D Capital-to-Assets	R&D
rd_at	R&D scaled by Assets	R&D
rd_sale	R&D scaled by Sales	R&D

Figure A.3.1 Number of Observations by GHG Reporting Status

This figure depicts the number of non-micro-cap global stocks with reported Scope 1 or 2 GHG emission levels ("Reported") and stocks with Scope 1 or 2 GHG emission levels estimated by Trucost ("Estimated") as well as stocks with either missing Scope 1 emission value or missing Scope 2 emission value over the 2003-2022 period. Stocks with missing values for both Scope 1 and Scope 2 emission levels are excluded from the sample.

Panel A: Scope 1 Emission Data

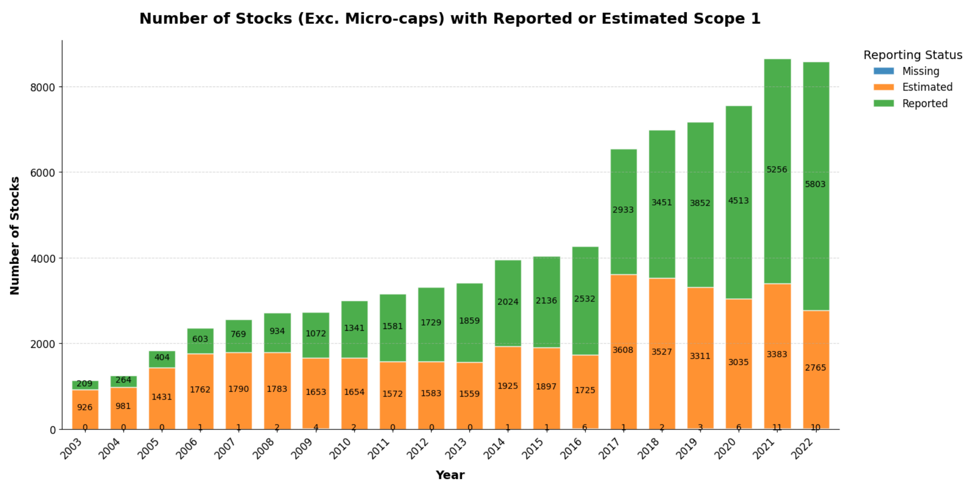

Panel B: Scope 2 Emission Data

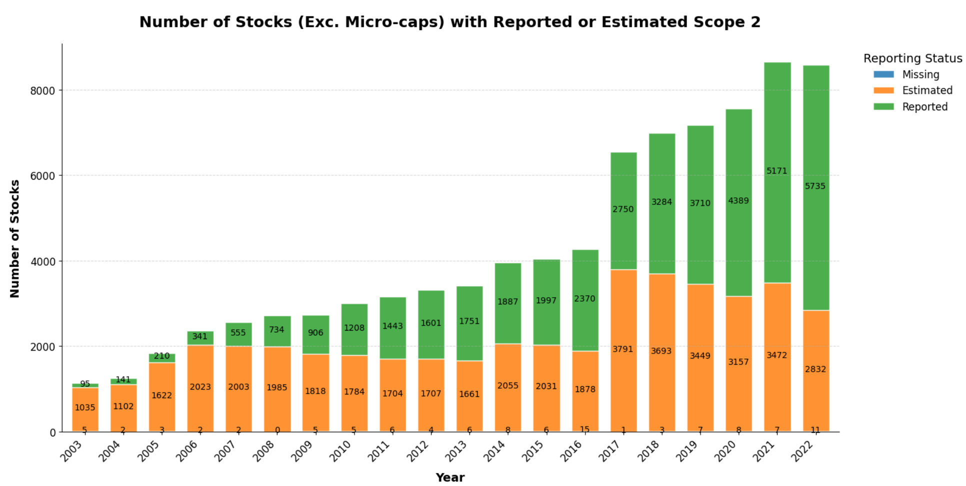

Table A.3.2 Hyperparameter Information

This table presents the hyperparameters involved in our training, their search spaces, their default values in LightGBM, and the selected values for three models (separated by '|').

Parameter	**min**	**max**	**step size**	**default**	**selected**
n_estimators	-	-	-	100	53|56|69
extra_trees	-	-	-	False	True
num_leaves	2	256	1	31	256|251|247
subsample	0.1	1.0	0.1	1.0	0.8|1.0|1.0
subsample_freq	1	7	1	0	3|3|6
min_child_samples	64	512	1	20	64|76|64
min_split_gain	0.0	0.6	0.1	0.0	0.0|0.0|0.0
max_bin	10	255	1	255	130|165|220
reg_alpha	10^{-8}	10.0	log, continuous	0.0	.0013|.000038|$1.46 * 10^{-7}$
reg_lambda	10^{-8}	10.0	log, continuous	0.0	.00048|.000012|0.67
min_child_weight	10^{-3}	10.0	log, continuous	10^{-3}	3.76|0.26|3.56
max_depth	-1	32	1	-1	19|21|31
path_smooth	0.0	1.0	continuous	0.0	0.55|0.50|0.44

Figure A.3.2 Feature Importance: U.S. Sample

This figure depicts the time-varying importance of Scope 1 emission intensity (left) and Scope 2 emission intensity (right) in determining log(M2B) for U.S. firms. The importance is measured by the rank of Scope 1 or 2 emission intensity in each subsample model according to absolute SHAP values, "split" and "gain," where "split" refers to the number of times a variable is used to make a split in decision trees, and "gain" refers to the reduction in training loss resulting from making a new split point.

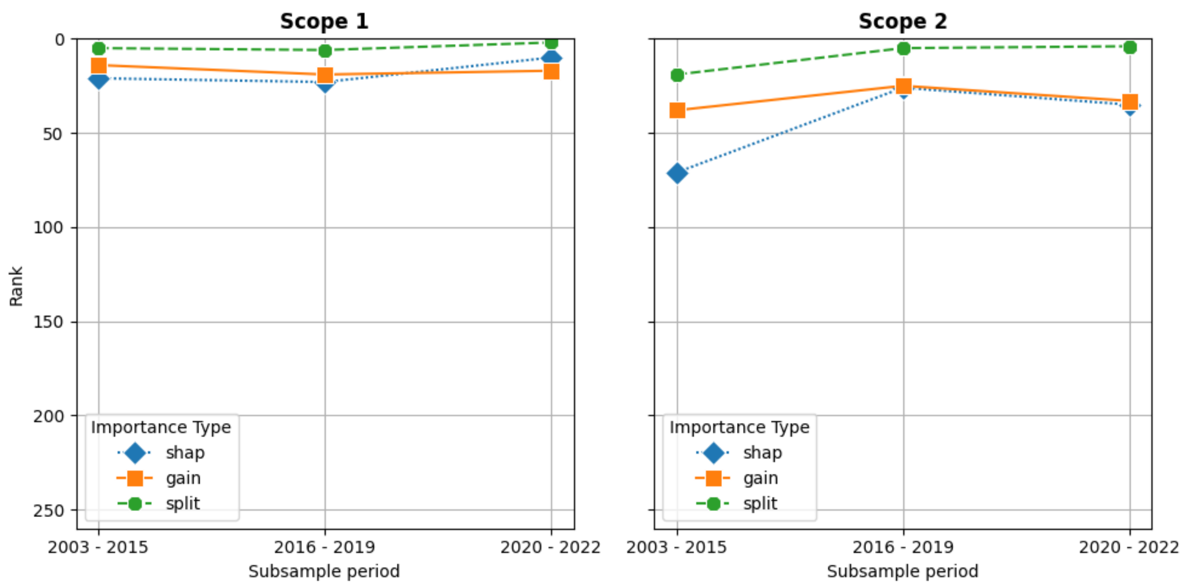

Figure A.3.3 Feature Importance: Reported GHG Values & Trucost Estimates

This figure depicts the time-varying importance of Scope 1 emission intensity (left) and Scope 2 emission intensity (right) in determining log(M2B). The importance is measured by the rank of Scope 1 or 2 emission intensity in each subsample model according to absolute SHAP values, "split" and "gain," where "split" refers to the number of times a variable is used to make a split in decision trees, and "gain" refers to the reduction in training loss resulting from making a new split point. We expand the GHG reporter sample by including stocks with GHG emission values estimated by Trucost.

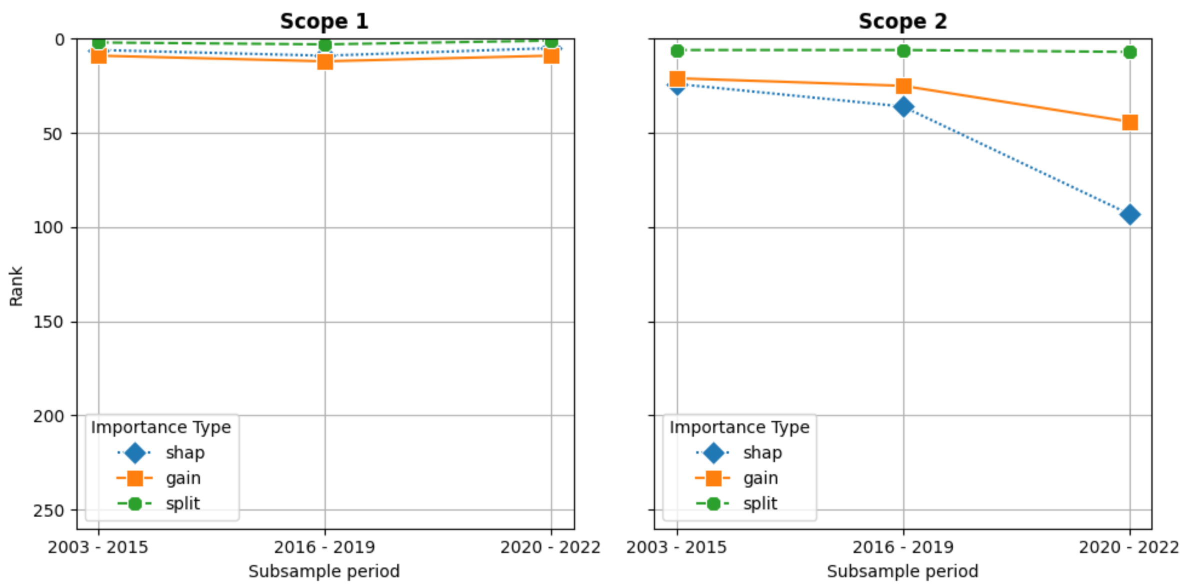

References

Apley, D. W., and J. Zhu. 2020. Visualizing the Effects of Predictor Variables in Black Box Supervised Learning Models. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 82:1059–86.

Aswani, J., A. Raghunandan, and S. Rajgopal. 2024. Are Carbon Emissions Associated with Stock Returns?*. *Review of Finance* 28:75–106.

Bao, Y., B. Ke, B. Li, Y. J. Yu, and J. Zhang. 2020. Detecting Accounting Fraud in Publicly Traded U.S. Firms Using a Machine Learning Approach. *Journal of Accounting Research* 58:199–235.

Barth, M. E., K. Li, and C. G. McClure. 2023. Evolution in Value Relevance of Accounting Information. *The Accounting Review* 98:1–28.

Bartram, S. M., and M. Grinblatt. 2018. Agnostic fundamental analysis works. *Journal of Financial Economics* 128:125–47.

———. 2021. Global market inefficiencies. *Journal of Financial Economics* 139:234–59.

Bhojraj, S., and C. M. C. Lee. 2002. Who Is My Peer? A Valuation-Based Approach to the Selection of Comparable Firms. *Journal of Accounting Research* 40:407–39.

Bianchi, D., M. Büchner, and A. Tamoni. 2021. Bond Risk Premiums with Machine Learning. *The Review of Financial Studies* 34:1046–89.

Bolton, P., and M. Kacperczyk. 2021a. Do investors care about carbon risk? *Journal of Financial Economics* 142:517–49.

———. 2023. Global Pricing of Carbon-Transition Risk. *The Journal of Finance* 78:3677–754.

Bolton, P., and M. T. Kacperczyk. 2021b. Carbon Disclosure and the Cost of Capital.

Chapple, L., P. M. Clarkson, and D. L. Gold. 2013. The Cost of Carbon: Capital Market Effects of the Proposed Emission Trading Scheme (ETS). *Abacus* 49:1–33.

Chen, C., B. Ke, and Q. Zhao. 2024. Enhancing Firm Quality Measurement Using Machine Learning: The Importance of Theory-motivated Data Engineering.

Chen, L., M. Pelger, and J. Zhu. 2024. Deep Learning in Asset Pricing. *Management Science* 70:714–50.

Chen, X., Y. H. t. Cho, Y. Dou, and B. Lev. 2022. Predicting Future Earnings Changes Using Machine Learning and Detailed Financial Data. *Journal of Accounting Research* 60:467–515.

Cong, L. W., G. Feng, J. He, and X. He. 2023. Growing the Efficient Frontier on Panel Trees.

Cong, L. W., T. Liang, and X. Zhang. 2019. Textual Factors: A Scalable, Interpretable, and Data-driven Approach to Analyzing Unstructured Information.

Cong, L. W., K. Tang, J. Wang, and Y. Zhang. 2021. AlphaPortfolio: Direct Construction Through Deep Reinforcement Learning and Interpretable AI.

Connors, E., H. H. Johnston, and L. S. Gao. 2013. The informational value of Toxics Release Inventory performance. *Sustainability Accounting, Management and Policy Journal* 4:32–55.

Cormier, D., and M. Magnan. 1997. Investors' assessment of implicit environmental liabilities: An empirical investigation. *Journal of Accounting and Public Policy* 16:215–41.

DeMiguel, V., J. Gil-Bazo, F. J. Nogales, and A. A. P. Santos. 2023. Machine learning and fund characteristics help to select mutual funds with positive alpha. *Journal of Financial Economics* 150:103737–.

Duan, T., F. W. Li, and Q. Wen. 2023. Is Carbon Risk Priced in the Cross Section of Corporate Bond Returns? *Journal of Financial and Quantitative Analysis* 1–35.

Erel, I., L. H. Stern, C. Tan, and M. S. Weisbach. 2021. Selecting Directors Using Machine Learning. *The Review of Financial Studies* 34:3226–64.

Freyberger, J., A. Neuhierl, and M. Weber. 2020. Dissecting Characteristics Nonparametrically. *The Review of Financial Studies* 33:2326–77.

Garveya, G. T., M. Iyera, and J. Nashb. 2018. Carbon footprint and productivity: Does the "E" in ESG capture efficiency as well as environment? *Journal Of Investment Management* 16:59–69.

Geertsema, P., and H. Lu. 2023. Relative Valuation with Machine Learning. *Journal of Accounting Research* 61:329–76.

Giglio, S., T. Kuchler, J. Stroebel, and X. Zeng. 2023. Biodiversity Risk.

Görgen, M., A. Jacob, M. Nerlinger, R. Riordan, M. Rohleder, and M. Wilkens. 2020. Carbon Risk.

Gu, S., B. Kelly, and D. Xiu. 2020. Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies* 33:2223–73.

———. 2021. Autoencoder asset pricing models. *Journal of Econometrics* 222:429–50.

Hou, K., C. Xue, and L. Zhang. 2020. Replicating Anomalies. *The Review of Financial Studies* 33:2019–133.

Hsu, P.-H., K. Li, and C.-Y. Tsou. 2023. The Pollution Premium. *The Journal of Finance* 78:1343–92.

Hughes, K. E. 2000. The Value Relevance of Nonfinancial Measures of Air Pollution in the Electric Utility Industry. *The Accounting Review* 75:209–28.

In, S. Y., K. Y. Park, and A. Monk. 2019. Is 'Being Green' Rewarded in the Market?: An Empirical Investigation of Carbon Emission Intensity and Stock Returns.

Jensen, T. I., B. Kelly, and L. H. Pedersen. 2023. Is There a Replication Crisis in Finance? *The Journal of Finance* 78:2465–518.

Johnston, D. M., S. E. Sefcik, and N. S. Soderstrom. 2008. The Value Relevance of Greenhouse Gas Emissions Allowances: An Exploratory Study in the Related United States SO2 Market. *European Accounting Review* 17:747–64.

Kaniel, R., Z. Lin, M. Pelger, and S. Van Nieuwerburgh. 2023. Machine-learning the skill of mutual fund managers. *Journal of Financial Economics* 150:94–138.

Ke, G., Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu. 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc.

Kelly, B. T., S. Pruitt, and Y. Su. 2019. Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134:501–24.

Lee, K.-H., B. Min, and K.-H. Yook. 2015. The impacts of carbon (CO2) emissions and environmental research and development (R&D) investment on firm performance. *International Journal of Production Economics* 167:1–11.

Lundberg, S. M., G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. 2020. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence* 2:56–67.

Lundberg, S. M., and S.-I. Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc.

Matsumura, E. M., R. Prakash, and S. C. Vera-Muñoz. 2014. Firm-Value Effects of Carbon Emissions and Carbon Disclosures. *The Accounting Review* 89:695–724.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski. 2021. Responsible investing: The ESG-efficient frontier. *Journal of Financial Economics* 142:572–97.

Rhodes–Kropf, M., D. T. Robinson, and S. Viswanathan. 2005. Valuation waves and merger activity: The empirical evidence. *Journal of Financial Economics* 77:561–603.

Shapley, L. S. 1951. Notes on the N-Person Game — II: The Value of an N-Person Game. Working Paper, RAND Corporation.

US EPA, OAR. 2020. Scope 1 and Scope 2 Inventory Guidance. https://www.epa.gov/climateleadership/scope-1-and-scope-2-inventory-guidance.

Weng, Q., and H. Xu. 2018. A review of China's carbon trading market. *Renewable and Sustainable Energy Reviews* 91:613–9.

Zhang, S. 2024. Carbon Returns Across the Globe.

Zhang, Y.-J., and Y.-M. Wei. 2010. An overview of current research on EU ETS: Evidence from its operating mechanism and economic effect. *Applied Energy* 87:1804–14.