Green vs. Brown: How the Paris Agreement Altered the Operating Performance of U.S. Firms

JINGQI HUAN, BBE

Department of Economics

Submitted in partial fulfillment
of the requirements for the degree of

Master of Business Economics

Faculty of Economics, Brock University
St. Catharines, Ontario

Abstract

Existing research shows that financial markets have increasingly priced climate-related risks, with low-emission ("green") firms outperforming high-emission ("brown") firms during periods of heightened climate concern, particularly following the Paris Agreement. However, financial market valuations are influenced by expectations, sentiment, and noise, and may not accurately reflect firms' real operating adjustments. Whether climate policy signals translate into observable changes in corporate operations therefore remains an open question.This paper examines whether the Paris Agreement led to differential changes in the operating performance of high- and low-emission firms in the United States. Using firm-level data for S&P 500 companies from 2012 to 2019, firms are classified as high- or low-emission based on industry-relative carbon intensity. A difference-in-differences framework is employed to compare operating outcomes before and after the Agreement's entry into force. The empirical results indicate limited evidence of short-run effects on operating profitability. No statistically significant differences are found in EBITDA or return on assets between high- and low-emission firms following the Paris Agreement. In contrast, asset turnover declines modestly for more emission-intensive firms, suggesting weaker asset utilization after the Agreement. Overall, the findings highlight a potential disconnect between climate-related financial market pricing and real operational adjustments, and contribute to bridging the climate finance and corporate operations literatures.

Contents

1 Introduction

Climate change represents one of the most critical economic and policy challenges of the twenty-first century. In response, governments worldwide signed the Paris Agreement in 2016, committing to emission reduction targets tailored to national circumstances. This policy shift has sparked extensive debate in financial economics regarding the valuation of carbon risk. Existing research indicates that capital markets have begun pricing these risks. For instance, Pástor et al. (2022) demonstrate that during periods of heightened climate concern, green" firms (those with lower emissions) outperform their brown" counterparts. Similarly, Hassan et al. (2019) find that political factors and environmental sentiment significantly influence market valuations, as investors incorporate climate-related information into their expectations.

However, financial market reactions do not necessarily mirror operational adjustments at the firm level. Capital market pricing incorporates forward-looking expectations, investor sentiment, and noise, which may diverge from corporate fundamentals. In contrast, real operating metrics—such as earnings, returns, and asset utilization—more directly reflect a firm's internal strategic shifts. As emphasized by Stein (1996), capital market signals may be imperfectly transmitted into managerial decision-making. Consequently, the actual operational impact of climate policies may differ substantially from their immediate effects on market valuations.

The environmental economics literature offers conflicting predictions regarding how environmental regulation affects firm competitiveness. On one hand, Greenstone et al. (2012) find that strict air quality regulations reduce manufacturing productivity, consistent with the view that compliance imposes significant costs. On the other hand, Porter and van der Linde (1995) argue that environmental regulations can stimulate innovation and improve efficiency, potentially enhancing firm performance. Importantly, much of this literature focuses on binding domestic regulations with clear enforcement mechanisms. By contrast, the Paris Agreement operates primarily through expectations about future policy rather than immediate regulatory constraints, making its effects on firm-level operations an open empirical question.

These considerations motivate my research. I investigate whether the Paris Agreement has had differential impacts on the operating performance of high-emission (brown") and low-emission (green") firms in the United States.

To address this question, I use firm-level data for S&P 500 companies from 2012 to 2019 to construct a carbon emission intensity index. Firms are classified as brown" or green" based on the industry-level median to ensure comparability within sectors. I then employ a difference-in-differences framework to compare changes in operating performance—measured by Earnings Before Interest, Taxes, Depreciation, and Amortization (EBITDA), Return on Assets (ROA), and asset turnover before and after the Paris Agreement's entry into force. The empirical specification includes firm and year fixed effects, and robustness checks are conducted using alternative emission thresholds.

The results indicate limited evidence of short-run changes in operating profitability following the Paris Agreement. Across all model specifications, I find no statistically significant differences in EBITDA or ROA between green and brown firms. In contrast, asset turnover exhibits a modest decline among brown firms relative to green firms after the Agreement's implementation. This pattern suggests that climate policy signals may have influenced firms' asset utilization decisions, even in the absence of immediate and widespread changes in profitability.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data and variable construction. Section 4 presents the empirical methodology. Section 5 reports the main results and robustness checks. Section 6 discusses and implications. Section 7 concludes.

2 Literature Review

At the core of this paper is the idea that the Paris Agreement represents a salient climate-policy shock that shifted expectations about the future direction of climate regulation and the low-carbon transition. Even if the Agreement did not immediately translate into a single domestic rule, it plausibly altered the perceived trajectory of transition policy and thereby affected firms through expectations-driven channels. In particular, if investors and creditors reprice climate transition risk, carbon-intensive (brown) firms may face higher required returns and tighter financing conditions than low-emission (green) firms, with potential spillovers into real decisions such as investment, production planning, and cost management. Consistent with this mechanism, Bolton and Kacperczyk (2021) show that carbon emissions exposure is reflected in expected returns, indicating that transition-related risks can be priced by financial markets. Complementary evidence from corporate bond markets suggests that climate regulatory risk affects firms' credit conditions, and Paris-related variation has been used to identify such repricing (Seltzer et al., 2022). These findings motivate treating the Paris Agreement as an economically meaningful shock and examining whether it is associated with differential changes in operating outcomes for brown versus green firms. This section reviews the related literature on (i) climate risk pricing, (ii) regulation and real effects, and (iii) mechanisms linking policy shocks to firm operations.

2.1 Climate Finance and the Pricing of Climate Risks

In recent years, some literature has documented how financial markets price climate-related risks. Pástor et al. (2022) show that green companies outperform brown companies during periods of greater climate concern. In other words, investors reward companies that are more resilient to environmental risks. The research of Hassan et al. (2019) supplements this view, developing indicators based on political and environmental sentiment and analyzing how these factors affect corporate valuations. Their research shows that public discourse on environmental and regulatory risks also has a substantial impact on stock prices. These researchers' findings indicate that financial markets have begun to actively incorporate climate-related signals into asset valuations.

However, these studies also highlight a limitation of climate finance: financial market pricing does not always reflect actual corporate operations. In financial markets, market reactions are driven by expectations, sentiment, and noise. These factors may deviate from actual corporate behavior. Therefore, while these studies have highlighted the importance of climate factors in capital markets, they do not answer whether this importance has simultaneously spilled over into corporate operations. This distinction is particularly important for the present study. If the Paris Agreement primarily moves expectations and discount rates, market prices may react quickly even when operational adjustments are gradual. By examining operating profitability (IHS(EBITDA) and ROA) and asset utilization (asset turnover), the analysis tests whether transition-risk repricing is accompanied by measurable changes in firms' real performance.

2.2 Environmental Regulation and Real Economic Effects

The actual impact of regulation on businesses is not always uniform. The second category of literature examines how environmental regulations affect a firm's productivity, costs, and competitiveness. Greenstone et al. (2012) used a difference-in-differences approach to study the impact of the U.S. Clean Air Act on manufacturing activities. Their research found that in areas designated as not yet compliant by the new regulations, both productivity and economic output declined. This implies that their research indicates that regulation is essentially a cost increase for businesses; stricter regulations increase the compliance burden, thereby reducing firm performance.

Conversely, Porter and van der Linde (1995) propose a different hypothesis: they argued that reasonable environmental regulations can stimulate innovation, prompting companies to develop their own technologies and improve the efficiency of resource utilization. In other words, they believed that environmental regulations do not necessarily reduce a company's competitiveness, but rather have a positive effect in the long run.

These perspectives imply competing predictions for operating outcomes. Under a compliance-cost view, expected tightening of climate policy raises effective production and adjustment costs for carbon-intensive firms, predicting a relative decline in operating profitability after Paris. Under the Porter Hypothesis, however, regulation can induce innovation and process re-optimization, implying that negative short-run effects may attenuate or reverse as firms improve efficiency. Empirically, examining profitability (EBITDA, ROA) alongside efficiency (Asset Turnover) is informative: cost pressure should depress profitability, while demand-side or utilization effects should be reflected more strongly in asset turnover.

In addition, Hsu et al. (2023) provide strong evidence that a firm's emission level is a key determinant of its cost of capital. They find that high-emission firms face a "pollution premium," resulting in significantly higher financing costs. This finding lends theoretical validity to the carbon emission-based classification method used in this paper, as it suggests that high and low emission firms operate under fundamentally different financial constraints. If higher financing costs translate into tighter investment and working-capital constraints, this channel provides a direct link from emissions exposure to operating performance, reinforcing the rationale for comparing brown and green firms around Paris.

2.3 Market Reactions Versus Internal Firm Operations

However, financial market reactions and internal corporate operations are not always perfectly aligned. Stein (1996) points out that capital markets exert distorting pressure on management's decision-making, especially when stock prices are influenced by investor sentiment rather than fundamentals. In order to cater to short-term market expectations, management may allocate resources inefficiently, leading to a discrepancy between market reaction and actual operating performance. Related empirical work also suggests that financial market responses to policy announcements do not always translate into immediate changes in firms' real decisions, particularly when policy signals affect expectations more than current production constraints. These findings support the view that examining firm-level operating outcomes provides a complementary perspective to studies focused solely on asset prices.

These theoretical frameworks explain why strong capital market reactions to climate policy—even with ample documentation—may not translate into measurable changes in corporate operations. Accordingly, this paper contributes by moving beyond market-based measures and testing for differential changes in firm operations around Paris, exploiting pre-period emissions exposure to identify whether the policy shock is associated with economically meaningful changes in profitability and asset utilization.

3 Data

3.1 Data Sources and Sample Construction

The empirical analysis is based on an unbalanced panel of U.S. firms included in the S&P 500 index over the period 2012–2019. Financial statement data, including total assets, net sales, and Earnings Before Interest, Taxes, Depreciation, and Amortization (EBITDA), are obtained from Compustat North America Fundamentals Annual. The sample period includes four years of pre-treatment observations (2012–2015) and four years of post-treatment observations (2016–2019), allowing for a symmetric comparison around the implementation of the Paris Agreement.

The construction of the analysis sample proceeds in several steps. First, I retain firm-year observations for S&P 500 firms with available data on total assets, net sales, and EBITDA. Second, I merge the financial data with the emissions database by firm identifier and fiscal year. For firm-year observations in which both reported and estimated CO_2 emissions are available, priority is given to reported values. When reported emissions are missing but estimated emissions are available, the estimated values are used; observations lacking both measures are excluded from the sample. The final sample consists of 3,500 firm-year observations.

To accommodate firm-year observations with non-positive values of total assets, net sales, or EBITDA, I apply the inverse hyperbolic sine (IHS) transformation, defined as $\operatorname{asinh}(x) = \ln\big(x + \sqrt{x^2+1}\big)$, which behaves similarly to the logarithmic transformation for large values while remaining well-defined at zero (Burbidge et al., 1988).

Finally, to mitigate the influence of extreme outliers, key continuous variables—including EBITDA, total assets, net sales, and carbon intensity—are winsorized at the 1st and 99th percentiles. Due to data availability, the resulting dataset forms an unbalanced annual panel of S&P 500 firms observed from 2012 to 2019 and provides sufficient variation across firms and over time for the subsequent empirical analysis.

Table 1: Summary Statistics: All Firms vs Green vs Brown

	All Firms		Green		Brown		
Variable	Mean	SD	Mean	SD	Mean	SD	p-value
IHS(EBITDA)	14.76	3.73	14.48	4.03	15.04	3.37	< 0.001
ROA (EBITDA/Assets)	0.13	0.09	0.13	0.09	0.14	0.09	0.016
Asset Turnover	0.73	0.66	0.78	0.74	0.68	0.55	< 0.001
In millions:							
Total Assets (USD)	53.41	123.22	46.57	108.92	60.46	136.07	0.001
Net Sales (USD)	20.42	32.95	19.77	33.84	21.09	32.01	0.240
EBITDA (USD)	4.08	6.72	3.24	5.27	4.94	7.84	< 0.001
CO_2 Emissions (tCO_2e)	3.90	11.07	1.85	6.83	6.02	13.86	< 0.001
Carbon Intensity (CO_2/Sales)	0.31	0.98	0.12	0.44	0.50	1.30	< 0.001
Log(Assets)	16.73	1.35	16.59	1.36	16.87	1.33	< 0.001

Notes: Green (Brown) firms are defined using the pre-2016 (2012 to 2015) industry-median carbon intensity (CO_2/Sales). p-values are from Welch two-sample t-tests for differences in means. Financial statement variables and CO_2 emissions are reported in millions. Key continuous variables are winsorized at the 1st and 99th percentiles within year. The sample covers 2012–2019 firm-year observations.

Table 1 reports summary statistics for the full sample and separately for green and brown firms.

3.2 Variables

3.2.1 Outcome Variables

To capture different dimensions of firms' operating performance, I construct three outcome variables measuring operating profitability, operating return on assets, and asset utilization efficiency.

The first outcome variable is operating profitability, measured as the inverse hyperbolic sine (IHS) transformation of Earnings Before Interest, Taxes, Depreciation, and Amortization (EBITDA), as described in Section 3.1. EBITDA captures the scale of operating cash flows generated by the firm while abstracting from financing and tax decisions. In accounting terms, EBITDA is earnings before interest, taxes, depreciation, and amortization, and is commonly used as a proxy for operating performance that is less sensitive to capital structure and depreciation policies than net income. The IHS transformation behaves similarly to the natural logarithm for large values of EBITDA while remaining well-defined at zero, facilitating an interpretation of regression coefficients comparable to approximate percentage changes.

The second outcome variable is operating return on assets (ROA), defined as

$$ROA_{it} = \frac{\text{EBITDA}_{it}}{\text{Total Assets}_{it}},$$

where i indexes firms and t indexes years. Total assets refer to the book value of total assets reported on the balance sheet at fiscal year-end (Compustat annual fundamentals). As a size proxy, total assets also serve as the denominator in both ROA and asset turnover, making these measures comparable across firms of different scale. Using EBITDA rather than net income isolates the performance of the firm's core operations from the effects of capital structure and tax policies.

The third outcome variable is asset turnover, which measures asset utilization efficiency and is defined as

$$AT_{it} = \frac{\text{Net Sales}_{it}}{\text{Total Assets}_{it}}.$$

Net sales (revenues) are measured as annual sales reported on the income statement, while total assets are measured as fiscal-year-end balance sheet total assets (both from Compustat annual fundamentals). Asset turnover therefore captures how much revenue a firm generates per dollar of assets employed. A decline in asset turnover indicates that a firm generates less revenue per unit of assets and may reflect underutilized or stranded assets in response to regulatory or market pressures.

3.2.2 Treatment Variables

The identification strategy is based on a difference-in-differences framework that exploits variation in firms' emission intensity and the timing of the Paris Agreement. The key treatment variables are defined as follows.

First, I construct firm-level CO_2 emissions and carbon intensity. Let CO2_{it} denote the final emissions measure for firm i in year t, defined as reported CO_2 emissions when available and estimated emissions otherwise. Carbon intensity is then given by

$$CI_{it} = \frac{\text{CO2}_{it}}{\text{Net Sales}_{it}}.$$

Scaling emissions by net sales captures emissions intensity relative to the scale of a firm's economic activity and allows for meaningful comparisons across firms of different sizes.

To account for structural differences in emissions across industries, I classify firms into high- and low-emission groups using an industry-adjusted measure. Specifically, for each industry, I compute the median carbon intensity in the pre-treatment period 2012–2015. For each firm, I then calculate its average pre-treatment carbon intensity and compare it to its industry's pre-treatment median. A firm is classified as high-emission if its average pre-period carbon intensity exceeds its industry's pre-period median, and as low-emission otherwise. Formally, the treatment indicator is defined as

$$Treat_i = \begin{cases} 1, & \text{if } \overline{CI}_{i,\text{pre}} > \text{Median}(CI_{\text{industry,pre}}), \\ 0, & \text{otherwise}, \end{cases}$$

where $\overline{CI}_{i,\text{pre}}$ is the firm's average carbon intensity over 2012–2015. This industry-adjusted classification ensures that the treatment group captures firms that are relatively "dirtier" than their direct peers, rather than simply reflecting differences in sector composition. In the full sample, this procedure yields 1,745 firm-year observations classified as low-emission and 1,682 firm-year observations classified as high-emission.

Second, I define the post-treatment period indicator $Post_t$, which equals one for years 2016–2019 and zero for 2012–2015. The year 2016 is chosen as the beginning of the post-treatment period to reflect the adoption and early implementation of the Paris Agreement.

3.2.3 Control Variables

To isolate the effect of the Paris Agreement from other time-varying firm characteristics, I include a set of standard control variables in the regressions. Specifically, firm size ($Size_{it}$) is defined as the natural logarithm of total assets. Firm size is included to account for the fact that larger firms often benefit from economies of scale and possess greater resources to absorb regulatory compliance costs compared to smaller peers.

In addition to these observable controls, all specifications include firm fixed effects to absorb time-invariant heterogeneity across firms and year fixed effects to capture aggregate shocks that affect all firms in a given year.

4 Methodology

This section outlines the empirical strategy used to estimate the impact of the Paris Agreement on the operating performance of high- and low-emission firms. The analysis relies on a difference-in-differences (DID) framework that exploits both cross-sectional variation in firms' emission intensity and the temporal variation around the implementation of the Paris Agreement. I first present the baseline specification, then discuss an event-study extension, and finally outline the key identification assumptions and potential threats.

4.1 Baseline Difference-in-Differences Specification

I estimate the effect of the Paris Agreement using a difference-in-differences (DID) design that compares outcomes for high-emission ("brown") and low-emission ("green") firms before versus after the Agreement. The intuition is straightforward. Suppose there are two groups of firms, brown and green, and two time periods, pre-Paris and post-Paris. Even if brown firms differ from green firms in levels (e.g., industry composition or production technology), DID removes any *time-invariant* differences between the groups by focusing on changes over time. It then subtracts the change for the control group (green firms) from the change for the treated group (brown firms) to isolate the differential shift around the policy shock.

Formally, in the 2×2 case the DID can be written as:

$$\text{DID} = \left(\overline{Y}^{Brown,Post} - \overline{Y}^{Brown,Pre}\right) - \left(\overline{Y}^{Green,Post} - \overline{Y}^{Green,Pre}\right).$$

This "difference of differences" is interpreted as the causal effect of the Paris Agreement on brown firms under the identifying assumption discussed below.

In the panel setting, the baseline regression specification is:

$$Y_{it} = \alpha + \beta(Treat_i \times Post_t) + \delta Treat_i + \theta Post_t + \gamma' X_{it} + \varepsilon_{it}, \tag{1}$$

where i indexes firms and t indexes years. $Treat_i$ equals one if firm i is classified as high-emission based on its *pre-treatment* (2012–2015) industry-adjusted carbon intensity, and zero otherwise. $Post_t$ equals one for the post-Paris period (2016–2019) and zero for the pre-treatment period (2012–2015). The interaction term $Treat_i \times Post_t$ captures the differential post-Paris change for high-emission firms relative to low-emission firms.

I estimate Equation (1) for three outcomes: IHS(EBITDA$_{it}$), ROA_{it}, and AT_{it} (asset turnover). The control vector X_{it} includes time-varying firm characteristics, in particular firm size measured as log(Total Assets$_{it}$).

accTo ount for unobserved heterogeneity, I then estimate the standard two-way fixed effects (TWFE) DiD model:

$$Y_{it} = \alpha + \beta(Treat_i \times Post_t) + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \tag{2}$$

where firm fixed effects μ_i absorb all time-invariant differences across firms (e.g., persistent differences in business models or baseline production technologies), and year fixed effects λ_t capture shocks common to all firms in a given year . Note that in Equation (2) the main effects $Treat_i$ and $Post_t$ are not separately identified: $Treat_i$ is absorbed by firm fixed effects and $Post_t$ is absorbed by year fixed effects. For this reason, I report both the "no fixed effects" specification in Equation (1)—which displays the coefficients on $Treat_i$ and $Post_t$ explicitly—and the preferred TWFE specification in Equation (2).

The key identifying assumption is that, absent the Paris Agreement, high- and low-emission firms would have followed parallel trends in the outcomes:

$$\mathbb{E}\left[Y_{it}(0) - Y_{i,t-1}(0) \mid Treat_i = 1\right] = \mathbb{E}\left[Y_{it}(0) - Y_{i,t-1}(0) \mid Treat_i = 0\right].$$

Intuitively, the green firms provide a counterfactual for how brown firms would have evolved in the absence of the policy shock. Classifying firms using pre-treatment average carbon intensity (2012–2015) helps ensure that treatment status is predetermined with respect to post-2015 outcomes, reducing concerns that treatment assignment is driven by post-period performance.

All regressions are estimated by OLS and standard errors are clustered at the firm level to account for serial correlation in outcomes within firms over time.

4.2 Limitations and Threats to Identification

While the DID framework is well-suited to studying policy shocks with common timing and heterogeneous exposure, its causal interpretation relies on assumptions that may be violated in this setting. I discuss the main threats and how the empirical design addresses them.

The key identifying assumption is that, absent the Paris Agreement, high- and low-emission firms would have followed parallel trends in operating outcomes. This may fail if brown firms are concentrated in industries with systematically different trends (e.g., commodity cycles or energy price sensitivity) that evolve differently over time. Firm fixed effects absorb time-invariant differences in industry composition and business models, and year fixed effects capture aggregate shocks common to all firms. Nevertheless, if shocks are industry-specific and time-varying, they could induce differential trends unrelated to Paris. To mitigate this concern, I classify firms using industry-adjusted pre-period carbon intensity and assess pre-trends using the event-study specification.

A second concern is anticipation: firms may adjust investment and operations before 2016 if Paris-related negotiations and information were incorporated earlier. In such a case, the estimated post-period effect may understate (or mis-time) the true adjustment path. The event-study helps diagnose anticipation by examining whether treated and control firms begin to diverge prior to 2016. More generally, because policy implementation is gradual and heterogeneous across jurisdictions, the estimated effect should be interpreted as the average differential change associated with the post-Paris period rather than a sharp one-time treatment.

Lastly, because the sample is an unbalanced panel, entry and exit may differ across brown and green firms. If Paris-related forces influence survival or reporting in a way that differs by emissions exposure, the estimates may reflect compositional change rather than within-firm operational adjustments. The use of firm fixed effects focuses identification on within-firm changes over time, but selection concerns remain a potential limitation.

Given these limitations, the DID estimates are best interpreted as evidence on whether operating outcomes for high-emission firms changed differentially relative to low-emission firms in the post-Paris period, under the maintained assumption that the event-study provides support for parallel pre-trends and that remaining differential shocks are not driving the results.

4.3 Event-Study

To examine the dynamics of the treatment effect and to provide evidence on the plausibility of the parallel trends assumption, I also estimate an event-study specification. Instead of a single post-period dummy, I interact the treatment indicator with a set of year dummies, excluding one pre-treatment year as the reference category. Specifically, the model is

$$Y_{it} = \alpha + \sum_{k \neq k_0} \beta_k \left(Treat_i \times D_t^k \right) + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \tag{3}$$

where D_t^k is an indicator equal to one if year t corresponds to calendar year k and zero otherwise, and $k_0 = 2015$ denotes the omitted base year, corresponding to the final pre-treatment year prior to the Paris Agreement. This normalization allows the coefficients β_k to be interpreted as deviations from the pre-treatment baseline and facilitates an assessment of the parallel trends assumption.

The pre-treatment coefficients (β_k for $k < 2016$) provide a visual and statistical check of the parallel trends assumption: if high- and low-emission firms followed similar trajectories prior to the Paris Agreement, these coefficients should be close to zero and statistically insignificant. The post-treatment coefficients (β_k for $k \geq 2016$) show whether and when any divergence in operating performance emerges following the Agreement. Because policy anticipation or early adjustment could begin before 2016, the event-study is particularly informative about whether any divergence is already present in the pre-period. In the absence of anticipation effects, coefficients for $k < 2016$ should be statistically indistinguishable from zero.

To visualize the dynamics of the treatment effect and to assess the plausibility of the parallel trends assumption, I plot the estimated coefficients β_k and their 0.95 confidence intervals in Figure. The coefficients are expressed in event time relative to 2016.

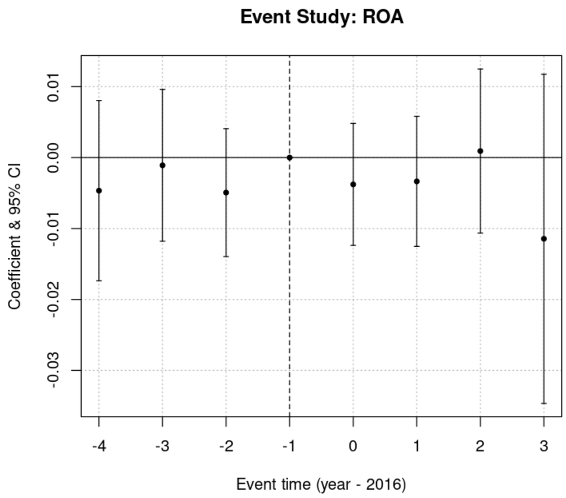

(a) Return on Assets (ROA)

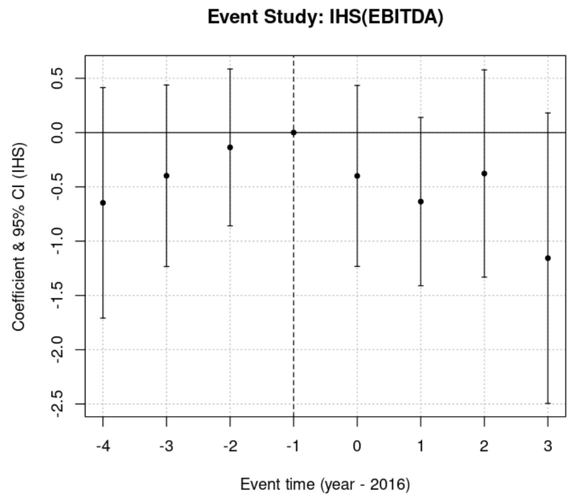

(b) IHS(EBITDA)

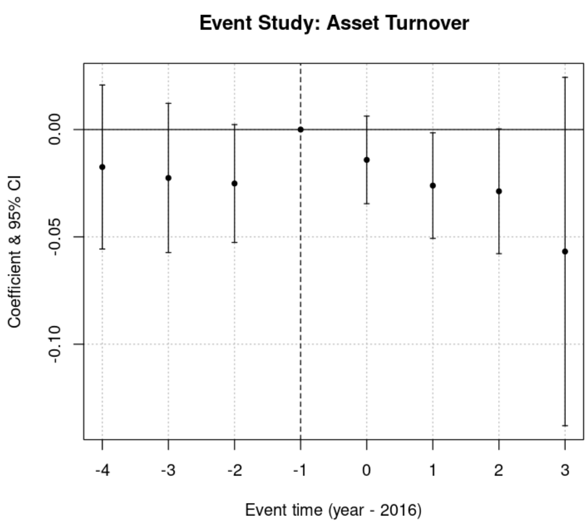

(c) Asset Turnover

Figure 1: Event-study estimates of the differential effect of the Paris Agreement on firm operating performance. The figure plots event-time coefficients from difference-in-differences specifications comparing high- and low-emission firms, with 2015 serving as the omitted base year. Vertical bars denote 95% confidence intervals.

The pre-treatment coefficients for the years 2012–2015 are small in magnitude and statistically indistinguishable from zero across all three outcome variables, which is consistent with the parallel trends assumption. Following 2016, the coefficients for IHS(EBITDA) and ROA remain close to zero, whereas the coefficients for asset turnover become modestly negative in later years, suggesting a mild deterioration in asset utilization for high-emission firms relative to low-emission firms.

5 Results

This section presents the empirical results from the difference-in-differences specifications described above. I first report the baseline estimates using the full sample of firms classified as high- and low-emission based on industry-adjusted carbon intensity. I then discuss the event-study evidence on pre-trends and dynamic effects. Finally, I examine the robustness of the findings using an extreme-group design and including industry-by-year fixed effects.

5.1 Baseline Difference-in-Differences Estimates

Table 2 reports a pooled (no fixed effects) version of the DID model for IHS(EBITDA) that explicitly displays the main effects of treatment status and the post-Paris period. This specification is useful for showing the "level" differences between high- and low-emission firms and how average outcomes shift after 2016, although it does not address time-invariant firm heterogeneity.

Column (1) indicates that high-emission firms have higher average IHS(EBITDA) than low-emission firms in the pre-period (coefficient 0.6025, s.e. 0.2424), and that average operating income rises in the post-2016 period (coefficient 0.3226, s.e. 0.1800). After controlling for firm size in column (2), the level difference between high- and low-emission firms attenuates (coefficient 0.3685, s.e. 0.2129), while the post-period shift also becomes smaller. Firm size, proxied by log(Assets), is strongly positively related to operating income (coefficient

Table 2: DID Results showing Level Effects (IHS(EBITDA))

	(1) No control	(2) Control size
High-emission (Treat)	0.6025*	0.3685+
	(0.2424)	(0.2129)
Post-2016 (Post)	0.3226+	0.2044
	(0.1800)	(0.1745)
DID (Treat × Post)	-0.1026	-0.0997
	(0.2436)	(0.2427)
Log(Assets)		0.8251***
		(0.0733)
Constant	14.3356***	0.6884
	(0.2014)	(1.3325)
Num. Obs.	3392	3392
R^2	0.007	0.093

Notes: Standard errors in parentheses (clustered at the firm level). Significance: + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

0.8251, s.e. 0.0733), as expected.

The DID coefficient of interest is the interaction term, $Treat \times Post$. In both columns (1) and (2), the estimated interaction is negative (around -0.10) but statistically indistinguishable from zero, suggesting no clear evidence that high-emission firms experienced a differential change in operating profitability relative to low-emission firms after 2016 in this pooled specification. Because this pooled model may confound persistent firm differences with the treatment indicator, the preferred estimates incorporate fixed effects to absorb time-invariant firm heterogeneity and common macro shocks.

Tables 5 and 6 report analogous DID estimates for ROA and asset turnover. Columns (1) and (2) present pooled specifications without fixed effects, first with only the main effects and then with the standard interaction term. Column (3) adds year fixed effects, which absorb the common post-period indicator; columns (4) and (5) further include firm fixed effects, which absorb the time-invariant treatment indicator. Accordingly, in the two-way fixed effects specifications the main effects of *High-emission* and *Post* are not separately reported, and identification comes from the interaction term *High-emission* × *Post*.

For ROA (Table 5), the estimated DID coefficients are small and statistically insignificant across all specifications. The interaction term is positive in the pooled and year-fixed-effects models (around 0.003–0.004 with standard errors of 0.005), and becomes essentially zero once firm and year fixed effects and controls are included (column (5)). Economically, these magnitudes imply no meaningful differential change in operating returns on assets for high-emission firms relative to low-emission firms in the post-2016 period.

For asset turnover (Table 6), the pooled estimates show that high-emission firms have lower average asset turnover in levels (the *High-emission* coefficient is about -0.11 and marginally significant), and that average turnover declines after 2016 for all firms (the *Post* coefficient is negative and statistically significant). However, the DID interaction terms are again close to zero and statistically insignificant across specifications, including the preferred two-way fixed effects model with controls (column (5)). Taken together, the turnover results suggest that while levels and aggregate time shifts are present, the Paris-period does not generate a robust *differential* change in asset utilization for high-emission firms relative to low-emission firms.

5.2 Event-Study Evidence and Parallel Trends

On page 11, the figures plots the coefficients from the event-study specification in equation (3), where the high-emission indicator is interacted with a full set of year dummies. The horizontal axis shows event time relative to 2016, while the vertical axis reports the estimated coefficients and their 95% confidence intervals for each of the three outcome variables.

The pre-treatment coefficients for the years 2012–2015 are close to zero and statistically indistinguishable from zero for all three outcomes. This pattern provides supportive evidence for the parallel trends assumption: prior to the Paris Agreement, high- and low-emission firms followed broadly similar trajectories in terms of operating profitability, return on assets, and asset utilization. There is no indication of systematic divergence in trends during the pre-period that could confound the interpretation of the difference-in-differences estimates.

In the post-treatment period, the event-study coefficients remain small for log EBITDA and ROA. The point estimates for these outcomes fluctuate slightly around zero, and the confidence intervals are wide enough to include zero in most years, which is consistent with the baseline finding of no robust differential effect of the Paris Agreement on operating profitability. For asset turnover, the post-2016 coefficients become modestly negative in the later years of the sample, suggesting some deterioration in asset utilization for high-emission firms relative to low-emission firms. However, the estimated effects are not uniformly large across all post-treatment years, and the overall picture remains one of relatively muted operational responses.

Taken together, the event-study evidence reinforces the baseline results. The absence of strong pre-trends supports the validity of the identification strategy, and the dynamic estimates do not reveal substantial or persistent post-Paris deviations in operating performance between high- and low-emission firms.

5.3 Robustness: Extreme Emission Groups and Alternative Cutoffs

To assess whether the baseline results are sensitive to how "high-emission" firms are defined, I re-estimate the DID model using a sequence of increasingly restrictive cutoffs in the pre-treatment (2012–2015) carbon-intensity distribution. For each cutoff $c \in \{50, 40, 30, 20, 10\}$, I redefine the treatment indicator based on whether a firm falls in the high-emission group under cutoff c and estimate the same specification with firm and year fixed effects and a size control.[1] The results are reported in Table 3.

Across outcomes, the estimated interaction terms are not statistically distinguishable from zero and do not display a stable or monotonic pattern as the cutoff becomes more restrictive. In Panel A, the IHS(EBITDA) estimates vary in sign across cutoffs, and confidence intervals widen substantially as the sample shrinks, making it difficult to draw precise conclusions about differential operating-income effects among the most emission-intensive firms.

[1] As the cutoff tightens, the effective sample size declines sharply, which mechanically reduces statistical power and inflates standard errors.

Table 3: Sensitivity Analysis Using Alternative Emission Cutoffs

	50%	40%	30%	20%	10%
Panel A: IHS(EBITDA)					
DID ($\text{High}_i \times \text{Post}_t$)	-0.2338	-0.0329	0.4237	0.3542	0.5065
	(0.2539)	(0.2693)	(0.2716)	(0.3468)	(0.4218)
ln(Assets)	0.9890**	0.7613*	1.0155**	1.2675**	1.5489+
	(0.3188)	(0.3122)	(0.3323)	(0.4649)	(0.8204)
Observations	3,357	2,698	2,040	1,391	735
R^2	0.375	0.358	0.372	0.392	0.309
Panel B: Return on Assets (ROA)					
DID ($\text{High}_i \times \text{Post}_t$)	-0.0001	-0.0004	0.0034	0.0009	0.0019
	(0.0046)	(0.0051)	(0.0062)	(0.0064)	(0.0089)
ln(Assets)	-0.0524***	-0.0541***	-0.0541***	-0.0514***	-0.0481***
	(0.0089)	(0.0096)	(0.0122)	(0.0086)	(0.0120)
Observations	3,357	2,698	2,040	1,391	735
R^2	0.801	0.813	0.808	0.837	0.826
Panel C: Asset Turnover					
DID ($\text{High}_i \times \text{Post}_t$)	-0.0100	-0.0140	-0.0144	-0.0262	-0.0463
	(0.0159)	(0.0177)	(0.0183)	(0.0221)	(0.0317)
ln(Assets)	-0.2587***	-0.2484***	-0.2235***	-0.2273***	-0.2212***
	(0.0275)	(0.0264)	(0.0267)	(0.0277)	(0.0416)
Observations	3,388	2,729	2,066	1,402	743
R^2	0.966	0.967	0.973	0.974	0.978

Notes: This table reports sensitivity analyses of the baseline difference-in-differences estimates using alternative emission cutoffs to define high- and low-emission firms. Columns correspond to cutoffs based on the top 50%, 40%, 30%, 20%, and 10% of firms ranked by pre-treatment emission intensity. All specifications include firm and year fixed effects and control for firm size. Standard errors are clustered at the firm level and reported in parentheses. $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Panel B shows that ROA effects remain tightly centered near zero for all cutoffs, providing little evidence of differential changes in operating returns. In Panel C, the asset-turnover interaction terms are uniformly negative and become somewhat larger in magnitude for tighter cutoffs, but they remain imprecisely estimated and statistically insignificant.

These patterns are informative for mechanism interpretation. If the Paris Agreement primarily increased compliance burdens or operating costs for the most carbon-intensive firms, one would expect the DID estimates to become more negative and more precisely estimated as attention is restricted to the extreme emitters. Instead, the results show no systematic amplification of negative operating effects in the extreme-group designs. Taken together, the sensitivity checks reinforce the baseline conclusion that, within the 2012–2019 window, the Paris-period is not associated with robust differential changes in profitability or operating efficiency for high-emission firms relative to low-emission firms.

5.4 Heterogeneous Exposure Analysis

To examine whether the baseline results depend on how emissions exposure is measured, I implement a four-way set of alternative treatment definitions that combine (i) industry-relative versus aggregate measures and (ii) intensity versus absolute emissions.

To examine whether the baseline estimates depend on how emissions exposure is measured, I implement a 2×2 set of alternative treatment definitions that vary along two dimensions: (i) whether exposure is measured relative to industry peers (industry-relative) or using an economy-wide benchmark (full-sample/aggregate), and (ii) whether exposure is defined by emissions intensity (CO_2/Sales) or by absolute emissions (CO_2 levels). These distinctions are economically meaningful: industry-relative measures capture within-sector positioning where regulatory and technological constraints are often industry-specific, whereas aggregate measures reflect exposure to economy-wide transition policy and investor screening. Likewise, intensity captures carbon efficiency per unit of output, while absolute emissions capture scale-related exposure to total emissions constraints. Table 4 reports the results.

Table 4: Robustness Checks across Alternative Emission and Treatment Definitions

	Industry-relative intensity	Aggregate intensity	Industry-relative emissions	Aggregate emissions
Panel A: IHS(EBITDA)				
High-emission × Post-Paris	-0.2311	-0.4601^{+}	-0.2878	-0.3425
	(0.2511)	(0.2529)	(0.2501)	(0.2424)
ln(Assets)	0.9865**	0.9543**	0.9644**	0.9375**
	(0.3177)	(0.3183)	(0.3169)	(0.3110)
Observations	3,392	3,392	3,392	3,392
R^2	0.375	0.376	0.375	0.375
Panel B: Return on Assets (ROA)				
High-emission × Post-Paris	-0.0001	-0.0105*	-0.0068	-0.0131**
	(0.0046)	(0.0044)	(0.0044)	(0.0043)
ln(Assets)	-0.0524***	-0.0535***	-0.0533***	-0.0548***
	(0.0089)	(0.0088)	(0.0088)	(0.0088)
Observations	3,392	3,392	3,392	3,392
R^2	0.802	0.803	0.802	0.803
Panel C: Asset Turnover				
High-emission × Post-Paris	-0.0090	-0.0697***	-0.0633***	-0.1008***
	(0.0157)	(0.0155)	(0.0159)	(0.0154)
ln(Assets)	-0.2583***	-0.2650***	-0.2657***	-0.2764***
	(0.0274)	(0.0278)	(0.0279)	(0.0278)
Observations	3,427	3,427	3,427	3,427
R^2	0.966	0.967	0.967	0.968

Notes: This table reports robustness checks using alternative definitions of emission exposure. Columns differ by whether emissions are measured using industry-relative versus aggregate metrics, and whether emission intensity or absolute emission levels are used. All specifications include firm and year fixed effects and control for firm size. Standard errors are clustered at the firm level and reported in parentheses. $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

For operating income measured by IHS(EBITDA), the estimated coefficients on the interaction term between high-emission status and the post-Paris period remain negative across all four definitions, but are generally imprecisely estimated. Only the specification based on aggregate emission intensity yields a marginally significant coefficient at the 10 percent level, while the remaining estimates are statistically indistinguishable from zero. Overall, these results indicate that the baseline finding of no robust differential effect on operating income is largely insensitive to how emission exposure is measured.

A somewhat different picture emerges for operating profitability as measured by ROA. The estimated DID effects are close to zero under the industry-relative definitions, but become more negative under the full-sample (aggregate) exposure measures. In particular, the interaction term is statistically significant when exposure is defined using aggregate intensity and aggregate emissions, indicating that high-emission firms experience a modest decline in accounting profitability relative to low-emission firms in the post-Paris period under these broader benchmarks. At the same time, the magnitude of these ROA effects remains economically small, and the sensitivity to the exposure definition suggests that profitability responses are present but not uniformly robust across all measurement choices.

In contrast, the asset-turnover results exhibit a more systematic and robust pattern. When exposure is defined using aggregate measures—both intensity-based and absolute—the estimated treatment effects are negative and statistically significant at conventional levels. Moreover, the coefficients become larger in magnitude as the definition shifts from industry-relative to aggregate benchmarks, consistent with the interpretation that firms with larger overall emissions footprints experienced a more pronounced decline in asset utilization following the Paris Agreement. Relative to the sample mean of asset turnover, the implied reductions are economically meaningful, pointing to an adjustment that operates through utilization or scaling of activity rather than purely through accounting reclassification.

Taken together, the heterogeneous exposure analysis suggests limited and definition-dependent evidence of profitability effects (EBITDA and ROA), but clearer evidence of efficiency adjustments concentrated in asset utilization. The fact that the strongest and most consistent effects appear for asset turnover—especially under aggregate exposure definitions—aligns with an expectations-driven transition channel in which carbon-intensive firms respond by curbing utilization, postponing expansion, or reallocating capital, even when short-run operating profits do not display a uniformly large decline. In this sense, post-Paris adjustment may manifest first in how intensively firms deploy their asset base, with profitability effects remaining modest or harder to detect within the 2012–2019 window.

6 Discussion

A natural question raised by the empirical results is why the estimated treatment effects are generally small and statistically insignificant. One plausible explanation relates to the institutional nature of the Paris Agreement in the U.S. context. The Agreement primarily operates as an international commitment that requires subsequent domestic implementation to become a binding constraint on firms. Over our sample period, however, there was no immediate introduction of a nationwide carbon price or a uniform federal emissions cap that would have sharply tightened regulatory pressure on U.S. corporations. Moreover, the political uncertainty surrounding U.S. climate policy with the country joining the Paris Agreement in late 2016 and a new administration soon announcing its intention to withdraw — may have further reduced firms' incentives to undertake costly, irreversible operational adjustments in response to the Agreement.

A second possible reason is that operating performance is inherently slow-moving. Large S&P 500 firms typically make investment and production decisions over long horizons, and changes in capital stock, production technologies, or supply chains take years to materialize. Relatedly, heightened policy uncertainty may further delay firms' real adjustments, leading to muted short-run effects on profitability even when longer-run impacts are expected (Gulen and Ion, 2016). Even if firms began to internalize climate considerations after the Paris

Agreement, the short post-treatment window from 2016 to 2019 may not be sufficient for such adjustments to translate into noticeable differences in EBITDA or ROA. This interpretation is consistent with the event-study evidence, which shows limited divergence in profitability measures but some deterioration in asset turnover for the most emission-intensive firms, suggestive of early signs of underutilized or stranded assets rather than an immediate collapse in earnings.

Finally, the weak real effects documented here contrast with the strong pricing reactions found in the climate finance literature. One possibility is that capital markets had already priced in climate risks before the Paris Agreement, as suggested by Pástor et al. (2022) and Hassan et al. (2019). In that case, the Agreement may have had a larger impact on expectations and valuations than on short-run operating outcomes, reinforcing the idea that market reactions and firm-level operations need not move in lockstep.

7 Conclusion

This paper examines the impact of the 2016 Paris Agreement on the operating performance of U.S. firms. Using a difference-in-differences approach, I selected S&P 500 companies from 2012 to 2019 as a sample to examine whether high-emission firms experienced structural changes in profitability and efficiency relative to their low-emission peers.

The empirical results yield two distinct sets of conclusions. First, in the complete sample, I found no statistical evidence of a significant negative impact of the Paris Agreement on profitability (EBITDA), asset turnover, or return on assets (ROA). This suggests that in the short term, due to the non-binding nature of the agreement and the lack of an immediately effective federal carbon pricing mechanism, high-emission firms have been able to maintain their financial profit margins, possibly through cost pass-through or a lack of rigorous enforcement.

This study contributes to the literature by highlighting the disconnect between financial market pricing and actual economic performance. While previous research Pástor et al. (2022) suggests that capital markets can reflect climate risks in an immediate manner, my results indicate that actual operational responses are more sticky and complex.

In summary, this paper argues that, unlike financial markets which have begun to incorporate environmental factors into pricing, a decoupling may exist between a firm's actual operating performance and its financial performance.

For policymakers, these results suggest that international commitments alone, without clear and stable domestic enforcement mechanisms, may be insufficient to drive a proactive green transition. Future research could extend this analysis by examining longer time spans or specific asset-level data to further reveal the dynamics of asset stranding in transition economies.

Limitations and issues of this paper are as follows: First, the sample is limited to large firms, neglecting the impact on small and medium-sized enterprises. Finally, I do not provide evidence that the decline in asset turnover is due to policy uncertainty.

Regarding data processing, I do not have a complete dataset; estimated data were used when emissions data were missing, which may introduce measurement errors.

References

Patrick Bolton and Marcin Kacperczyk. Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549, 2021. doi: 10.1016/j.jfineco.2021.05.008.

John B. Burbidge, Lonnie Magee, and A. Leslie Robb. Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association*, 83(401):123–127, 1988.

Michael Greenstone, John A List, and Chad Syverson. The effects of environmental regulation on the competitiveness of U.S. manufacturing. Working Paper 18392, National Bureau of Economic Research, 2012.

Huseyin Gulen and Mihai Ion. Policy uncertainty and corporate investment. *The Review of Financial Studies*, 29(3):523–564, 2016.

Tarek A Hassan, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. Firm-level political risk: Measurement and effects. *The Quarterly Journal of Economics*, 134(4): 2135–2202, 2019.

Po-Hsuan Hsu, Kai Li, and Chi-Yang Tsou. The pollution premium. *The Journal of Finance*, 78(3):1343–1392, 2023.

Ľuboš Pástor, Robert F Stambaugh, and Lucian Taylor. Dissecting green returns. *Journal of Financial Economics*, 146(2):403–424, 2022.

Michael E Porter and Claas van der Linde. Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives*, 9(4):97–118, 1995.

Lee H. Seltzer, Laura T. Starks, and Qifei Zhu. Climate regulatory risk and corporate bonds. NBER Working Paper 29994, National Bureau of Economic Research, 2022. NY Fed Staff Report No. 1014; revised March 2025.

Jeremy C Stein. Rational capital budgeting in an irrational world. *Journal of Business*, 69 (4):429–455, 1996.

A Additional Difference-in-Differences Results

Table 5: Baseline Difference-in-Differences Results (ROA)

	(1)	(2)	(3)	(4)	(5)
High-emission	0.007	0.006	0.006		
	(0.007)	(0.007)	(0.007)		
Post	-0.000	-0.002			
	(0.003)	(0.003)			
High-emission × Post		0.004	0.004	0.003	-0.000
		(0.005)	(0.005)	(0.005)	(0.005)
Num. Obs.	3,392	3,392	3,392	3,392	3,392
R^2	0.002	0.002	0.005	0.783	0.802
Firm FE	No	No	No	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Controls	No	No	No	No	Yes

Standard errors in parentheses.

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 6: Baseline Difference-in-Differences Results (Asset Turnover)

	(1)	(2)	(3)	(4)	(5)
High-emission	-0.106^{+}	-0.107^{+}	-0.107^{+}		
	(0.058)	(0.061)	(0.061)		
Post	-0.038***	-0.038*			
	(0.011)	(0.018)			
High-emission × Post		0.001	0.005	0.003	-0.009
		(0.022)	(0.022)	(0.017)	(0.016)
Num. Obs.	3,427	3,427	3,427	3,427	3,427
R^2	0.007	0.007	0.009	0.958	0.966
Firm FE	No	No	No	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Controls	No	No	No	No	Yes

Standard errors in parentheses.

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 7: Baseline Difference-in-Differences Results with Firm and Year Fixed Effects

	IHS(EBITDA)	ROA	Asset Turnover
(high emission$_i$ × Post$_t$)	-0.2311	-0.0001	-0.0090
	(0.2511)	(0.0046)	(0.0157)
ln(Assets)	0.9865**	-0.0524***	-0.2583***
	(0.3177)	(0.0089)	(0.0274)
Observations	3,392	3,392	3,427
R^2	0.375	0.802	0.966
Adjusted R^2	0.268	0.768	0.961
Within R^2	0.006	0.086	0.197
Within Adj. R^2	0.006	0.085	0.197
AIC	17,962.3	-11,308.3	-3,809.6
BIC	21,020.8	-8,249.9	-733.8
RMSE	2.95	0.04	0.12
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Std. Errors	Clustered by firm	Clustered by firm	Clustered by firm

Standard errors in parentheses.

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001.

B Code

```
2 sheets used
# ==============================================================================
# Step 0: (first)
# ==============================================================================
library(tidyverse)
library(janitor)
library(readr)
library(fixest)
library(modelsummary)
library(ggplot2)
# ==============================================================================

winsor_function <- function(x, probs = c(0.01, 0.99)) {
  if (all(is.na(x))) return(x)

  q <- quantile(x, probs, na.rm = TRUE)

  x[x < q[1]] <- q[1]

  x[x > q[2]] <- q[2]

  return(x)
}

# ==============================================================================
# Step 1:
# ==============================================================================
raw_data <- read_csv("data_2012_2019.csv")
df_clean <- raw_data %>%
  janitor::clean_names()

# ==============================================================================
# Step 2:
# ==============================================================================
df_numeric <- df_clean %>%
  mutate(
    # EBITDA
    ebitda_raw = parse_number(as.character(ebit_depreciation)),

    #
    assets     = parse_number(as.character(total_assets)),

    #
    sales      = parse_number(as.character(net_sales_or_revenues)),

    #
    co2_real   = parse_number(as.character(co2_equivalents_emission_total)),
    co2_est    = parse_number(as.character(estimated_co2_equivalents_emission_total)),
    co2_final  = coalesce(co2_real, co2_est),

    #
    year       = as.numeric(year)
  ) %>%

  filter(assets > 0, sales > 0)
```

```
# ================================================================================
# Step 2.5: Winsorization
# ================================================================================
winsor_vars <- c("ebitda_raw", "assets", "sales", "co2_final")

df_winsor <- df_numeric %>%
  group_by(year) %>%
  mutate(across(all_of(winsor_vars), ~ winsor_function(.))) %>%
  ungroup()

# ================================================================================
# Step 3:  (Outcome Variables)
# ================================================================================
df_outcomes <- df_winsor %>%
  mutate(
    # 1. ROA
    roa = ebitda_raw / assets,

    # 2. Asset Turnover
    asset_turnover = sales / assets,

    # 3. IHS(EBITDA) -  Log
    # asinh(x) = ln(x + sqrt(x^2 + 1))

    ihs_ebitda = asinh(ebitda_raw)
  )

# ================================================================================
# Step 4: (Treat)
# ================================================================================

# 4.1
df_intensity <- df_outcomes %>%
  mutate(carbon_intensity = co2_final / sales) %>%
  filter(!is.na(carbon_intensity))

# 4.2  Pre-period (2012-2015)
firm_pre_avg <- df_intensity %>%
  filter(year >= 2012, year <= 2015) %>%
  group_by(name, icbic) %>%
  summarize(
    firm_avg_intensity = mean(carbon_intensity, na.rm = TRUE),
    .groups = "drop"
  )

# 4.3
industry_benchmarks <- firm_pre_avg %>%
  group_by(icbic) %>%
  summarize(
    industry_median = median(firm_avg_intensity, na.rm = TRUE),
    .groups = "drop"
  )

# 4.4  Treat (1=high, 0=low)
firm_treat_status <- firm_pre_avg %>%
  left_join(industry_benchmarks, by = "icbic") %>%
  mutate(
    treat = if_else(firm_avg_intensity > industry_median, 1, 0)
```

```
  ) %>%
  select(name, icbic, treat)

# ==============================================================================
# Step 5:
# ==============================================================================
df_final <- df_intensity %>%
  inner_join(firm_treat_status, by = c("name", "icbic")) %>%
  mutate(
    post = if_else(year >= 2016, 1, 0),
    did  = treat * post,

    ln_assets = log(assets),

    event_time = year - 2016
  )

# ==============================================================================
# Step 7: DID
# ==============================================================================

# 1. IHS(EBITDA)
did_ihs_ebitda <- feols(
  ihs_ebitda ~ did + ln_assets | name + year,
  data = df_final,
  vcov = ~name
)

# 2. ROA
did_roa_nq <- feols(
  roa ~ did + ln_assets | name + year,
  data = df_final,
  vcov = ~name
)

# 3. Asset Turnover
did_turnover_nq <- feols(
  asset_turnover ~ did + ln_assets | name + year,
  data = df_final,
  vcov = ~name
)

# ==============================================================================
# Step 8:
# ==============================================================================

modelsummary(
  list(
    "IHS(EBITDA)"    = did_ihs_ebitda,
    "ROA"            = did_roa_nq,
    "Asset Turnover" = did_turnover_nq
  ),
  stars = TRUE,
  fmt = 4,
  statistic = "std.error",
  title = "DID Results (Using IHS for EBITDA)"
)

# ==============================================================================
```

```
# Step 9: Event Study ROA
# ==============================================================================

es_roa_nq <- feols(
  roa ~ i(event_time, treat, ref = -1) + ln_assets | name + year,
  data = df_final,
  vcov = ~name
)

iplot(
  es_roa_nq,
  main = "Event Study: ROA",
  xlab = "Event time (year - 2016)",
  ylab = "Coefficient & 95% CI"
)
# ------------------------------------------------------------------------------
# 2. Event Study: IHS(EBITDA)
# ------------------------------------------------------------------------------
es_ihs_ebitda <- feols(
  ihs_ebitda ~ i(event_time, treat, ref = -1) + ln_assets | name + year,
  data = df_final,
  vcov = ~name
)
# IHS(EBITDA)
iplot(
  es_ihs_ebitda,
  main = "Event Study: IHS(EBITDA)",
  xlab = "Event time (year - 2016)",
  ylab = "Coefficient & 95% CI (IHS)"
)
# 3. Event Study: Asset Turnover
# ------------------------------------------------------------------------------
es_turnover <- feols(
  asset_turnover ~ i(event_time, treat, ref = -1) + ln_assets | name + year,
  data = df_final,
  vcov = ~name
)
# Asset Turnover
iplot(
  es_turnover,
  main = "Event Study: Asset Turnover",
  xlab = "Event time (year - 2016)",
  ylab = "Coefficient & 95% CI"
)

# ==============================================================================
#  Sensitivity Check: All Outcomes
# ==============================================================================

#  (Top/Bottom X%)
cutoffs <- c(0.5, 0.4, 0.3, 0.2, 0.1)

#
models_ebitda   <- list()
models_roa      <- list()
models_turnover <- list()

#
for (k in cutoffs) {
```

```
# --- 1.
p_lower <- k
p_upper <- 1 - k

#
industry_thresholds <- firm_pre_avg %>%
  group_by(icbic) %>%
  summarize(
    th_lower = quantile(firm_avg_intensity, p_lower, na.rm = TRUE),
    th_upper = quantile(firm_avg_intensity, p_upper, na.rm = TRUE),
    .groups = "drop"
  )

#
current_treat_status <- firm_pre_avg %>%
  left_join(industry_thresholds, by = "icbic") %>%
  mutate(
    treat = case_when(
      firm_avg_intensity > th_upper ~ 1,
      firm_avg_intensity < th_lower ~ 0,
      TRUE ~ NA_real_
    )
  ) %>%
  filter(!is.na(treat)) %>%
  select(name, icbic, treat)

#
df_loop <- df_intensity %>%
  inner_join(current_treat_status, by = c("name", "icbic")) %>%
  mutate(
    post = if_else(year >= 2016, 1, 0),
    did  = treat * post,
    ln_assets = log(assets),
    ihs_ebitda = asinh(ebitda_raw)
  )

# --- 2.

#
col_name <- paste0(k*100, "%")

# Model A: IHS(EBITDA)
models_ebitda[[col_name]] <- feols(
  ihs_ebitda ~ did + ln_assets | name + year,
  data = df_loop,
  vcov = ~name
)

# Model B: ROA
models_roa[[col_name]] <- feols(
  roa ~ did + ln_assets | name + year,
  data = df_loop,
  vcov = ~name
)

# Model C: Asset Turnover
models_turnover[[col_name]] <- feols(
  asset_turnover ~ did + ln_assets | name + year,
```

```
    data = df_loop,
    vcov = ~name
  )
}

# ==============================================================================
#
# ==============================================================================

#  1: IHS(EBITDA)
modelsummary(
  models_ebitda,
  stars = TRUE,
  fmt = 4,
  statistic = "std.error",
  title = "Panel A: Sensitivity Analysis for IHS(EBITDA)",
  gof_map = c("nobs", "r.squared")
)

#  2: ROA
modelsummary(
  models_roa,
  stars = TRUE,
  fmt = 4,
  statistic = "std.error",
  title = "Panel B: Sensitivity Analysis for ROA",
  gof_map = c("nobs", "r.squared")
)

#  3: Asset Turnover
modelsummary(
  models_turnover,
  stars = TRUE,
  fmt = 4,
  statistic = "std.error",
  title = "Panel C: Sensitivity Analysis for Asset Turnover",
  gof_map = c("nobs", "r.squared")
)

# ==============================================================================
# Step 0: Second
# ==============================================================================
library(tidyverse)
library(janitor)
library(readr)
library(fixest)
library(modelsummary)
library(ggplot2)
library(dplyr)
library(tidyr)
library(purrr)
library(knitr)
winsor_function <- function(x, probs = c(0.01, 0.99)) {
  if (all(is.na(x))) return(x)
  q <- quantile(x, probs, na.rm = TRUE)
  x[x < q[1]] <- q[1]
  x[x > q[2]] <- q[2]

  return(x)
```

```
}

# ==============================================================================
# Step 1: read
# ==============================================================================
raw_data <- read_csv("data_2012_2019.csv")
df_clean <- raw_data %>%
  janitor::clean_names()

# ==============================================================================
# Step 2: clean
# ==============================================================================
df_numeric <- df_clean %>%
  mutate(
    # EBITDA
    ebitda_raw = parse_number(as.character(ebit_depreciation)),


    assets     = parse_number(as.character(total_assets)),


    sales      = parse_number(as.character(net_sales_or_revenues)),


    co2_real   = parse_number(as.character(co2_equivalents_emission_total)),
    co2_est    = parse_number(as.character(estimated_co2_equivalents_emission_total)),
    co2_final  = coalesce(co2_real, co2_est),

    year       = as.numeric(year)
  ) %>%
  filter(assets > 0, sales > 0)

# ==============================================================================
# Step 2.5:  (Winsorization)
# ==============================================================================
# 1% to 99%
winsor_vars <- c("ebitda_raw", "assets", "sales", "co2_final")

df_winsor <- df_numeric %>%
  group_by(year) %>%
  mutate(across(all_of(winsor_vars), ~ winsor_function(.))) %>%
  ungroup()

# ==============================================================================
# Step 3: Outcome Variables
# ==============================================================================
df_outcomes <- df_winsor %>%
  mutate(
    # 1. ROA
    roa = ebitda_raw / assets,

    # 2. Asset Turnover
    asset_turnover = sales / assets,

    # 3. IHS(EBITDA) -  Log
    # asinh(x) = ln(x + sqrt(x^2 + 1))
    ihs_ebitda = asinh(ebitda_raw)
  )
# ==============================================================================
```

```
# Step 4:  2x2  Treatment variable
# ============================================================================

# 4.1
df_metrics <- df_outcomes %>%
  mutate(
    #  (Intensity)
    metric_intensity = co2_final / sales,
    #  (Absolute)
    metric_absolute  = co2_final
  ) %>%
  filter(!is.na(metric_intensity), !is.na(metric_absolute))

# 4.2  Pre-period (2012-2015)
firm_pre_avg <- df_metrics %>%
  filter(year >= 2012, year <= 2015) %>%
  group_by(name, icbic) %>%
  summarize(
    avg_int = mean(metric_intensity, na.rm = TRUE),
    avg_abs = mean(metric_absolute, na.rm = TRUE),
    .groups = "drop"
  )

# 4.3  4  Treat variable
# ----------------------------------------------------------------------------

#  (Full Sample)
median_full_int <- median(firm_pre_avg$avg_int, na.rm = TRUE)
median_full_abs <- median(firm_pre_avg$avg_abs, na.rm = TRUE)

#  (Within Industry)
industry_benchmarks <- firm_pre_avg %>%
  group_by(icbic) %>%
  summarize(
    median_ind_int = median(avg_int, na.rm = TRUE),
    median_ind_abs = median(avg_abs, na.rm = TRUE),
    .groups = "drop"
  )

#  4  Treat
firm_treat_4dim <- firm_pre_avg %>%
  left_join(industry_benchmarks, by = "icbic") %>%
  mutate(
    # 1.  (Industry-Adjusted Intensity)
    treat_ind_int  = if_else(avg_int > median_ind_int, 1, 0),

    # 2.  (Unadjusted Intensity)
    treat_full_int = if_else(avg_int > median_full_int, 1, 0),

    # 3.  (Industry-Adjusted Absolute Emissions)
    treat_ind_abs  = if_else(avg_abs > median_ind_abs, 1, 0),

    # 4. (Unadjusted Absolute Emissions)
    treat_full_abs = if_else(avg_abs > median_full_abs, 1, 0)
  ) %>%
  select(name, icbic, starts_with("treat_"))

# ============================================================================
# Step 5:
```

```
# ==============================================================================
df_final_4dim <- df_metrics %>%
  inner_join(firm_treat_4dim, by = c("name", "icbic")) %>%
  mutate(
    post = if_else(year >= 2016, 1, 0),
    ln_assets = log(assets),

    # 4  DID
    did_ind_int  = treat_ind_int * post,
    did_full_int = treat_full_int * post,
    did_ind_abs  = treat_ind_abs * post,
    did_full_abs = treat_full_abs * post,

    # IHS
    ihs_ebitda = asinh(ebitda_raw)
  )

# ==============================================================================
# Step 7: variables x des
# ==============================================================================


# --- IHS(EBITDA)---
m_ebitda_1 <- feols(ihs_ebitda ~ did_ind_int + ln_assets | name + year, data = df_final_4dim, vc
m_ebitda_2 <- feols(ihs_ebitda ~ did_full_int + ln_assets | name + year, data = df_final_4dim, v
m_ebitda_3 <- feols(ihs_ebitda ~ did_ind_abs + ln_assets | name + year, data = df_final_4dim, vc
m_ebitda_4 <- feols(ihs_ebitda ~ did_full_abs + ln_assets | name + year, data = df_final_4dim, v

# --- ROA  ---
m_roa_1 <- feols(roa ~ did_ind_int + ln_assets | name + year, data = df_final_4dim, vcov = ~name
m_roa_2 <- feols(roa ~ did_full_int + ln_assets | name + year, data = df_final_4dim, vcov = ~nam
m_roa_3 <- feols(roa ~ did_ind_abs + ln_assets | name + year, data = df_final_4dim, vcov = ~name
m_roa_4 <- feols(roa ~ did_full_abs + ln_assets | name + year, data = df_final_4dim, vcov = ~nam

# ---  Asset Turnover  ---
m_turnover_1 <- feols(asset_turnover ~ did_ind_int + ln_assets | name + year, data = df_final_4d
m_turnover_2 <- feols(asset_turnover ~ did_full_int + ln_assets | name + year, data = df_final_4
m_turnover_3 <- feols(asset_turnover ~ did_ind_abs + ln_assets | name + year, data = df_final_4d
m_turnover_4 <- feols(asset_turnover ~ did_full_abs + ln_assets | name + year, data = df_final_4

# ==============================================================================
# Step 8: three tables
# ==============================================================================

#
coef_map_custom <- c(
  "did_ind_int"  = "DID",
  "did_full_int" = "DID",
  "did_ind_abs"  = "DID",
  "did_full_abs" = "DID",
  "ln_assets"    = "Log(Assets)"
)

# ---  1: IHS(EBITDA) ---
modelsummary(
  list(
    "Ind x Int"  = m_ebitda_1,
    "Full x Int" = m_ebitda_2,
    "Ind x Abs"  = m_ebitda_3,
```

```
    "Full x Abs" = m_ebitda_4
  ),
  coef_map = coef_map_custom,
  stars = TRUE,
  fmt = 4,
  statistic = "std.error",
  title = "Panel A: Robustness Check for IHS(EBITDA) - 4 Dimensions",
  gof_map = c("nobs", "r.squared")
)

# ---  2: ROA ---
modelsummary(
  list(
    "Ind x Int"  = m_roa_1,
    "Full x Int" = m_roa_2,
    "Ind x Abs"  = m_roa_3,
    "Full x Abs" = m_roa_4
  ),
  coef_map = coef_map_custom,
  stars = TRUE,
  fmt = 4,
  statistic = "std.error",
  title = "Panel B: Robustness Check for ROA - 4 Dimensions",
  gof_map = c("nobs", "r.squared")
)

# ---  3: Asset Turnover ---
modelsummary(
  list(
    "Ind x Int"  = m_turnover_1,
    "Full x Int" = m_turnover_2,
    "Ind x Abs"  = m_turnover_3,
    "Full x Abs" = m_turnover_4
  ),
  coef_map = coef_map_custom,
  stars = TRUE,
  fmt = 4,
  statistic = "std.error",
  title = "Panel C: Robustness Check for Asset Turnover - 4 Dimensions",
  gof_map = c("nobs", "r.squared")
)

# EBITDA

# (1) Treat + Post
m1 <- feols(
  ihs_ebitda ~ treat_ind_int + post,
  data = df_final_4dim,
  vcov = ~name
)

# (2) Treat + Post + Treat×Post
m2 <- feols(
  ihs_ebitda ~ treat_ind_int * post,
  data = df_final_4dim,
  vcov = ~name
)

# (3) + Year FE
```

```
m3 <- feols(
  ihs_ebitda ~ treat_ind_int * post | year,
  data = df_final_4dim,
  vcov = ~name
)

# (4) + Firm FE + Year FE
m4 <- feols(
  ihs_ebitda ~ treat_ind_int * post | name + year,
  data = df_final_4dim,
  vcov = ~name
)

# (5) + Controls + Firm FE + Year FE
m5 <- feols(
  ihs_ebitda ~ treat_ind_int * post + ln_assets | name + year,
  data = df_final_4dim,
  vcov = ~name
)
modelsummary(
  list(
    "(1)" = m1,
    "(2)" = m2,
    "(3)" = m3,
    "(4)" = m4,
    "(5)" = m5
  ),
  coef_map = c(
    "treat_ind_int"      = "High-emission",
    "post"               = "Post",
    "treat_ind_int:post" = "High-emission × Post"
  ),
  gof_map = c("nobs", "r.squared"),
  stars = TRUE,
  statistic = "std.error",
  add_rows = tibble(
    term = c("Firm FE", "Year FE", "Controls"),
    `(1)` = c("No",  "No",  "No"),
    `(2)` = c("No",  "No",  "No"),
    `(3)` = c("No",  "Yes", "No"),
    `(4)` = c("Yes", "Yes", "No"),
    `(5)` = c("Yes", "Yes", "Yes")
  ),
  title = "Table 1: Baseline Difference-in-Differences Results (IHS(EBITDA))"
)

# RoA

m1_roa <- feols(roa ~ treat_ind_int + post, data = df_final_4dim, vcov = ~name)
m2_roa <- feols(roa ~ treat_ind_int * post, data = df_final_4dim, vcov = ~name)
m3_roa <- feols(roa ~ treat_ind_int * post | year, data = df_final_4dim, vcov = ~name)
m4_roa <- feols(roa ~ treat_ind_int * post | name + year, data = df_final_4dim, vcov = ~name)
m5_roa <- feols(roa ~ treat_ind_int * post + ln_assets | name + year, data = df_final_4dim, vcov

modelsummary(
  list("(1)" = m1_roa, "(2)" = m2_roa, "(3)" = m3_roa, "(4)" = m4_roa, "(5)" = m5_roa),
  coef_map = c(
    "treat_ind_int"      = "High-emission",
    "post"               = "Post",
```

```
    "treat_ind_int:post" = "High-emission × Post"
  ),
  gof_map = c("nobs", "r.squared"),
  stars = TRUE,
  statistic = "std.error",
  add_rows = tibble(
    term = c("Firm FE", "Year FE", "Controls"),
    `(1)` = c("No",  "No",  "No"),
    `(2)` = c("No",  "No",  "No"),
    `(3)` = c("No",  "Yes", "No"),
    `(4)` = c("Yes", "Yes", "No"),
    `(5)` = c("Yes", "Yes", "Yes")
  ),
  title = "Table 1: Baseline Difference-in-Differences Results (ROA)"
)

# asset turnover

m1_at <- feols(asset_turnover ~ treat_ind_int + post, data = df_final_4dim, vcov = ~name)
m2_at <- feols(asset_turnover ~ treat_ind_int * post, data = df_final_4dim, vcov = ~name)
m3_at <- feols(asset_turnover ~ treat_ind_int * post | year, data = df_final_4dim, vcov = ~name)
m4_at <- feols(asset_turnover ~ treat_ind_int * post | name + year, data = df_final_4dim, vcov =
m5_at <- feols(asset_turnover ~ treat_ind_int * post + ln_assets | name + year, data = df_final_

modelsummary(
  list("(1)" = m1_at, "(2)" = m2_at, "(3)" = m3_at, "(4)" = m4_at, "(5)" = m5_at),
  coef_map = c(
    "treat_ind_int"      = "High-emission",
    "post"               = "Post",
    "treat_ind_int:post" = "High-emission × Post"
  ),
  gof_map = c("nobs", "r.squared"),
  stars = TRUE,
  statistic = "std.error",
  add_rows = tibble(
    term = c("Firm FE", "Year FE", "Controls"),
    `(1)` = c("No",  "No",  "No"),
    `(2)` = c("No",  "No",  "No"),
    `(3)` = c("No",  "Yes", "No"),
    `(4)` = c("Yes", "Yes", "No"),
    `(5)` = c("Yes", "Yes", "Yes")
  ),
  title = "Table 1: Baseline Difference-in-Differences Results (Asset Turnover)"
)


# summary stats
# outcomes + size + emissions/intensity
# -------------------------
vars_sum <- c(
  "ihs_ebitda", "roa", "asset_turnover",
  "assets", "sales", "ebitda_raw",
  "co2_final", "metric_intensity",   # intensity = CO2/sales
  "ln_assets"
)

#
var_labels <- c(
  ihs_ebitda      = "IHS(EBITDA)",
```

```
  roa             = "ROA (EBITDA/Assets)",
  asset_turnover  = "Asset Turnover (Sales/Assets)",
  assets          = "Total Assets",
  sales           = "Net Sales",
  ebitda_raw      = "EBITDA",
  co2_final       = "CO2 Emissions (tCO2e)",
  metric_intensity= "Carbon Intensity (CO2/Sales)",
  ln_assets       = "Log(Assets)"
)

# ------------------------
#  mean/sdAll/Green/Brown+ p-valueGreen vs Brown
# Green = treat_ind_int==0, Brown = treat_ind_int==1
# ------------------------
make_balance_table <- function(data, vars, treat = "treat_ind_int") {

  mean_sd <- function(d){
    d %>%
      summarise(across(all_of(vars),
                       list(mean = ~mean(.x, na.rm = TRUE),
                            sd   = ~sd(.x, na.rm = TRUE)),
                       .names = "{.col}__{.fn}")) %>%
      pivot_longer(everything(),
                   names_to = c("vname","stat"),
                   names_sep = "__",
                   values_to = "value") %>%
      pivot_wider(names_from = stat, values_from = value)
  }

  all_stats <- mean_sd(data) %>% rename(All_Mean = mean, All_SD = sd)

  green_stats <- mean_sd(data %>% filter(.data[[treat]] == 0)) %>%
    rename(Green_Mean = mean, Green_SD = sd)

  brown_stats <- mean_sd(data %>% filter(.data[[treat]] == 1)) %>%
    rename(Brown_Mean = mean, Brown_SD = sd)

  # p value
  pvals <- tibble(vname = vars) %>%
    mutate(
      p_value = map_dbl(vname, function(vn){
        xg <- data[[vn]][data[[treat]] == 0]
        xb <- data[[vn]][data[[treat]] == 1]
        tryCatch(t.test(xg, xb)$p.value, error = function(e) NA_real_)
      })
    )

  all_stats %>%
    left_join(green_stats, by = "vname") %>%
    left_join(brown_stats, by = "vname") %>%
    left_join(pvals, by = "vname")
}


tab_sumstats <- make_balance_table(df_final_4dim, vars_sum, treat = "treat_ind_int")
tab_sumstats

#
```

```
label_map <- c(
  ihs_ebitda       = "IHS(EBITDA)",
  roa              = "ROA (EBITDA / Assets)",
  asset_turnover   = "Asset Turnover (Sales / Assets)",
  assets           = "Total Assets",
  sales            = "Net Sales",
  ebitda_raw       = "EBITDA",
  co2_final        = "CO2 Emissions (tCO2e)",
  metric_intensity = "Carbon Intensity (CO2 / Sales)",
  ln_assets        = "Log(Assets)"
)

tab_sumstats_pretty <- tab_sumstats %>%
  mutate(Variable = recode(vname, !!!label_map)) %>%
  select(Variable, All_Mean, All_SD, Green_Mean, Green_SD, Brown_Mean, Brown_SD, p_value)

# LaTeX
modelsummary::datasummary_df(
  tab_sumstats_pretty,
  output = "summary_stats.tex",
  fmt = 3,
  title = "Summary Statistics: All Firms vs Green vs Brown"
)
#---------------------------------------------------------------------------
# drop fixed
# (1)
m1_new <- feols(
  ihs_ebitda ~ treat_ind_int * post,
  data = df_final_4dim,
  vcov = ~name
)

# (2)
m2_new <- feols(
  ihs_ebitda ~ treat_ind_int * post + ln_assets,
  data = df_final_4dim,
  vcov = ~name
)
# ==============================================================================
# new table
# ==============================================================================

coef_map_final <- c(
  "treat_ind_int"       = "High-emission (Treat)",
  "post"                = "Post-2016 (Post)",
  "treat_ind_int:post"  = "DID (Treat × Post)",
  "ln_assets"           = "Log(Assets)",
  "(Intercept)"         = "Constant"
)

# 2.
modelsummary(
  list(
    "No control"   = m1_new,
    "control size" = m2_new
  ),
  coef_map = coef_map_final,
  stars = TRUE,
  fmt = 4,
```

```
  statistic = "std.error",
  add_rows = tibble(
    term = c("Firm FE", "Year FE"),
    "No control"   = c("No", "No"),
    "control size" = c("No", "No")
  ),
  title = "Table 1: DiD Results showing Level Effects (IHS(EBITDA))"
)
```