Legislator Tweets About the Green Transition and the Returns of Green versus Brown Stocks

Victor DeMiguel Javier Gil-Bazo Milind Goel*

March 27, 2025

Abstract

To study the impact of green-transition regulation on firm value, we analyze stock returns around legislator tweets about climate change. Green stocks significantly outperform brown stocks in the one-to-ten-minute window around pro-transition tweets. The cumulative daily-average green-minus-brown portfolio return around pro-transition tweets is 6.9% higher than around anti-transition ones. For tweets that mention environmental regulations, the spread increases in the weeks preceding a congressional vote and is larger before close votes and when Congress is split. Our findings suggest that the green transition impacts the relative performance of green and brown stocks—at least partly—via a regulatory channel.

Keywords: climate change, sustainability, transition regulation, social media, high-frequency identification
JEL Classification: D72, G10, G12, G14

*DeMiguel: London Business School, avmiguel@london.edu. Gil-Bazo: Universitat Pompeu Fabra and Barcelona School of Economics, javier.gil-bazo@upf.edu. Goel: London Business School, mgoel@london.edu. We are grateful to Roberto Gómez-Cram for extensive discussions and suggestions. We are also grateful for comments from Juan Imbet, Alberto Martín-Utrera, Markus Pelger, and seminar participants at WHU-Otto Beisheim School of Management and ESCP Business School.

1 Introduction

Climate change has prompted legislators worldwide to regulate carbon emissions, a process known as the green transition. To study the impact of green-transition regulation on firm value, we compare the returns of green and brown stocks in the minutes around tweets about climate change posted by members of U.S. Congress. We find that green stocks outperform brown stocks in the one-to-ten-minute window around pro-transition tweets, with the return spread statistically significant at the 1% level. The effect is also economically significant: accumulating the daily-average one-to-ten-minute return of a green-minus-brown portfolio around pro-transition tweets yields a return 6.9% higher than around anti-transition tweets. Moreover, we find that the market correctly anticipates which high-emission firms can decarbonize their operations in response to green-transition shocks. To focus on the regulatory channel, we consider legislator tweets that mention specific environmental regulations, and find that the spread between the returns of green and brown stocks around tweets that support pro-transition regulations increases in the weeks leading to a congressional vote and vanishes afterward. Moreover, the spread is larger before close votes and when Congress is split, with the two chambers controlled by different parties. Our findings suggest that the green transition impacts the relative performance of green and brown stocks, at least in part, via a regulatory channel.

Legislators often use social media platforms such as Twitter (now X) to express their views on the appropriate pace and specific policies required to transition to a green economy. For example, consider the Climate Action Now Act, voted on in the House of Representatives in 2019, which was intended to prevent the president from withdrawing the United States from the Paris Agreement and require the administration to develop a plan for meeting the agreement's emissions reduction goals.[1] As shown in Figure 1, while Democratic Representative Darren Soto tweeted the day before the vote to support the regulation, emphasizing that it would prevent the Trump administration's withdrawal from the Paris Agreement, Republi-

[1]The bill passed the House of Representatives on May 2, 2019, but was never brought to a vote in the Senate, as the Republican majority, led by then-Senate Majority Leader Mitch McConnell, opposed the measure and declined to advance it, aligning with the Trump administration's decision to exit the Paris Agreement.

Figure 1: Legislator tweets about Climate Action Now Act

This figure contains two legislator tweets about the the Climate Action Now Act, voted on in the House of Representatives in 2019, which was intended to prevent the president from withdrawing the United States from the Paris Agreement and require the administration to develop a plan for meeting the agreement's emissions reduction goals. While Democratic Representative Darren Soto tweeted the day before the vote to support the regulation, Republican Representative Doug LaMalfa tweeted against it.



can Representative Doug LaMalfa tweeted to oppose the regulation, urging his colleagues to vote against the "partisan attack on American energy production." These statements contain information about the probability that legislators will support pro-transition regulations in Congress, and thus, may affect firm value.

One would expect transition regulation to have a negative impact on the value of brown firms because of the cost of adapting their production and distribution processes to comply with new regulation, as well as the risk of fines for non-compliance. In contrast, the impact of transition regulation on the value of green firms is likely to be less negative because they are less vulnerable to such regulatory burdens, compared to brown firms. Indeed, green firms may even benefit from climate policies through subsidies or tax incentives. Ultimately, whether and how the effect of transition regulation on brown and green firms differs is an empirical question. However, measuring the impact of transition regulation on firm value is challenging. Consider, for instance, a difference-in-differences study that exploits the

staggered adoption of environmental regulation in different jurisdictions. The validity of causal estimates from such a study is questioned by the endogenous nature of the decision to adopt the regulation in a particular jurisdiction and the strategic response of firms to anticipated regulation.

To isolate the impact of transition regulation on firm value from that of other potential confounding factors, we leverage the high-frequency identification approach of Bianchi, Gómez-Cram, and Kung (2024) and study the returns of green and brown stocks in the one-to-ten-minute window around legislator tweets about climate change. The rationale for this approach is that legislator tweets contain new information about their views on the green transition, and thus about the probability that transition regulations will pass future congressional votes. Upon observing the legislator tweets, investors update their expectations about future regulations and their valuations of green and brown firms. Thus, studying stock returns in the minutes around legislator tweets about climate change allows us to isolate the effect of news about green-transition regulation from other confounders. The identifying assumption is that no other information is released over such a short window around the tweet that systematically affects the prices of green and brown stocks. To perform the test, we download all tweets posted between January 1, 2014 and March 31, 2023 by the official Twitter accounts of U.S. Congress members. We identify 38,055 legislator tweets about climate change posted during trading hours and use ChatGPT to measure their sentiment toward the green transition, defined as the extent to which a tweet is in favor or against the green transition.[2] Each day, we classify the 1,000 largest U.S. stocks as green, gray, or brown based on whether their most recently disclosed greenhouse gas emissions intensity falls in the bottom, middle, or top tercile of the cross section. We then run a panel regression of stock returns around each tweet on tweet sentiment toward the green transition as well as the interaction of tweet sentiment with green and brown dummy variables.

Our main finding is that a one standard deviation increase in legislator-tweet sentiment toward the green transition is associated with a decrease in the one-to-ten-minute return around the tweet of -0.057 basis points (bps) for green stocks, -0.140 bps for gray stocks,

[2]We validate our ChatGPT sentiment measure by showing in Section IA.1 of the Internet Appendix that it is highly correlated with those generated by five researchers for a random sample of 100 legislator tweets.

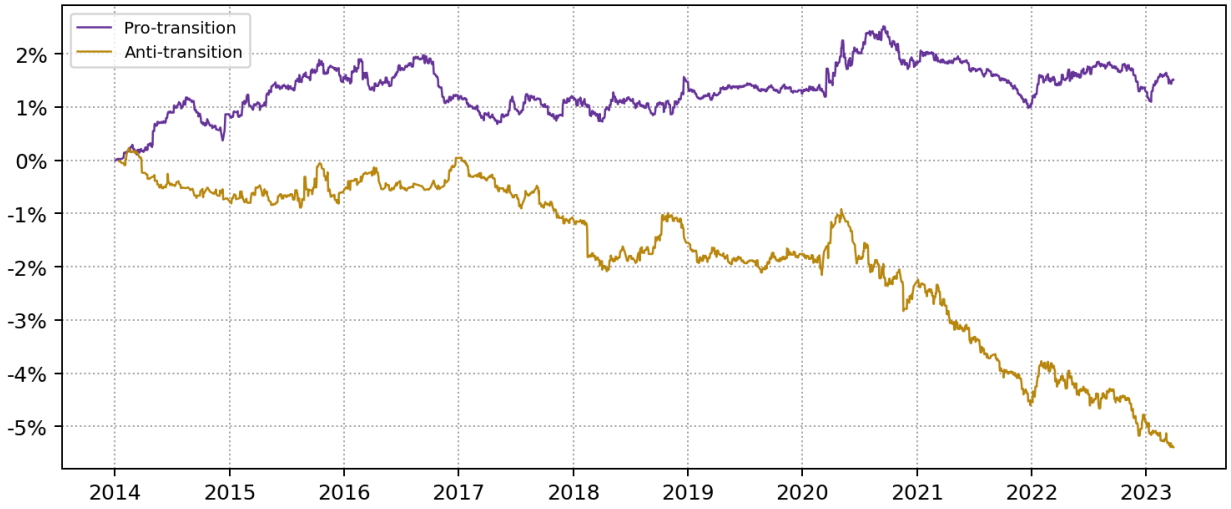
and −0.194 bps for brown stocks, with the differences between these three returns significant at the 1% level. Thus, legislator support toward the green transition reduces the value of all firms, a result confirming the economic insight that regulation imposes a financial burden or operational constraints on all firms. However, brown firms are more negatively impacted by the pro-transition tweets because of their higher exposure to environmental regulation.

The economic significance of the effect of legislator tweets on stock returns increases as we extend the horizon from 10 minutes after the tweet to one day. For instance, the return spread between the returns of green and brown stocks increases from 0.137 bps to 0.319 bps with each standard deviation increment in tweet sentiment toward the green transition when we extend the event horizon from 10 to 90 minutes. The spread further increases more than tenfold to 1.471 bps one day after the tweet. These findings suggest that it takes at least one day for stock prices to incorporate the information contained in legislator tweets. We further demonstrate that our baseline analysis remains robust to contracting the event horizon around which stock returns are computed, adjusting for market returns, including more granular time fixed effects, using alternative sources of carbon emissions data, and using an alternative carbon intensity measure.

To further gauge the economic relevance of our findings, we compute value-weighted green-minus-brown portfolio returns in the one-to-ten-minute window around legislator tweets. On each day, we compute the average green-minus-brown portfolio return, separately, around pro-transition and around anti-transition tweets, which we define as those categorized by ChatGPT as being supportive or unsupportive toward the green transition. We then accumulate these average returns over the duration of our sample. Figure 2 depicts the resulting cumulative daily-average returns of a value-weighted green-minus-brown portfolio around pro-transition tweets (purple line) and around anti-transition tweets (yellow line). The aggregate green-minus-brown portfolio return accumulates to 1.5% around pro-transition tweets and −5.4% around anti-transition tweets. This 6.9% spread is economically significant because we effectively accumulate the returns around a single tweet on each of the 2,270 trading days in our sample.

Figure 2: Green-minus-brown returns around pro- and anti-transition tweets

This figure reports the cumulative daily-average returns of a value-weighted green-minus-brown portfolio around pro-transition tweets (purple line) and around anti-transition tweets (yellow line). We first compute value-weighted green-minus-brown portfolio returns in the one-to-ten-minute event horizon around each tweet using the market capitalization of firms at the end of the previous day. We then identify pro- and anti-transition tweets using a categorical sentiment measure obtained from ChatGPT. Because the number of pro-transition tweets in our sample is 3.54 times greater than that of anti-transition tweets, simply accumulating returns around tweets falling under each category would result in an unfair comparison. Thus, on each day, we first compute the *average* of the one-to-ten-minute green-minus-brown portfolio returns around pro-transition tweets and around anti-transition tweets, and then accumulate these average daily high-frequency returns. We exclude days on which we do not have *both* pro-transition and anti-transition tweets, resulting in a time series of 2,270 trading days.



Certain high-emission firms are likely to be more able or willing to adapt their carbon emissions to news about transition regulation than others. If markets can predict which firms will adapt their carbon emissions in response to transition-regulation news, then one would expect that firm returns around pro-transition tweets should be informative about their future carbon emissions. To test this hypothesis, we regress the change in carbon emissions intensity between consecutive firm disclosures on the aggregate stock return around pro-transition legislator tweets during the year leading to the first disclosure. We find that there is no significant relation between the price reaction of *low-emission* firms around legislator tweets and changes in their future carbon emissions intensity, which could be explained by low-emission firms being already carbon efficient. However, high-emission firms with higher returns around pro-transition tweets decarbonize more in the future. This finding reveals that the market correctly identifies and rewards high-emission firms that can better adapt

to transition regulation, and thus suggests that the green transition impacts not only firm value, but also real outcomes such as carbon emissions. In addition, our results indicate that brown firms that are not able or willing to decarbonize shoulder a disproportionately large share of the burden imposed by transition regulations.

We then examine the mechanism through which the legislator tweets impact stock prices. First, we show that the legislator tweets about climate change contain information that is distinct from that in mainstream media. In particular, we construct a daily ProTransition Sentiment (PTS) index by adding the numerical legislator sentiment toward the green transition across all tweets posted each day. Reassuringly, the PTS index spikes around days on which climate-related regulations are introduced and voted on in Congress. We then study the lead-lag relation between the PTS index and the Media Climate Change Concerns (MCCC) index of Ardia, Bluteau, Boudt, and Inghelbrecht (2023), which proxies climate-change concerns expressed in mainstream media. To do this, we estimate a vector autoregressive model and find that both indices significantly predict future values of themselves and each other, which demonstrates that at least part of the information in legislator tweets is distinct from that in mainstream news. This suggests the impact of legislator tweets on the relative performance of green and brown stocks occurs via an information channel.

Second, we focus on the regulatory channel by using a dictionary approach to identify a sample of 40,821 legislator tweets that mention specific environmental regulations. Reproducing our baseline analysis on this complementary sample also reveals a spread in the returns of green versus brown stocks around legislator tweets that mention environmental regulations. Consistent with the hypothesis that these regulatory tweets affect the relative value of green and brown stocks because they contain information about the probability that regulation passes a congressional vote, we find that the spread increases in the weeks leading up to a congressional vote and becomes statistically insignificant afterward. Moreover, the spread is larger before votes that eventually pass by a close margin and when Congress is split, with a different party controlling each of the two chambers. These findings suggest that the green transition impacts the relative performance of green and brown stocks, at least partly, via a regulatory channel.

Our results have important implications for policymakers, firms, and investors. Policymakers can alleviate the impact of the green transition on brown firms by adopting a consistent approach to the green transition, thus avoiding abrupt shifts in policy that amplify the impact of transition regulation on brown firms. For instance, consider a transition regulation that is approved during a particular presidential or congressional term, but may be repealed or amended during the next term. Firms have to decide whether and how to adapt their operations to the new regulation, knowing that it may be revised or repealed in the future. This reduces firm incentives to adapt and introduces additional costs if Congress eventually revises the regulation in the future. Our work has implications also for high-emission firms—our analysis of the association between stock returns around pro-transition tweets and future carbon emissions shows that high-emission firms can add value by enhancing their ability to adjust their emissions in response to new regulations. Finally, our manuscript demonstrates that investors should measure asset exposure to transition regulation and manage this exposure in their portfolio—for example, by holding stocks of green firms or brown firms with credible plans to adapt their emissions.

Our work contributes to the climate-finance literature. Pástor, Stambaugh, and Taylor (2021) propose an equilibrium model that incorporates sustainability criteria and generates two predictions. First, "green assets have low expected returns because investors enjoy holding them and because green assets hedge climate risk." Second, "green assets nevertheless outperform" when there are positive shocks to "customers' tastes for green products and investors' tastes for green holdings." A growing literature investigates the first prediction by examining the effect of firm-level carbon emissions on the cross section of stock returns to test whether investors demand compensation for their exposure to carbon emission risk (Bolton and Kacperczyk, 2021, 2023; Eskildsen et al., 2024). In contrast, we focus on the second prediction of Pástor et al. (2021), similar to Ardia et al. (2023), who construct a daily MCCC index using mainstream media news and show that green firms outperform brown firms when media concerns about climate change increase unexpectedly. However, green stocks should also outperform brown stocks when pro-transition *regulation* becomes more likely, because it burdens non-complying firms. Our main contribution is to test this prediction directly by studying stock returns around climate-change tweets posted by legislators.

Moreover, Ardia et al. (2023, p. 7608) acknowledge that endogeneity may be driving the relation between their MCCC index and stock returns: "endogeneity may arise due to a potential feedback loop that an exogenous shock leading to increased concern in climate change affects the pricing of green and brown stocks," and thus they "abstain from any causal interpretation" of their results. Similar endogeneity concerns affect other papers that also estimate the correlation between daily or weekly climate-related news indices and stock returns, such as Pástor, Stambaugh, and Taylor (2022) and Acharya, Giglio, Pastore, Stroebel, Tan, and Yong (2024).[3] Another contribution of our manuscript to the climate-finance literature is to address these endogeneity concerns by leveraging the high-frequency identification approach of Bianchi et al. (2024) to study price movements in the minutes around legislator tweets.

Bianchi et al. (2024) examine legislator tweets that mention specific firms and find that the firm stock returns increase with positive tweet sentiment. Cookson, Fox, Gil-Bazo, Imbet, and Schiller (2024) use a similar identification strategy to analyze the effect of tweets on the distress of U.S. banks during the banking crisis of 2023. In contrast, we shift the focus from firm-specific tweets to *climate-change* tweets, investigating their impact on the prices of a *cross-section* of firms ranked by their carbon emissions intensity.[4]

Our work is also related to the emerging literature that uses textual analysis in the context of climate risk. Li, Shan, Tang, and Yao (2024) produce firm-level measures of transition climate risk by applying a dictionary approach to earnings conference calls, and find that transition risk is associated with negative returns for non-proactive firms. Sautner, Van Lent, Vilkov, and Zhang (2023) use machine learning to analyze earnings calls and

[3] Pástor et al. (2022) use the climate change concerns index of Ardia et al. (2023) to show that although green assets delivered high returns in recent years, their outperformance reflects unexpectedly strong increases in environmental concerns, not high expected returns, consistent with the predictions of Pástor et al. (2021, p. 550). Acharya et al. (2024) develop a theoretical model to examine the impact of climate transition risk on the performance of various types of energy firms. To empirically test the model's predictions, they construct three weekly climate transition risk indices based on New York Times news, and analyze the correlation between energy prices, energy firm stock returns, and index innovations. Their empirical results provide strong support for their model's predictions.

[4] Our focus on the effect of social media posts on the cross section of stock returns also distinguishes our manuscript from the literature that leverages high-frequency identification in macroeconomics (Bernanke and Kuttner, 2005; Nakamura and Steinsson, 2018; Bianchi, Comin, Kung, Kind, and Matusche, 2019; Gómez-Cram and Grotteria, 2022; Känzig, 2021), which primarily studies aggregate economic quantities such as market returns, inflation, interest rates, and output.

produce firm-level measures of climate-change exposure, and find that the measures predict asset prices and real outcomes. Kölbel, Leippold, Rillaerts, and Wang (2024) apply BERT to 10K filings to generate firm-level measures of physical and transition climate risks, and show that disclosure of transition (physical) risk is associated with an increase (decrease) in credit default swap spreads. A distinguishing feature of our work is that we leverage a high-frequency approach to identify the impact of legislator tweets about climate change on the returns of green versus brown stocks.

Empirically, we find that, while the conventional dictionary and machine-learning approaches used in the literature adequately identify climate-related keywords and the overall tone of a post, they do not capture its sentiment toward green transition. For instance, although the simple phrase—"carbon tax has decreased"—expresses a negative sentiment toward the green transition, conventional approaches may classify it as positive because it combines the words *decrease* and *tax*. In contrast, ChatGPT captures this distinction. Thus, one contribution of our manuscript is to show how ChatGPT can be used to obtain 'context-aware' sentiment measures. This positions our work within the nascent literature that harnesses ChatGPT for applications in economics and finance (Lopez-Lira and Tang, 2023; Cowen and Tabarrok, 2023; Korinek, 2023; Ko and Lee, 2024; Bianchi et al., 2024; Acharya et al., 2024).

Finally, our work is related to the literature that measures the impact of political cycles on how firm value depends on carbon emissions (Santa-Clara and Valkanov, 2003; Pástor and Veronesi, 2020; Addoum and Kumar, 2016). Ramelli, Wagner, Zeckhauser, and Ziegler (2021) show that both carbon-intensive and climate-responsible firms gained following the 2016 election of Donald Trump. They attribute the gains of climate-responsible firms to investors betting on a reversal in climate policy after the 2020 election. Baz, Cathcart, Michaelides, and Zhang (2023) use 10K filings to generate a firm-specific metric of climate-change exposure. They show that the returns of high-exposure firms increased during the twelve months after President Trump's election, a finding they attribute to a subsequent decline in climate-change concern. Mathurin (2024) uses firm press releases to measure firm exposure to environmental policies and green technologies. She shows that the risk premia

associated with these two exposures are negatively correlated and their sign and magnitude depend on the political orientation of the government. Our approach based on a real-time stream of legislator tweets provides a more granular view of the effect of politician viewpoints on the value of green versus brown firms.

2 Data

This section describes the data used in the manuscript. Section 2.1 describes how we download legislator tweets, Section 2.2 how we identify climate-change tweets, Section 2.3 how we measure legislator-tweet sentiment toward the green transition, Section 2.4 how we compile minute-level stock prices, and Section 2.5 how we compute carbon emissions intensity.

2.1 Compiling legislator tweets

We obtain the official legislative Twitter handles of 489 out of 541 members of the U.S. Congress that served as of March 31, 2023 from `@unitedstatesproject`, as well as information on their house, duration, and party membership. Of the 489 legislators, 244 were Democrats, 242 Republicans, and three independents. For each legislator, we use the Twitter Developer API to download all tweets posted between January 1, 2014 and March 31, 2023 during each term the legislator served. In total, we download 2,830,646 tweets. Of the 2,830,646 tweets, 62.16% were posted by Democrats, 37.12% by Republicans, and only 0.72% by independents.

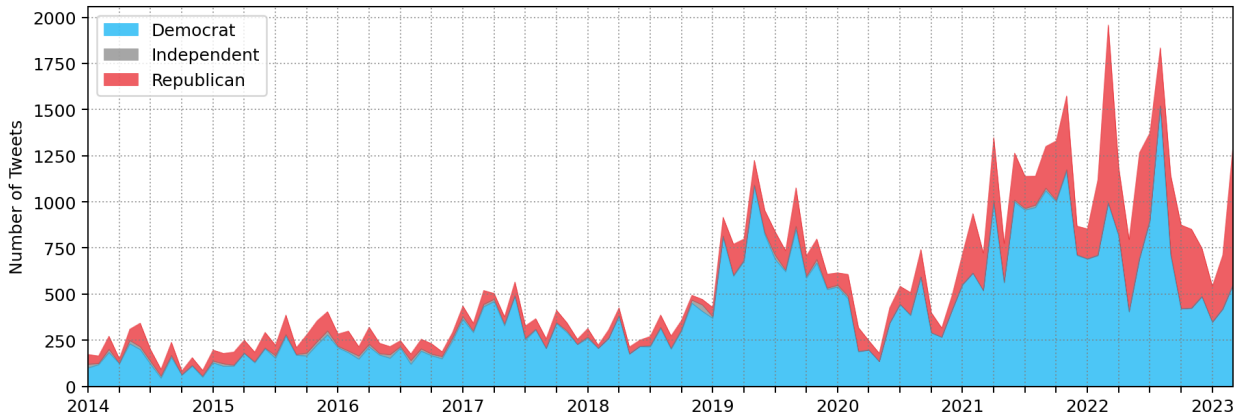
2.2 Identifying climate-change tweets

To identify climate-change tweets, we use the ESG4 and ESG9 versions of the large-language model (LLM) FinBERT, which are specifically fine tuned to classify text into ESG categories.[5] The ESG4 model labels text into four categories (Environmental, Social, Gov-

[5] Huang, Wang, and Yang (2020) developed FinBERT by using financial texts to fine tune Google's BERT LLM. An advantage of using FinBERT to identify climate-change tweets is that we have a very large number

Figure 3: Monthly number of climate-change tweets by party

This figure depicts the breakdown of the monthly number of tweets posted by Republicans (red), Democrats (blue), and Independents (gray) between January 2014 and March 2023.



ernance or None categories) and the ESG9 model into nine categories (Climate Change, Natural Capital, Pollution & Waste, Human Capital, Product Liability, Community Relations, Corporate Governance, Business Ethics & Values, and Non-ESG). The labels produced by the two models are further associated with a confidence score between zero and one.

We select 62,067 out of 2,830,646 tweets that are simultaneously labeled as Environmental by the ESG4 model and as Climate Change by the ESG9 model with a confidence score of at least 90%.[6] Figure 3 depicts the time series of the monthly number of climate-change legislator tweets posted by members of each party. Over our sample, 70.7% of climate-change tweets come from Democrats, 27.8% from Republicans, and 1.5% from Independents.

2.3 Evaluating tweet sentiment toward the green transition

We use ChatGPT-3.5 to measure tweet sentiment toward the green transition, leveraging ChatGPT's ability to produce contextually relevant responses to specific questions. We

of over 2.8 million legislator tweets and FinBERT is much less computationally demanding compared to other LLM models such as ChatGPT.

[6]We use both the ESG4 and ESG9 models to ensure that we select tweets that contain climate-change content. One would expect that tweets labeled as Environmental by the ESG4 model would fall into either the Climate Change, Natural Capital, or Pollution & Waste categories of the ESG9 model. However, this is not always the case; tweets classified as Environmental are also sometimes classified into the other six ESG9 categories. This suggests that there is some noise in the classification or that the two models are picking up on slightly different contexts. The latter is plausible because ESG4 FinBERT is trained on 2,000 labeled sentences, while ESG9 FinBERT is separately trained on 14,000 labeled sentences.

query ChatGPT to answer two *mirror* prompts for each tweet. The first prompt asks: "*You are a professor of finance. Please produce a single percentage value between 0% and 100% to capture the extent to which the following tweet is in favor of green transition.*" The second prompt asks: "*You are a professor of finance. Please produce a single percentage value between 0% and 100% to capture the extent to which the following tweet is against green transition.*" We set the temperature to zero to obtain deterministic responses.[7] Finally, we compute the numerical sentiment measure (ranging between −100 and 100) by subtracting the *against*- from the *in-favor*-measure. We use mirror prompts because a single prompt typically leads the model to interpret or frame sentiment in one direction, potentially skewing the response toward either positive or negative sentiment.[8]

Figure 4 graphs the distribution of the numerical sentiment measure. We observe that the distribution is left-skewed and that there are more climate-change legislator tweets in favor of the green transition than against. Also, Figure IA.1 in the Internet Appendix depicts the time series of the average monthly numerical sentiment score by party affiliation over the duration of our sample. We observe that Democrats are consistently more supportive of the green transition, compared to Republicans.

We also query ChatGPT to obtain a categorical measure of sentiment toward the green transition—"*You are a professor of finance. Please determine whether the following tweet has a favorable, unfavorable, or neutral stance on green transition. Output only a single word*"—with temperature set to zero. Consistent with the results in Figure 4, ChatGPT labels 62.4% tweets as favorable, 18.8% as neutral, and 17.8% as unfavorable.

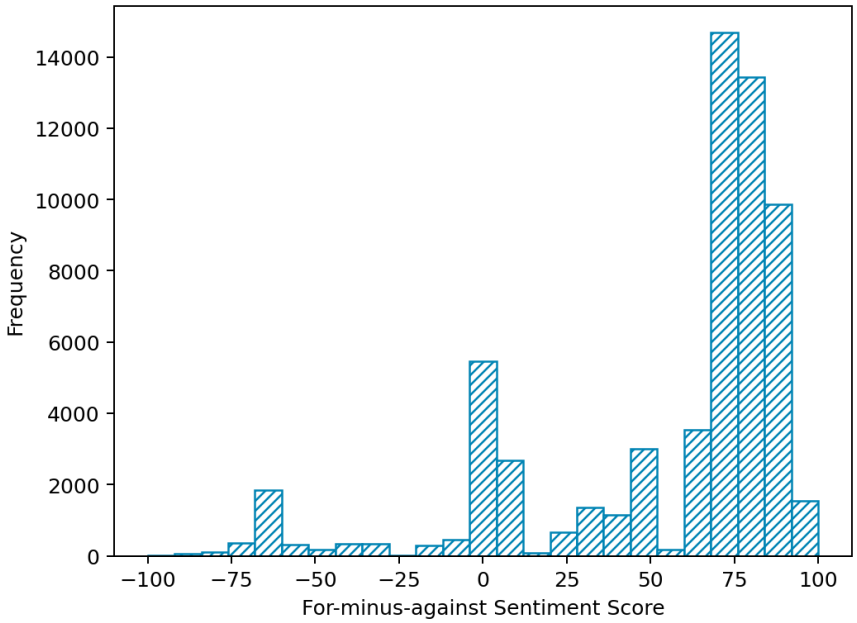
We validate the numerical and categorical ChatGPT sentiment measures by comparing them with those generated by five researchers for a random sample of 100 tweets. Section IA.1

[7] Temperature is a ChatGPT parameter ranging between zero and two that controls the creativity, and thus stochasticity, of responses.

[8] Because the model is obligated to return a value, neutral or ambiguous sentiments may not be accurately captured by a single prompt. For example, a tweet neutral toward green transition might receive similar scores for both directions. Using both in-favor and against prompts helps ensure that the final sentiment close to zero, which better reflects neutrality. Furthermore, language models may interpret prompts differently depending on wording, context, or other subtle variations. Using both in favor and against prompts helps create symmetry in phrasing, ensuring that the overall sentiment is not skewed by the model's inherent interpretation biases in just one direction.

Figure 4: Distribution of ChatGPT sentiment

This histogram depicts the empirical distribution of the numerical tweet sentiment toward the green transition computed by subtracting the ChatGPT-generated *against*-measure from the *for*-measure.



of the Internet Appendix shows that ChatGPT sentiment measure is highly correlated with those generated by the researchers.

While conventional dictionary (Loughran and McDonald, 2011) and machine-learning approaches (Kvam, Molnár, Wankel, and Ødegaard, 2022) are likely to adequately measure the overall tone of a tweet, they are unlikely to capture its sentiment toward green transition. In contrast, ChatGPT performs well at measuring tweet sentiment toward the green transition. For instance, Table IA.1 in the Internet Appendix shows the numerical sentiment measure (obtained from ChatGPT) for a randomly selected sample of tweets classified as 'Negative' by the Tone FinBERT model, a version of the FinBERT LLM fine tuned to measure text sentiment. We find that while the sentiment measure obtained from ChatGPT is not perfect, it correctly identifies several tweets in the table that while containing negative keywords, are overall positive toward the green transition.

2.4 *Compiling high-frequency stock prices*

We select 38,055 out of 62,067 climate-change legislator tweets posted during normal NYSE trading hours. We obtain minute-level stock prices from FirstRate and monthly stock market capitalization from the Center for Research in Security Prices (CRSP). For each calendar year, we sort stocks based on their market capitalization during the last month of the previous calendar year, and then select the largest 1,000 stocks. We must query the high-frequency price data for *each* stock and *each* tweet, and thus, selecting the largest 1,000 stocks helps to reduce computational time.

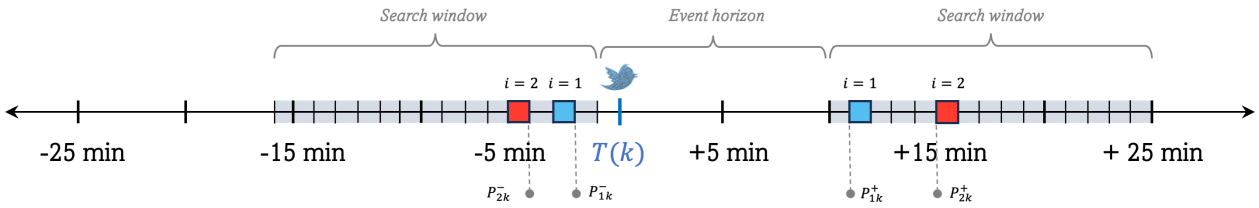
2.5 *Obtaining carbon emissions data*

Like Ardia et al. (2023), we use the Refinitiv Eikon API to obtain firms' annual greenhouse gas (GHG) emissions and revenue data. GHG emissions are reported in metric tonnes of carbon dioxide equivalent (tCO_2e). Around 97% of the GHG emissions data obtained from Refinitiv is disclosed by firms. We discard the remaining entries that are estimated. We then compute the carbon emissions intensity of a firm in a given fiscal year by dividing the sum of its absolute Scope 1 and Scope 2 GHG emissions by total revenue (expressed in million USD). Scope 1 emissions are direct GHG emissions that originate from sources owned or controlled by the firm, such as company vehicles and on-site fuel combustion. Scope 2 emissions are indirect GHG emissions that originate from the generation of purchased electricity, steam, heating, and cooling consumed by a firm. Like Ardia et al. (2023), we shift carbon emissions intensity values by one year because emissions are reported with a one-year delay. To perform robustness checks, we also compute carbon emissions intensity using S&P Trucost data.[9]

[9]Refinitiv also provides data on Scope 3 emissions, which originate from a firm's upstream and downstream activities. We exclude Scope 3 emissions when computing carbon emissions intensity because only 60.7% of the firms that report Scope 1 and Scope 2 emissions also report Scope 3 emissions. Thus, using Scope 3 emissions substantially reduces our sample size. However, as we discuss in Section 3.5, Column (6) of Table 2 shows that our main findings are robust to including Scope 3 emissions when computing carbon emissions intensity. This is consistent with the findings of Ardia et al. (2023), who note that inclusion of Scope 3 emissions is inconsequential. Table IA.2 in the Internet Appendix ranks the Fama-French 48 industries by their average carbon intensity measures, and reports the average absolute Scope 1, 2, and 3 emissions for each industry.

Figure 5: Minute-level stock prices around legislator tweets

This figure illustrates how we query minute-level stock prices around legislator tweets. The symbol $T(k)$ denotes the timestamp of the kth tweet. The event horizon is the time interval from one minute before to ten minutes after the tweet. The search windows are the fifteen-minute lookup intervals before and after the event horizon (grey shaded regions). For the search window before the event horizon, we query the closing price for the last minute for which there is price data available. For the search window after the event horizon, we query the opening price for the first minute for which there is price data available. The blue and the red boxes represent minute-level prices of two different stocks.



3 Empirical analysis

This section examines how legislator-tweet sentiment toward the green transition impacts the relative prices of green and brown stocks.

3.1 High-frequency identification

We first describe how we compute high-frequency stock returns around legislator tweets following the approach in Bianchi et al. (2024). Let $T(k)$ be the timestamp of the kth tweet. For the kth tweet and ith stock, we download the following minute-level prices from FirstRate: (i) closing price for the *last* minute in which a trade was recorded within the window $[T(k)-16 \text{ min}, T(k)-1 \text{ min}]$ *before* the tweet, denoted by P^-_{ik}, and (ii) opening price for the *first* minute in which a trade was recorded within the window $[T(k)+10 \text{ min}, T(k)+25 \text{ min}]$ *after* the tweet, denoted by P^+_{ik}. Figure 5 illustrates this process.

We term the time interval from one minute before to ten minutes after the tweet as the *event horizon* and the fifteen-minute lookup intervals before and after the event horizon as the *search windows*. Because we are working with minute-level prices, querying stock prices one minute before a tweet provides the tightest event horizon achievable.[10] Extending the

[10]To ensure that P^-_{ik} occurs at least one minute before the tweet is published, we obtain the closing price associated with the timestamp two minutes before that tweet is published.

event horizon to ten minutes after a tweet provides reasonable time for market participants to assimilate the information and react, while being sufficiently short to avoid the confounding impact of other events that may occur around the time of publication of the tweet. We then calculate the change in stock price of firm i in the event horizon around the kth tweet as:

$$\Delta p_{ik} = \log(P_{ik}^{+}) - \log(P_{ik}^{-}).$$

We set $\Delta p_{ik} = 0$ if there are no minute-level prices recorded in FirstRate during either the search window before or after the tweet, and winsorize Δp_{ik} at the 0.5^{th} and 99.5^{th} percentiles.[11] Throughout the paper, returns are expressed in basis points.

We now describe our baseline panel regression. At the beginning of each day, we sort firms by their carbon emissions intensity, and categorize firms that fall into the bottom, middle, and top terciles as *green*, *gray*, and *brown*, respectively.[12] We then estimate the following panel regression:

$$\begin{aligned} \Delta p_{ik} = \; & \gamma_G \cdot \mathbb{1}_{Green(i,k)} + \gamma_B \cdot \mathbb{1}_{Brown(i,k)} \\ & + \tau \; \cdot \textit{Sentiment(k)} \\ & + \tau_G \cdot \textit{Sentiment(k)} \times \mathbb{1}_{Green(i,k)} \\ & + \tau_B \cdot \textit{Sentiment(k)} \times \mathbb{1}_{Brown(i,k)} \\ & + \theta_i + \theta_a + \theta_q + \epsilon_{ik}, \end{aligned} \tag{1}$$

where $\mathbb{1}_{Green(i,k)}$ and $\mathbb{1}_{Brown(i,k)}$ are indicator variables equal to one when the ith firm is green and brown, respectively, at the timestamp of the kth tweet, and *Sentiment(k)* is the kth tweet numerical sentiment toward the green transition, obtained as described in Section 2.3, min-max normalized annually between zero and one. The terms θ_i, θ_a, and θ_q capture firm, legislator, and year-quarter fixed effects, respectively.

The coefficients γ_G and γ_B measure the average returns of green and brown stocks around legislator tweets with zero normalized sentiment (strongly against) toward the green transition, relative to observations of the same firms when classified as gray, the omitted category.

[11] According to the data provider, missing minute prices imply that no trade occurred during that minute for that stock. Thus, we assume that the tweet did not impact the price of the stock.

[12] At the beginning of each day, we consider the most recent carbon intensity measure available for each stock. If, for a given stock, the available carbon intensity measure is older than one year, then we assume that the information is not available for that stock on that day.

The coefficient τ measures the marginal effect of legislator-tweet sentiment toward the green transition on the prices of gray firms. The coefficients τ_G and τ_B capture the incremental impact of legislator-tweet sentiment on the returns of green and brown stocks, relative to gray stocks. The error term is ϵ_{ik}.

3.2 Impact of legislator-tweet sentiment on stock prices

Table 1 reports the results obtained from estimating the panel regression in Equation (1) for various event horizons. Column (1) of Table 1, in particular, reports the results for our base-case event horizon from one minute before to ten minutes after each tweet.

We focus our discussion on the coefficients on the third, fourth, and fifth rows in Table 1, which characterize the impact of legislator-tweet sentiment toward the green transition on the returns of gray, green, and brown firms.[13] The third row reports the slope coefficient for legislator sentiment toward the green transition, and it shows that a one standard deviation increase in sentiment (0.206) leads to a reduction in the returns of all stocks (green, gray, and brown) around the tweet of -0.140 basis points ($= 0.206 \times 0.678$ bps), significant at the 1% level, which translates to a decrease in the market value of an average firm in our sample of \$302,820.[14] Thus, the market reacts negatively to pro-transition legislator tweets, presumably because investors believe that environmental regulation imposes a financial burden or operational constraints on all businesses.

The fourth and fifth rows in Table 1 report the slope coefficients for the interactions between the green and brown indicator variables and the legislator sentiment variable. We observe that when legislator-tweet sentiment toward the green transition increases by one standard deviation, the prices of brown firms decrease by an additional 0.054 basis points ($= 0.206 \times 0.260$ bps) relative to those of gray firms, significant at the 1% level. In contrast,

[13] The first two rows in Table 1 report the slope coefficients for the brown and green indicator variables. The average returns of green firms around legislator tweets with zero normalized sentiment (strongly against) are 0.284 basis points lower than those of gray firms, significant at the 1% level, while those brown firms are 0.128 basis points higher but insignificant.

[14] The average market capitalization of a firm in our sample is 21.63 billion.

Table 1: Impact of legislator tweets on the returns of green versus brown stocks

This table reports the results obtained from estimating the panel regression in Equation (1) for various event horizons. Columns (1)–(6) report the results for event horizons from one minute before the tweet to $s = 10, 15, 30, 60, 90$ minutes and one day after each tweet, respectively, where the label $\Delta p(-01, +e)$ indicates an event horizon from one minute before to e minutes after each tweet. The sentiment of tweets toward the green transition is obtained using ChatGPT and min-max normalized annually between zero and one. Stocks that fall into the top and bottom tercile of carbon emissions intensity—obtained by using the most recent Scope 1 and 2 emissions data available on Refinitiv Eikon at the time each tweet was posted—are classified as green and brown, with the rest classified as gray. We report all intercept and slope coefficients in basis points. Standard errors are robust to heteroscedasticity and clustered at the firm level. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) stars, respectively.

	(1) $\Delta p(-01,+10)$	(2) $\Delta p(-01,+15)$	(3) $\Delta p(-01,+30)$	(4) $\Delta p(-01,+60)$	(5) $\Delta p(-01,+90)$	(6) $\Delta p(-01,+1\text{d})$
Brown	0.128	0.0986	0.110	-0.128	-0.182	4.245***
	(1.27)	(0.84)	(0.60)	(-0.43)	(-0.43)	(2.62)
Green	-0.284***	-0.361***	-0.579***	-0.639**	-1.178***	-4.183*
	(-2.68)	(-2.83)	(-2.99)	(-2.06)	(-2.59)	(-1.91)
Sentiment	-0.678***	-0.524***	-0.908***	-1.531***	-4.145***	0.491
	(-9.34)	(-6.06)	(-6.49)	(-7.17)	(-14.13)	(0.66)
Brown × Sentiment	-0.260***	-0.245**	-0.418**	-0.137	-0.269	-3.121**
	(-2.59)	(-2.02)	(-2.18)	(-0.47)	(-0.64)	(-2.57)
Green × Sentiment	0.402***	0.470***	0.742***	0.965***	1.549***	4.019***
	(4.13)	(4.07)	(4.08)	(3.56)	(4.05)	(3.48)
Constant	0.532***	0.366***	0.556***	0.952***	3.314***	3.326***
	(8.12)	(4.64)	(4.50)	(4.85)	(12.32)	(3.39)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Legislator FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R2	0.00707	0.00759	0.00735	0.00803	0.00901	0.0139
Observations	12131416	12065147	11719374	10909570	10019410	12037363

the prices of green firms increase by 0.083 basis points ($= 0.206 \times 0.402$ bps) relative to gray firms, significant at a 1% level.

Overall, a one standard deviation increase in legislator-tweet sentiment toward the green transition is associated with a decrease in the one-to-ten-minute return around the tweet of $0.057 = 0.140 - 0.083$ basis points (bps) for green stocks, 0.140 bps for gray stocks, and $0.194 = 0.140 + 0.054$ bps for brown stocks. Thus, while legislator support toward the green transition reduces the value of all firms, brown firms are more negatively impacted because of their higher exposure to green-transition regulation. The spread between the returns of green and brown stocks around a legislator tweet is 0.137 bps (0.054 bps + 0.083 bps), which

translates into a $296,331 higher market valuation of a green firm relative to an otherwise similar brown firm around a legislator tweet.

3.3 Persistence over time

Does the relative effect of legislator tweets about climate change on the prices of green and brown stocks persist or revert over time? To answer this question, we report in Columns (1)–(6) of Table 1 the results from estimating the panel regression in Equation (1) for event horizons from one minute before the tweet to $s = 10, 15, 30, 60, 90$ minutes and one day after each tweet, respectively.[15] We must note, however, that a longer event horizon weakens the high-frequency identification because of an increased likelihood of the estimates being contaminated by other events resulting in price movements.

As we extend the event horizon, we find that the spread between the returns of green and brown stocks remains significant at a 1% level, and its economic magnitude increases substantially. For instance, when the event horizon is extended from ten to ninety minutes after the tweet, the spread associated with one standard deviation increase in the sentiment toward the green transition increases 2.3 times from 0.137 basis points ($0.206 \times (0.260 + 0.402)$ bps) to 0.319 basis points (0.206×1.549 bps). When the event horizon is extended to one day after the tweet, the coefficient associated with the sentiment measure becomes insignificant. However, the spread between the returns of green versus brown stocks remains significant at a 1% level and increases more than tenfold. This finding suggests that it takes at least one day for market prices to incorporate the information contained in the tweets.[16]

[15] We conserve the length of the search windows, while varying the length of the event horizon.

[16] Note that the number of observations decreases as we expand the event horizon from ten to ninety minutes, that is, from 12.13 million in Column (1) to 10.02 million in Column (5). This is because, for an event horizon ending at $T(k) + x$ minutes, a tweet must be published at least $x + 15$ minutes before the market closes on that day (that is, 16:00 EST) in order for us to query the opening price of the first minute recorded during the fifteen-minute search window after the tweet. Similarly, for event horizons starting at $T(k) - 1$ minute, tweets must be posted at least sixteen minutes after the market opens at 09:30 EST in order for us to query the closing price of the last minute recorded in a 15-minute search window. Hence, we drop entries with missing price information for which the search windows fall outside trading hours.

3.4 Economic significance

To further gauge the economic relevance of our findings, Figure 2 depicts the cumulative daily-average return of a value-weighted green-minus-brown portfolio around pro-transition tweets (purple line) and around anti-transition tweets (yellow line). We first compute value-weighted green-minus-brown portfolio returns in the one-to-ten-minute event horizon around each tweet using the market capitalization of firms at the end of the previous day.[17] We then identify pro- and anti-transition tweets using the categorical sentiment measure obtained as described in Section 2.3. Because the number of pro-transition tweets in our sample is 3.54 times greater than that of anti-transition tweets, simply accumulating returns around tweets falling under each category would result in an unfair comparison. Thus, on each day, we first compute the *average* of the one-to-ten-minute green-minus-brown portfolio returns around pro-transition tweets and around anti-transition tweets, and then accumulate these average daily high-frequency returns. We exclude days on which we do not have *both* pro-transition and anti-transition tweets, resulting in a time series of 2,270 trading days.

Figure 2 shows that the cumulative one-to-ten minute green-minus-brown portfolio returns around pro-transition tweets is about 1.5% and around anti-transition tweets is about −5.4%. This 6.9% spread is substantial from an economic perspective because we effectively accumulate the average return around only a *single* tweet on each of the 2,270 trading days in our sample. To provide context for these figures, the cumulative one-to-ten minute small-minus-big portfolio returns around pro-transition and anti-transition tweets—computed in a manner analogous to the green-minus-brown portfolio returns—are 0.57% and 1.03%, respectively. Similarly, the cumulative one-to-ten minute high-minus-low book-to-market (value) portfolio returns around pro-transition and anti-transition tweets are 0.15% and 0.04%, respectively. The relative magnitude of our findings suggests that legislator tweets contain material information that foreshadows future environmental regulation, leading to a differential and economically meaningful impact on the prices of green versus brown stocks.

[17] We avoid using the market capitalization of firms at the end of the same day that a tweet is published because it is impacted by the price reaction of stocks to the tweets posted on that day. Nonetheless, using the contemporaneous market capitalization instead of its lagged counterpart does not alter our conclusions.

Table 2: Robustness checks for impact of legislator tweets on stock returns

Column (1)–(6) of this table report various robustness checks for the results of the panel regression in Equation (1) reported in Column (2) of Table 1. Column (1) reports the results obtained by using an event horizon starting *five* minutes before the tweet (instead of one minute), Column (2) using a contracted event horizon that ends *five* minutes after the tweet (instead of ten minutes), Column (3) using market-adjusted returns as the dependent variable, computed by subtracting the high-frequency one-to-ten-minute returns of the S&P 500 index from those of the stocks, Column (4) using day-level fixed effects instead of quarter-level fixed effects, Column (5) categorizing green and brown firms using carbon intensity measures obtained from S&P Trucost (instead of Refinitiv Eikon), and Column (6) categorizing green and brown firms using carbon intensity measures that include Scope 3 emissions (instead of only Scope 1 and 2 emissions). We report all intercept and slope coefficients in basis points. Standard errors are robust to heteroscedasticity and clustered at the firm level. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) stars, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta p(-05,+10)$	$\Delta p(-01,+05)$	Excess Rets	Day Level FE	S&P Trucost	S3 Emission
Brown	0.127	0.124*	0.181*	0.142	0.119	0.0814
	(1.07)	(1.79)	(1.79)	(1.45)	(1.27)	(0.75)
Green	-0.297**	-0.0725	-0.299***	-0.313***	-0.153	-0.427***
	(-2.36)	(-0.91)	(-2.81)	(-2.98)	(-1.49)	(-3.58)
Sentiment	-1.353***	-0.355***	0.157**	-0.260***	-0.538***	-0.781***
	(-16.04)	(-6.51)	(2.23)	(-3.70)	(-8.33)	(-8.02)
Brown × Sentiment	-0.276**	-0.179**	-0.310***	-0.261***	-0.226**	-0.110
	(-2.37)	(-2.38)	(-3.15)	(-2.69)	(-2.34)	(-0.86)
Green × Sentiment	0.408***	0.131*	0.395***	0.401***	0.254***	0.596***
	(3.57)	(1.83)	(4.05)	(4.19)	(2.69)	(4.96)
Constant	1.085***	0.263***	-0.133**	0.217***	0.404***	0.591***
	(14.19)	(5.36)	(-2.04)	(3.38)	(7.04)	(7.24)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Legislator FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R2	0.00704	0.00649	0.00105	0.0348	0.00696	0.00729
Observations	12093244	12171876	12131054	12131416	11440155	7509912

3.5 Robustness checks

Table 2 shows that the results of the panel regression in Equation (1) reported in Column (1) of Table 1 are robust to using an event horizon starting *five* minutes before the tweet (instead of one minute), using a contracted event horizon that ends *five* minutes after the tweet (instead of ten minutes), using market-adjusted returns as the dependent variable, computed by subtracting the high-frequency one-to-ten-minute returns of the S&P 500 index from those of the stocks, using day-level fixed effects instead of quarter-level fixed effects, categorizing green and brown firms using carbon intensity measures obtained from S&P Trucost (instead

of Refinitiv Eikon), categorizing green and brown firms using carbon intensity measures that include Scope 3 emissions (instead of only Scope 1 and 2 emissions).

We now provide a detailed discussion of these robustness checks. We focus our discussion on the slopes for the interaction between tweet sentiment and green and brown indicator variables ($\hat{\tau_G}$ and $\hat{\tau_B}$), which characterize the differential impact of legislator tweets on the returns of green and brown stocks. Column (1) of Table 2 reports the results when we obtain stock returns using a *five*-to-ten-minute horizon. Although we expect tweet timestamps to be accurate, modifying the event horizon to begin five minutes (instead of one minute) before the tweet buffers against small discrepancies in the recorded timestamps. The coefficients on the interactions between sentiment and the brown and green indicators remain significant at the 5% and 1% levels, respectively, and comparable in magnitude to those obtained using one-to-ten-minute returns. Column (2) reports the estimates for a contracted event horizon that ends five minutes (instead of ten minutes) after the tweet. Owing to the tighter event horizon, $\hat{\tau_B}$ and $\hat{\tau_G}$ diminish in magnitude, but they remain significant at the 5% and 10% levels, respectively.

Column (3) reports the estimates obtained by using market-adjusted returns as the target variable by subtracting the one-to-ten-minute returns of the S&P 500 index from those of the stocks. The high-frequency returns of the market index are obtained from FirstRate analogously to those of stocks, as described in Section 3.1. The differential impact of sentiment on the returns of green versus brown stocks remains almost unchanged in comparison to our baseline results reported in Column (1) of Table 1. However, the estimate associated with the sentiment measure now becomes positive at the 5% significance level. The explanation for this is that a higher sentiment toward green transition is associated with lower market returns, subtracting negative market returns from stock returns has a positive net impact on market-adjusted returns when sentiment is high. Column (4) shows that our estimates remain robust to the introduction of day-level fixed effects, although the coefficient $\hat{\tau_B}$ is now significant only at the 5% level. This finding suggests that our results are not driven by concurrent unobserved events that simultaneously affect both legislator tweets and stock returns. Column (5) reports estimates obtained using the carbon intensity measure obtained

from S&P Trucost, an alternative data provider widely used in the literature. Our findings remain virtually unchanged.

Finally, Column (6) reports estimates obtained by including Scope 3 emissions to construct the carbon emissions intensity used to sort stocks into green and brown portfolios. Our main findings are robust to including Scope 3 emissions: the slope for sentiment is significantly negative and, more importantly, the slope for the interaction between sentiment and the green indicator variable is large and significant at the 1% confidence level. However, we observe that the interaction term between the sentiment measure and the brown indicator variable is insignificant. This may occur because when we include Scope 3 emissions to compute carbon intensity, we lose a disproportionate number of observations for brown firms. For instance, the median carbon intensity measure—computed using Scope 1 and 2 emissions only—of firms that report Scope 3 emissions is 32.04 tCO2e/$M, whereas that for firms that do not report Scope 3 emissions is 55.45 tCO2e/$M (42% higher). Thus, several gray firms are misclassified as brown when using the alternative intensity measure in Column (6).

3.6 Stock returns and decarbonization

Certain high-emission firms are likely to be more able or willing to adapt their carbon emissions to transition regulation than others. For instance, certain firms may rely on easily adaptable production processes, their institutional shareholders may be supportive of the green transition, or may own technology that can be altered to reduce emissions at a reasonable cost. If markets can predict which firms will adapt their carbon emissions in response to transition regulation, then one would expect that firm returns around pro-transition tweets should be informative about future carbon emissions. To test this hypothesis, we study whether firm stock returns around pro-transition tweets predict the future carbon emissions intensity of low and high emission firms.

Let C_{it} denote the carbon emissions intensity of the ith firm reported at disclosure date t. We compute the change in carbon emissions intensity between subsequent disclosures as:

$$CarbonDelta_{i,t+1} = C_{i,t+1} - C_{it}. \tag{2}$$

Carbon emissions are reported at an annual frequency, but the reporting cycles are heterogeneous across firms, and thus, the disclosure date t for C_{it} is specific to firm i. For each disclosure C_{it}, we aggregate the one-to-ten-minute returns of firm i around the pro-transition tweets that were posted during the year preceding date t:

$$AggReturn_{it} = \sum_{k \in \mathcal{K}_{it}} \Delta p_{ik}(-1, +10), \tag{3}$$

where $\mathcal{K}_{it}$ contains all tweets with timestamp in the year preceding date t. In each calendar year, we sort firms by their carbon emissions intensity and categorize firms that fall below and above the median as low- and high-emission firms. We then estimate the following panel regression separately for low- and high-emission firms:

$$CarbonDelta_{i,t+1} = \gamma\, AggReturn_{it} + \theta_i + \theta_y + \varepsilon_{it}, \tag{4}$$

where α and γ are the intercept and slope, θ_i are firm fixed effects, θ_y are time fixed effects for the calendar year corresponding to the disclosure date t, and ε_{it} is the error term.

Panels (a) and (b) of Table 3 report the results from estimating Equation (4) separately for low- and high-emission firms. Column (1) reports results for a specification without fixed effects. Columns (2) and (3) report results obtained when firm and year fixed effects are included, respectively. Column (4) reports results obtained when both firm and year fixed effects are included. We observe that the results for the slope coefficient γ are consistent across the four columns.

For instance, Column (4) of Panel (a) shows that price movements around pro-transition legislator tweets do not predict subsequent changes in the carbon emissions intensity of low-emission firms. The coefficient estimate associated with *AggReturn* is statistically and economically insignificant. This is consistent with low-emission firms not responding to future climate regulation by further reducing their emissions, perhaps because they are already carbon efficient. In contrast, Column (4) of Panel (b) shows that a one standard deviation increase in the *AggReturn* measure (that is, 20.81%) of high-emission firms is

Table 3: Legislator tweets and firm decarbonization

Panels (a) and (b) report the results from estimating Equation (4) separately for low- and high-emission firms. The dependent variable is *CarbonDelta* obtained from Equation (2). The independent variable, *AggReturn*, is obtained from Equation (3). Column (1) reports results for a specification without any fixed effects. Columns (2) and (3) report results obtained when firm- and year-fixed effects are included, respectively. Column (4) reports results obtained when both firm- and year-fixed effects are included. Standard errors heteroscedasticity robust and clustered at the firm level. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) starts, respectively.

Panel (a): Low-emission firms				
	(1)	(2)	(3)	(4)
AggReturn	0.000130	0.000131	0.0000761	0.0000983
	(1.02)	(1.04)	(0.59)	(0.73)
Constant	-0.653***	-0.653***	-0.762***	-0.761***
	(-4.32)	(-4.31)	(-259.56)	(-250.29)
Firm FE	No	No	Yes	Yes
Time FE	No	Yes	No	Yes
Clustering	Yes	Yes	Yes	Yes
Adjusted R2	0.000543	0.000713	0.0816	0.0837
Observations	1099	1099	1029	1029

Panel (b): High-emission firms				
	(1)	(2)	(3)	(4)
AggReturn	-0.0134**	-0.0125**	-0.0188**	-0.0183**
	(-2.54)	(-2.41)	(-2.31)	(-2.29)
Constant	-34.91***	-34.93***	-33.50***	-33.51***
	(-4.67)	(-4.66)	(-291.52)	(-297.33)
Firm FE	No	No	Yes	Yes
Time FE	No	Yes	No	Yes
Clustering	Yes	Yes	Yes	Yes
Adjusted R2	0.00685	0.0299	0.0139	0.0462
Observations	1067	1067	1009	1009

associated with a decrease in subsequent carbon emissions intensity by 38.08 tCO2/$M, significant at the 5% level. Thus, high-emission firms that decarbonize more in the future are associated with higher returns around pro-transition legislator tweets. This finding reveals that the market correctly identifies and rewards high-emission firms that can better adapt to transition regulation, and thus suggests that the green transition impacts not only firm value, but also real outcomes such as carbon emissions. In addition, our results indicate that brown firms that are not able or willing to decarbonize shoulder a disproportionately large share of the burden imposed by transition regulations.

4 Mechanism

We now investigate the mechanisms through which legislator tweets impact the prices of green versus brown stocks. In Section 4.1, we examine the information channel by showing that a daily pro-transition sentiment index based on legislator tweets contains information that is distinct from that communicated by climate-related news in mainstream media. In Section 4.2, we study the regulatory channel showing that, for tweets that mention specific environmental regulations, the spread between the return of green and brown stocks increases in the weeks leading to a congressional vote and disappears afterward. Moreover, we find that the spread is larger before close votes and when Congress is split.

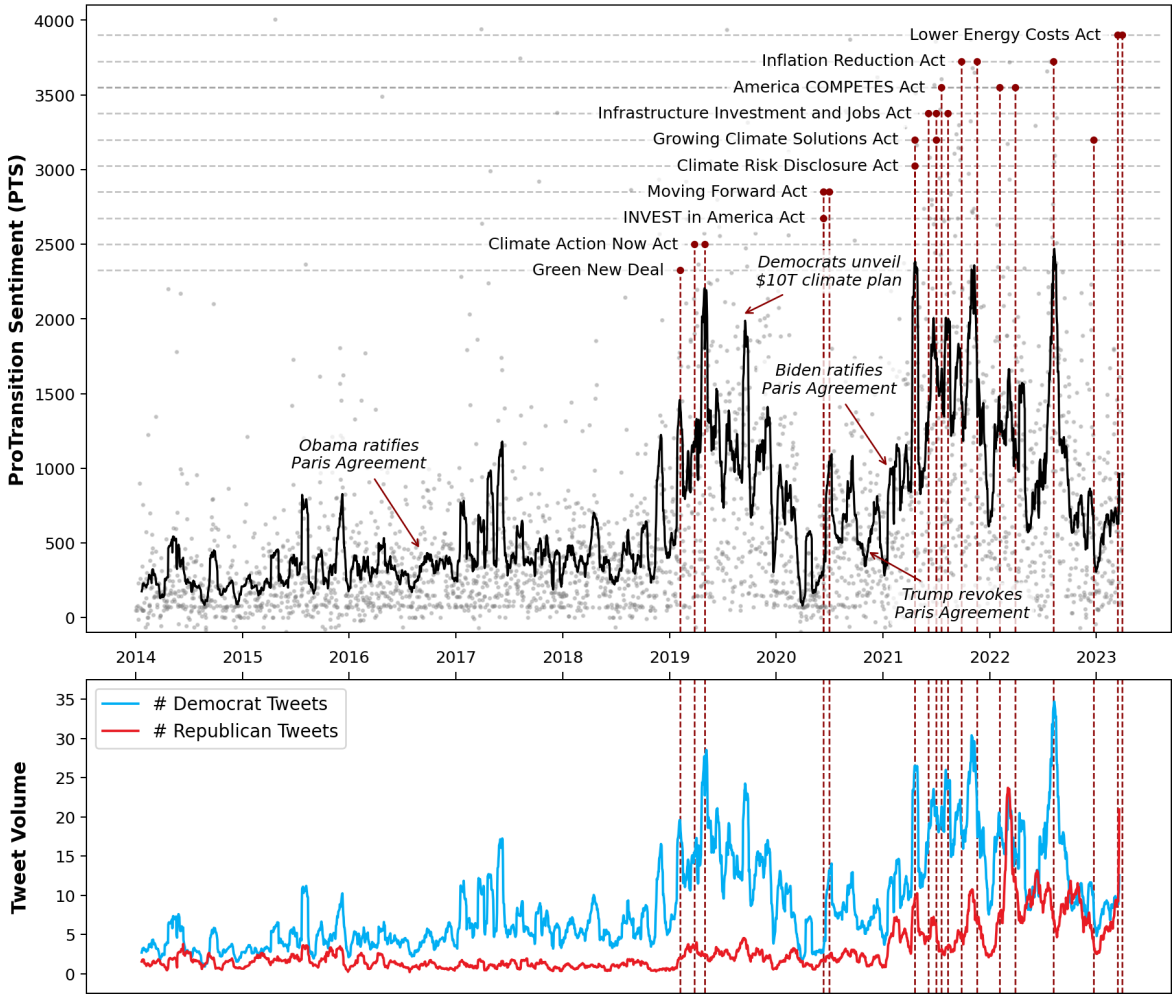
4.1 The information channel

In this section, we examine whether the information contained in legislator tweets about climate change is distinct from that in mainstream media news. To do this, we first construct a daily ProTransition Sentiment (PTS) index based on the legislator tweets, and then study whether there is a lead-lag relation between the Media Climate Change Concerns (MCCC) index of Ardia et al. (2023) and the PTS index.

We construct a daily PTS index by adding the numerical sentiment toward the green transition across all legislator tweets related to climate change posted during a given day. The top panel of Figure 6 depicts day-specific PTS values (light gray markers) and the twenty-day moving average of the PTS index (solid black line) shifted forward by ten days so that the time series is aligned with calendar events. The figure also annotates the ten environmental regulations introduced to Congress during our sample period that received most legislator tweets. For each regulation, the vertical dashed red lines—from left to right—demarcate the dates on which the regulation was (i) introduced, (ii) voted on in the sponsoring chamber, and (iii) voted on in the other chamber. For example, the 'Infrastructure Investment and Jobs Act' was introduced and voted on in the House of Representatives on June 4, 2021 and

Figure 6: ProTransition Sentiment (PTS) index

This figure shows the ProTransition Sentiment (PTS) index constructed by adding the numerical sentiment toward the green transition of all legislator tweets posted during a given day. The top panel depicts day-specific PTS values (light gray markers) and the twenty-day moving average of the PTS index (solid black line) shifted forward by ten days so that the time series is aligned with calendar events. The figure also annotates the ten environmental regulations introduced to Congress during our sample period that received most legislator tweets. For each regulation, the vertical dashed red lines—from left to right—demarcate the dates on which the regulation was (i) introduced, (ii) voted on in the sponsoring chamber, and (iii) voted on in the other chamber. Missing vertical lines imply that a congressional vote did not take place. The bottom panel shows the twenty-day moving average of the volume of Democrat (blue line) and Republican (red line) climate-change tweets posted each day. The moving averages are shifted forward by ten days to center the values. The vertical dashed red lines in the bottom panel correspond to those in the top panel.



July 1, 2021, respectively, and subsequently voted on in the Senate on August 10, 2021.[18] The absence of a second and third vertical line implies that the regulation was not put to

[18]Sometimes, regulations are introduced or voted on within days of each other. For example, the 'Climate Risk Disclosure Act' and the 'Growing Climate Solutions Act' are introduced five days apart on April 15, 2021 and April 20, 2021, respectively. In such instances, we overlap the vertical dashed red lines to improve visual clarity.

a vote in the sponsoring chamber and the other chamber, respectively. For example, the 'Climate Action Now Act' was introduced and voted on in the House of Representatives on March 27, 2019 and May 2, 2019, respectively. However, it was never put to a vote in the Senate. Reassuringly, the graph reveals consistent spikes in the PTS index around the days on which the regulations are introduced or voted on.

The bottom panel of Figure 6 shows the twenty-day moving average of the volume of Democrat (blue line) and Republican (red line) climate-change tweets posted each day. As in the top panel, the moving averages are shifted forward by ten days to center the values. The vertical dashed red lines in the bottom panel correspond to those in the top panel. We note a dramatic spike in the volume of Democratic tweets around the passage of the 'Inflation Reduction Act' in the Senate on August 7, 2022—a key regulation championed by the Biden administration. In contrast, we note a dramatic spike in the volume of Republican tweets around the passage of the 'Lower Energy Costs Act' on March 14, 2023—a regulation introduced by Republicans in strong opposition to the climate-related policies introduced during the Biden administration. In summary, Figure 6 shows that the PTS index provides a quantitative measure of the sentiment toward the green transition in the legislative discourse on climate-related topics, as evidenced by the interplay between the index and key regulatory milestones.

If legislator tweets contain information that is distinct from that in mainstream media news, we would expect current values of the PTS index to be able to forecast future values of the Media Climate Change Concerns (MCCC) index of Ardia et al. (2023).[19] To test this hypothesis, we study whether there is a lead-lag relation between the MCCC index of Ardia et al. (2023) and the PTS index. We standardize the PTS and MCCC indices so that their time series have zero mean and unit standard deviation.[20] To determine whether

[19] Ardia et al. (2023) construct the Media Climate Change Concerns (MCCC) index by using risk and sentiment lexicon to calculate the level of concern expressed in climate-related articles from twelve major U.S. news outlets. The article-specific measures are then carefully aggregated to produce a daily climate-change concerns index. The resultant index measures and combines the degree of risk and negativity expressed in mainstream media with respect to climate change.

[20] We verify that the time series are stationary using the augmented Dickey-Fuller test (ADF) implemented using `statsmodels.tsa.stattools.adfuller` package in Python with default settings. The null hypothesis is that there is a unit root, with the alternative that there is no unit root. The ADF statistic for PTS (MCCC) is -4.7 (-5.25), which rejects the null hypothesis. We further validate that each series is stationary by

there is a lead-lag relation between the two indices, that is, whether past values of one index can predict the other index, we follow DeMiguel, Nogales, and Uppal (2014, p. 1036) and estimate the following vector autoregressions:

$$\text{MCCC}_t = \alpha_1 + \beta_{11}\text{MCCC}_{t-1} + \beta_{12}\text{PTS}_{t-1} + \varepsilon_{1t}, \tag{5a}$$

$$\text{PTS}_t = \alpha_2 + \beta_{21}\text{MCCC}_{t-1} + \beta_{22}\text{PTS}_{t-1} + \varepsilon_{2t}, \tag{5b}$$

where MCCC_t and PTS_t are the indices on day t, α_1 and α_2 are the intercepts, β_{ij} for $i, j \in \{1, 2\}$ are the slopes for the lagged indices, and ε_1 and ε_2 are the error terms. If β_{12} is significant, we infer that the PTS index leads the MCCC index and, if β_{21} is significant, then the MCCC index leads the PTS index.

Columns (2) and (4) of Table 4 report the results obtained from estimating the vector autoregressions in Equations (5a) and (5b). We find that both $\hat{\beta}_{12}$ and $\hat{\beta}_{21}$ are significant, that is, both the MCCC and PTS indices lead each other. This demonstrates that neither index fully subsumes the information in the other and each index captures unique aspects of the green transition.

To further gauge the importance of each index in predicting the other, we also estimate autoregressions for each of the two indices separately, that is:

$$\text{MCCC}_t = \alpha_1 + \beta_1\text{MCCC}_{t-1} + \varepsilon_{1t}, \tag{6a}$$

$$\text{PTS}_t = \alpha_2 + \beta_2\text{PTS}_{t-1} + \varepsilon_{2t}. \tag{6b}$$

Columns (1) and (3) of Table 4 report the results obtained from estimating the separate autoregressions in Equations (6a) and (6b). Comparing Columns (1) and (2) in Table 4, we observe that, when we predict the MCCC index, introducing the lagged PTS index increases the adjusted R-squared by 2.7%. Similarly, comparing Columns (3) and (4) in Table 4, we find that, when we predict the PTS index, introducing the lagged MCCC index increases the adjusted R-squared by 2%. Thus, each index helps to predict the other even controlling for lagged values of the target index.

estimating the following equation: $\Delta y_t = \alpha + \beta t + \theta y_{t-1} + \delta_1 \Delta y_{t-1} + \delta_2 \Delta y_{t-2} + \epsilon_t$, where y_t represents the value of the series at time t, $\Delta y_t := y_t - y_{t-1}$, α captures drift, β is the parameter associated with time index t to capture trends, and ϵ_t is the error term. We observe that $\hat{\theta}$ is negative and statistically significant, indicating that the time series is stationary.

Table 4: Lead-lag relation between PTS and MCCC indices

This table reports the results from estimating vector autoregressive models for our ProTransition Sentiment (PTS) index and the Media Climate Change Concerns (MCCC) index of Ardia et al. (2023). Columns (2) and (4) report the results obtained from estimating the vector autoregressive model for the two indices formulated in Equations (5a) and (5b). Columns (1) and (3) report the results obtained from estimating the separate autoregressive models for the PTS and MCCC indices formulated in Equations (6a) and (6b). We construct a daily ProTransition Sentiment (PTS) index by aggregating the numerical sentiment toward the green transition of all legislator tweets posted during a given day. We standardize the PTS and MCCC indices so that their time series has zero mean and unit standard deviation. Standard errors are robust to heteroscedasticity. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) starts, respectively.

	(1)	(2)	(3)	(4)
	MCCC(t)	MCCC(t)	PTS(t)	PTS(t)
MCCC(t-1)	0.474***	0.410***		0.154***
	(26.41)	(21.40)		(7.24)
PTS(t-1)		0.176***	0.416***	0.362***
		(7.97)	(7.84)	(6.62)
Constant	0.0254	0.0211	0.0208	0.0164
	(1.54)	(1.30)	(1.25)	(1.00)
Adjusted R2	0.225	0.252	0.175	0.195
Observations	2799	2799	2799	2799

Our results suggest that both indices contain distinct information. While it is natural for legislators to discuss the green transition following heightened media concerns about climate change, the latter can also follow the former for several reasons. For instance, the news might echo the support expressed by legislators toward environmental regulation. Legislators would also be better positioned to anticipate future events such as environmental summits or international partnerships aimed toward mitigating climate change. Moreover, legislator tweets could highlight material information that the mainstream media has missed. Overall, we find that the legislator tweets contain distinct information and do not simply track that contained in mainstream news.

4.2 The regulatory channel

In this section, we focus on the regulatory channel by looking at the high-frequency returns of green versus brown stocks around legislator tweets that mention specific regulations introduced to Congress during our sample period.

4.2.1 Identifying regulatory tweets and measuring their sentiment

We combine ChatGPT with a dictionary approach to identify 40,821 tweets posted during NYSE trading hours that mention one of the 19 regulations about climate change listed in Table 5. These regulations are the most frequently mentioned regulations in legislator tweets among those introduced to Congress during our sample. Section IA.2.1 of the Internet Appendix gives a detailed description of the approach we use to identify these tweets.

For each tweet identified as mentioning a specific regulation, we measure legislator support (LSupport) toward the regulation using the following ChatGPT prompt: "*The following tweet was posted by a member of the United States Congress on* {`postdate`} *and mentions the following regulation:* {`bill`} *with code* {`code`}*. Please output a single continuous value between 0 and 1 that represents the likelihood that the tweet supports the bill. A value of 0 indicates strong opposition, a value of 1 indicates strong support and a value of 0.5 indicates neutrality. In case the tweet does not mention or is not related to the regulation, please output a value of 0.5,*" where {`postdate`} is the date on which the tweet was posted, {`bill`} the name the regulation mentioned in the tweet, and {`code`} its bill number.

We also ascertain whether each regulation is pro- or anti-transition, using the following ChatGPT prompt: "*Does the following regulation considered by the United States Congress support or oppose the green transition? The 'green transition' refers to the comprehensive shift from traditional, carbon-intensive energy sources and practices to sustainable, low-carbon alternatives. This encompasses adopting renewable energy technologies, enhancing energy efficiency, promoting sustainable agriculture, and implementing policies that reduce greenhouse gas emissions to mitigate climate change. The United Nations describes it as a fundamental, system-wide reorganization across technological, economic and social factors, including paradigms, goals, and values.*" We find that 18 out of the 19 regulations support the green transition and only one ("Lower Energy Costs Act," H.R.1) opposes it.

Table 5: Regulations discussed in environmental tweets

This table enumerates the 19 regulations with the largest number of mentions in our sample of legislator tweets about climate change. Regulations are listed in descending order of number of mentions. The first column titled 'Regulations' reports the name of the regulation. The second column titled 'Code' reports the corresponding bill number. Bill numbers starting with "H.R." are initiated (introduced) by members of the the House of Representatives (lower chamber). Bill numbers started with "S." are initiated by the members of the Senate (upper chamber). The third column titled 'Congress' reports the sequential number of the two-year term of the legislative body. The fourth column titled 'Introduction' reports the date on which the regulation was introduced in the sponsoring house. The fifth and sixth columns titled 'House' and 'Senate' report the dates on which the regulation was deliberated and voted on in the lower and upper chamber, respectively. Missing dates imply that the bill has either not been deliberated or voted upon in a given chamber, or that it failed to pass preceding vote in the other chamber. Regulations ratified by both chambers are sent to the President, who signs it into law. The dates on which the regulations are signed into law are reported by the last column titled 'President'. *S.1251 was passed as part of the Consolidated Appropriations Act 2023 and S.2754 was passed as part of the Consolidated Appropriations Act 2021.

Regulation	**Code**	**Congress**	**Introduction**	**House**	**Senate**	**President**
Inflation Reduction Act (Build Back Better Act)	H.R.5376	117th	September 27, 2021	November 19, 2021	August 7, 2022	August 16, 2022
Infrastructure Investment and Jobs Act (Bipartisan Infrastructure Law)	H.R.3684	117th	June 4, 2021	July 1, 2021	August 10, 2021	November 15, 2021
Climate Action Now Act	H.R.9	116th	March 27, 2019	May 2, 2019		
Lower Energy Costs Act	H.R.1	118th	March 14, 2023	March 30, 2023		
Green New Deal	H.Res.109	116th	February 7, 2019			
Growing Climate Solutions Act*	S.1251	117th	April 20, 2021	December 23, 2022	June 24, 2021	December 29, 2022
Climate Risk Disclosure Act	H.R.2570	117th	April 15, 2021			
America COMPETES Act	H.R.4521	117th	July 19, 2021	February 4, 2022	March 28, 2022	
Moving Forward Act	H.R.2	116th	June 11, 2020	July 1, 2020		
INVEST in America Act	H.R.7095	116th	June 4, 2020			
American Innovation and Manufacturing Act*	S.2754	116th	October 30, 2019	December 21, 2020	December 21, 2020	December 27, 2020
CLEAN Future Act	H.R.1512	117th	March 2, 2021			
Growing Renewable Energy and Efficiency Now Act	H.R.848	117th	February 4, 2021			
Trillion Trees Act	H.R.2639	117th	April 19, 2021			
Energy Policy Modernization Act	S.2012	114th	September 9, 2015	May 25, 2016	April 20, 2016	
CHIPS and Science Act	H.R.4346	117th	July 1, 2021	July 28, 2021	June 22, 2022	August 9, 2022
100% Clean Economy Act	H.R.5221	116th	November 21, 2019			
Energy Innovation and Carbon Dividend Act	H.R.763	116th	January 24, 2019			
USE IT Act	S.383	116th	February 7, 2019	December 21, 2020	December 21, 2020	December 27, 2020

In Section IA.2.2 of the Internet Appendix, we validate our ChatGPT-based measure of regulation support toward the green transition by comparing it with that implied by the classification of the League of Conservation Voters (LCV), used by Lueck et al. (2022), and find that the two approaches yield similar results.

Finally, we define pro-transition legislative support (PTLSupport) as the support of a legislator tweet toward pro-transition regulation. PTLSupport is equal to the continuous legislator support (LSupport) when the underlying regulation is pro-transition, and one minus the legislator support (Lsupport) otherwise.[21] Figure IA.2 in the Internet Appendix shows the distribution of PTLSupport.

4.2.2 Analyzing the impact of regulatory tweets on stock prices

Armed with a sample of regulatory legislator tweets and a measure of their pro-transition legislative support (PTLSupport), we are ready to study the impact of regulatory-tweet sentiment on stock returns. Column (1) of Table 6 reports the results from estimating panel regression (1) over the 40,821 regulatory tweets after replacing the sentiment measure with PTLSupport. The third row of Column (1) shows that one standard deviation (0.33) increase in the legislator support toward pro-transition regulation leads to a 0.161 bps ($= 0.33 \times -0.489$) reduction in the returns of gray stocks around the tweet, significant at the 1% level. However, the fifth row shows that the reduction in the returns of green stocks is 0.073 bps ($= 0.33 \times 0.220$) lower than that of gray firms, significant at the 1% level. Thus, consistent with our findings based on legislator tweets about climate change, using our sample of tweets that mention environmental regulations we find that legislator-tweet support toward environmental regulations has a differential impact on the returns of green versus brown firms. Moreover, this impact depends on whether the regulation is pro- or anti-transition.

If regulatory tweets affect stock prices because they contain information about the probability that regulation passes a congressional vote, then we would expect their impact to be

[21] If a tweet mentions multiple regulations, we query ChatGPT to obtain the legislator support toward each regulation separately. We then compute PTLSupport corresponding to each regulation and take the average value for that tweet.

Table 6: Impact of legislator support toward pro-transition regulation

This table reports the results from estimating the panel regression in Equation (1) on the sample of legislator tweets that mention environmental regulations obtained as described in Section 4.2. The 'Sentiment' measure in the original specification is replaced by pro-transition legislative support (PTLSupport). PTLSupport measures the stance of a tweet toward pro-transition regulation between zero (strong opposition) and one (strong support). The dependent variable is the one-to-ten minute return stock return around each tweet, $\Delta p(-01, +10)$. Stocks that fall into the top and bottom tercile of carbon emissions intensity—obtained by using the most recent Scope 1 and 2 emissions data available on Refinitiv Eikon at the time each tweet was posted—are classified as green and brown, with the rest classified as gray. Column (1) reports the estimates obtained over the full sample of 40,821 tweets with direct regulatory mentions. Columns (2)–(6) report the results from estimating the panel regression for tweets posted within the four, three, two, and one weeks leading up to a congressional vote by members of the house in which the vote took place, and one week after the vote, respectively. Standard errors are robust to heteroscedasticity and clustered at the firm level. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) starts, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Full sample	Leading 4wks	Leading 3wks	Leading 2wks	Leading 1wk	Lagging 1wk
Brown	-0.228	0.228	0.173	0.149	0.0503	-0.322
	(-1.55)	(0.83)	(0.57)	(0.44)	(0.11)	(-0.41)
Green	-0.0232	0.454*	0.472*	0.406	0.111	-0.173
	(-0.19)	(1.76)	(1.82)	(1.36)	(0.27)	(-0.29)
PTLSupport	-0.489***	-1.410***	-1.976***	-2.833***	-1.984***	0.538
	(-8.71)	(-10.81)	(-12.74)	(-15.34)	(-8.32)	(1.42)
Brown × PTLSupport	-0.0683	-0.614***	-0.511***	-0.687***	-0.722***	-0.258
	(-0.86)	(-3.94)	(-2.94)	(-3.76)	(-3.06)	(-0.61)
Green × PTLSupport	0.220***	-0.168	-0.223	-0.247	-0.0273	0.441
	(3.00)	(-1.02)	(-1.19)	(-1.29)	(-0.11)	(1.06)
Constant	0.274***	1.009***	1.435***	2.108***	1.820***	-1.503***
	(3.94)	(7.16)	(9.26)	(10.87)	(7.16)	(-3.41)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Legislator FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R2	0.00685	0.0249	0.0304	0.0380	0.0564	0.0866
Observations	9475082	2347107	1984091	1502572	928762	434399

stronger for tweets posted immediately before the vote. To test this hypothesis, Columns (2)–(6) of Table 6 report the results from estimating the panel regression for tweets posted within the four, three, two, and one weeks leading up to a congressional vote by members of the house in which the vote took place, and one week after the vote. The third row of Column (2) shows that one standard deviation increase in the legislator support toward pro-transition regulation leads to a 0.465 bps ($= 0.33 \times -1.410$) reduction in gray stock returns around tweets posted within the four weeks leading to a congressional vote, significant at

the 1% level. The fourth row shows that the reduction in brown stock returns is 0.203 bps ($= 0.33 \times -0.614$) higher than that of gray stocks, significant at the 1% level. Consistent with the hypothesis that regulatory tweets contain information about the probability that a regulation will pass a congressional vote, Column (5) of Table 6 shows that the reduction in gray stock returns increases when we focus on regulatory tweets posted in the week preceding the vote (compared to four weeks) from 0.465 bps to 0.655 bps ($= 0.33 \times -1.984$). Similarly, the incremental reduction in the expected returns of brown stocks appreciates from 0.203 bps to 0.238 bps ($= 0.33 \times -0.722$).[22] These estimates are statistically significant at the 1% level. Lastly, Column (6) shows that the impact of PTLSupport on stock prices becomes insignificant during the week following the congressional vote.

If regulatory tweets contain information about the probability that regulation passes a congressional vote, then we would also expect tweets posted before the vote to have a larger impact on stock prices when the uncertainty about the vote outcome is higher. To test this hypothesis, we split our sample between votes that eventually passed by a *close* margin and those that passed by a *clear* margin. Columns (1) and (2) of Table 7 report the results from estimating the panel regression for tweets posted within the two and one weeks preceding a close vote and Columns (3) and (4) for tweets posted within the two and one weeks preceding a clear vote. We compute the voting margin by dividing the difference between the *Yay* (for) and *Nay* (against) votes by the total number of votes cast for a given regulation. A congressional vote is considered close when the voting margin is less than 3%, and clear otherwise.[23] Consistent with the hypothesis that regulatory tweets contain information about the probability that a regulation passes a congressional vote, Table 7 shows that the impact of regulatory tweets on stock returns is substantially larger in weeks leading to a close vote, compared to weeks leading to a clear vote. For instance, the third rows of Columns (1) and

[22] Overall, we observe that an increase in PTLSupport leads to an increase in the prices of green firms, relative to gray and brown firms, when all tweets with legislative mentions are considered. However, when we focus on tweets posted during the weeks leading up to congressional votes, an increase in PTLS support leads to a decrease in the prices of brown firms, relative to gray and green firms. Albeit complementary, the results are consistent with one another, and with our baseline analysis in Section 3.2.

[23] We classify votes into close or clear based on a 3% threshold because it results in a reasonable number of observations in the subsample associated with "close" votes. However, Table IA.9 in the Internet Appendix shows that our findings are robust to using a 5% threshold to categorize close votes. Also, we do not decrease the threshold beyond 3% because it would result in all votes in the upper chamber being classified as "clear".

Table 7: Voting margin and legislator support toward pro-transition regulation

This table reports the impact of legislator support toward pro-transition regulation (PTLSupport) on green versus brown stock prices in the weeks preceding congressional votes based on the tightness of the voting margin. We compute the voting margin by dividing the difference between the *Yay* (for) and *Nay* (against) votes by the total number of votes cast for a given regulation. A congressional vote is considered "close" when the voting margin is less than 3%, and "clear" otherwise. Columns (1) and (2) reproduce Columns (4) and (5) of Table 6 for tweets posted in the weeks preceding congressional votes that were close. Columns (3) and (4) reproduce Columns (4) and (5) of Table 6 for tweets posted in the weeks preceding congressional votes that were clear. The dependent variable is the one-to-ten minute return stock return around each tweet, $\Delta p(-01, +10)$. Stocks that fall into the top and bottom tercile of carbon emissions intensity—obtained by using the most recent Scope 1 and 2 emissions data available on Refinitiv Eikon at the time each tweet was posted—are classified as green and brown, with the rest classified as gray. Standard errors are robust to heteroscedasticity and clustered at the firm level. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) stars, respectively.

	(1) CLOSE (leading 2wks)	(2) CLOSE (leading 1wk)	(3) CLEAR (leading 2wks)	(4) CLEAR (leading 1wk)
Brown	0.625	1.285	-0.317	-0.173
	(0.93)	(1.32)	(-0.96)	(-0.40)
Green	-0.220	-0.910	0.918**	0.842
	(-0.33)	(-1.18)	(2.51)	(1.62)
PTLSupport	-7.674***	-5.269***	-0.0419	-3.109***
	(-21.48)	(-10.42)	(-0.20)	(-9.97)
Brown × PTLSupport	-0.924***	-1.554***	-0.523**	-0.605**
	(-3.32)	(-4.14)	(-2.36)	(-2.21)
Green × PTLSupport	-0.256	1.061***	0.0115	-0.410
	(-0.90)	(2.64)	(0.05)	(-1.46)
Constant	5.767***	3.959***	1.286***	3.037***
	(13.96)	(7.04)	(6.49)	(11.10)
Firm FE	Yes	Yes	Yes	Yes
Legislator FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Adjusted R2	0.0470	0.0646	0.0880	0.108
Observations	973296	581436	471423	315152

(2) show that one standard deviation increase in the legislator support toward pro-transition regulation leads to a 2.532 bps ($= 0.33 \times -7.674$) and 1.739 bps ($= 0.33 \times -5.269$) reduction in gray stock returns around tweets posted within the two and one weeks leading to a *close* vote, whereas the third rows of Columns (3) and (4) show it leads to a reduction of only 0.014 bps ($= 0.33 \times -0.04190$) and 1.026 bps ($= 0.33 \times -3.109$) around tweets posted within the two and one weeks leading to a *clear* vote. Similarly, the fourth rows of Columns (1) and (2) show that the reduction in brown stock returns is 0.305 bps ($= 0.33 \times -0.924$) and 0.513 bps ($= 0.33 \times -1.554$) higher than that of gray stocks around tweets posted in the two

Table 8: Congressional control and legislator support toward pro-transition regulation

This table reports the impact of legislator support toward pro-transition regulation (PTLSupport) on green versus brown stock prices depending on whether the Congress is unified or split. We define the Congress as being unified when both chambers—the House of Representatives and the Senate—are controlled by the same party. We define the Congress as being split when the two chambers are controlled by different parties. Columns (1) and (2) reproduce Column (1) of Table 6 for the sample of regulatory tweets posted when the Congress is unified and split, respectively. The dependent variable is the one-to-ten minute return stock return around each tweet, $\Delta p(-01, +10)$. Stocks that fall into the top and bottom tercile of carbon emissions intensity—obtained by using the most recent Scope 1 and 2 emissions data available on Refinitiv Eikon at the time each tweet was posted—are classified as green and brown, with the rest classified as gray. Standard errors are robust to heteroscedasticity and clustered at the firm level. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) stars, respectively.

	(1) Unified Congress	(2) Split Congress
Brown	-0.338*	-0.184
	(-1.75)	(-0.76)
Green	0.282*	-0.317
	(1.80)	(-1.30)
PTLSupport	-0.430***	-0.721***
	(-6.71)	(-6.72)
Brown × PTLSupport	-0.165*	-0.0150
	(-1.83)	(-0.10)
Green × PTLSupport	0.0429	0.421***
	(0.55)	(2.71)
Constant	0.289***	0.196
	(3.28)	(1.58)
Firm FE	Yes	Yes
Legislator FE	Yes	Yes
Time FE	Yes	Yes
Adjusted R2	0.00799	0.0361
Observations	7736836	1738237

and one weeks leading to a close vote, but only 0.173 bps ($= 0.33 \times -0.523$) and 0.200 bps ($= 0.33 \times -0.605$) around tweets posted in the two and one weeks leading to a clear vote.[24]

For the same reason, if regulatory tweets contain information about the probability that regulation passes a congressional vote, then we would expect that the impact of the regulatory tweets on stock prices would be less significant for tweets posted when Congress is unified (with the same party controlling both chambers), compared to tweets posted when

[24]When translating the regression coefficients in Table 7 to price movements, we use the standard deviation of PTLSupport over the full sample (0.33) to facilitate the comparison with the results from Table 6. Nevertheless, the standard deviation of PTLSupport across the unified and split Congress subsamples is very similar to that for the full sample.

Congress is split. To test this hypothesis, Columns (1) and (2) of Table 8 report the results from estimating the panel regression for tweets posted when Congress is unified and split, respectively. Consistent with the hypothesis, Table 8 shows that the impact of regulatory tweets on stock returns is substantially larger when Congress is split. For instance, comparing the third rows of Columns (1) and (2), we observe that while one standard deviation increase in the legislator support toward pro-transition regulation leads to only a 0.142 bps ($= 0.33 \times -0.430$) reduction in gray stock returns around tweets posted when Congress is unified, it leads to a 0.238 bps ($= 0.33 \times -0.721$) reduction around tweets posted when Congress is split. Similarly, comparing the fourth rows of Columns (1) and (2), we find that the return spread between green and brown stocks is only 0.069 ($= 0.33 \times (0.0429 + 0.165)$) around tweets posted while Congress is unified and 0.144 ($= 0.33 \times (0.015 + 0.421)$) around tweets posted when Congress is split. Moreover, while the impact of PTLSupport on the relative performance of green and brown stocks is statistically significant only at the 10% level when Congress is unified, it is significant at the 1% level when Congress is split.

By narrowing the focus to regulation-specific tweets, in this section we demonstrate that the impact of legislator tweets on stock returns is both more pronounced and statistically robust when tweets are more informative about the probability that regulation passes congressional votes. Taken together, the results in this section suggest that the green transition impacts the relative value of green and brown stocks, at least partly, via a regulatory channel.

5 Conclusion

We study the impact of green-transition regulation on firm value by comparing the returns of green and brown stocks in the minutes around tweets about climate change posted by members of U.S. Congress. We find that green stocks outperform brown stocks in the one-to-ten-minute window around pro-transition tweets, with the return spread statistically significant at the 1% level. The effect is also economically significant: accumulating the daily-average one-to-ten-minute return of a green-minus-brown portfolio around pro-transition tweets yields a return 6.9% higher than around anti-transition tweets. This spread is an order

of magnitude larger than that of small-minus-big size and high-minus-low value portfolios. We further link stock returns around legislator posts to the future carbon emissions of firms and find that investors correctly identify and reward high-emission firms that can decarbonize in response to transition-regulation news.

We construct a daily ProTransition Sentiment index which spikes around key legislative events and show that it contains information distinct from that communicated in mainstream media concerns related to climate change. To focus on the regulatory channel, we consider legislator tweets that mention environmental regulations, and find that the spread between the returns of green and brown stocks around tweets that support pro-transition regulations increases in the weeks leading to a congressional vote and vanishes afterward. Moreover, the spread is larger before close votes and when Congress is split. Our findings suggest that the green transition impacts the relative performance of green and brown stocks, at least partly, via a regulatory channel.

Our results have important implications for policymakers, firms, and investors. Policymakers can alleviate the impact of the green transition on brown firms by adopting a consistent approach to the green transition, thus avoiding abrupt shifts in policy that amplify the impact of green-transition regulation on brown firms. High-emission firms can add value by enhancing their ability to adjust their emissions in response to new transition regulation. Finally, investors can manage their exposure to green-transition regulation by holding stocks of green firms or brown firms with the ability to adapt their emissions.

References

V. V. Acharya, S. Giglio, S. Pastore, J. Stroebel, Z. Tan, and T. Yong. Climate transition risks and the energy sector. NYU Stern Working Paper, 2024.

J. M. Addoum and A. Kumar. Political sentiment and predictable returns. *The Review of Financial Studies*, 29(12):3471–3518, 2016.

D. Ardia, K. Bluteau, K. Boudt, and K. Inghelbrecht. Climate change concerns and the performance of green vs. brown stocks. *Management Science*, 69(12):7607–7632, 2023.

S. Baz, L. Cathcart, A. Michaelides, and Y. Zhang. Firm-level climate regulatory exposure. *Available at SSRN 3873886*, 2023.

B. S. Bernanke and K. N. Kuttner. What explains the stock market's reaction to federal reserve policy? *The Journal of Finance*, 60(3):1221–1257, 2005.

F. Bianchi, D. Comin, H. Kung, T. Kind, and A. Matusche. Slow recoveries through fiscal austerity: New insights in the effects of fiscal austerity. Technical report, ZEW Policy Brief, 2019.

F. Bianchi, R. Gómez-Cram, and H. Kung. Using social media to identify the effects of congressional viewpoints on asset prices. *The Review of Financial Studies*, 37(7):2244–2272, 2024.

P. Bolton and M. Kacperczyk. Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549, 2021.

P. Bolton and M. Kacperczyk. Global pricing of carbon-transition risk. *The Journal of Finance*, 78(6):3677–3754, 2023.

J. A. Cookson, C. Fox, J. Gil-Bazo, J. F. Imbet, and C. Schiller. Social media as a bank run catalyst. *Available on SSRN 4422754*, 2024.

T. Cowen and A. T. Tabarrok. How to learn and teach economics with large language models, including GPT. *Available on SSRN 4391863*, 2023.

V. DeMiguel, F. J. Nogales, and R. Uppal. Stock return serial dependence and out-of-sample portfolio performance. *The Review of Financial Studies*, 27(4):1031–1073, 2014.

M. Eskildsen, M. Ibert, T. I. Jensen, and L. H. Pedersen. In search of the true greenium. *Available at SSRN*, 2024.

R. Gómez-Cram and M. Grotteria. Real-time price discovery via verbal communication: Method and application to Fedspeak. *Journal of Financial Economics*, 143(3):993–1025, 2022.

A. Huang, H. Wang, and Y. Yang. FinBERT – A large language model approach to extracting information from financial text. 2020. Available on SSRN at https://papers.ssrn.com/sol3/papers.cfm?abstract/_id=3910214.

D. R. Känzig. The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125, 2021.

H. Ko and J. Lee. Can ChatGPT improve investment decisions? From a portfolio management perspective. *Finance Research Letters*, 64:105433, 2024.

J. F. Kölbel, M. Leippold, J. Rillaerts, and Q. Wang. Ask BERT: How regulatory disclosure of transition and physical climate risks affects the CDS term structure. *Journal of Financial Econometrics*, 22(1):30–69, 2024.

A. Korinek. Language models and cognitive automation for economic research. Technical report, National Bureau of Economic Research, 2023.

E. Kvam, P. Molnár, I. Wankel, and B. A. Ødegaard. Do sustainable company stock prices increase with ESG scrutiny? Evidence using social media. *Available at SSRN 4057988*, 2022.

Q. Li, H. Shan, Y. Tang, and V. Yao. Corporate climate risk: Measurements and responses. *The Review of Financial Studies*, 37(6):1778–1830, 2024.

A. Lopez-Lira and Y. Tang. Can ChatGPT forecast stock price movements? Return predictability and large language models. *arXiv preprint arXiv:2304.07619*, 2023.

T. Loughran and B. McDonald. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, 66(1):35–65, 2011.

D. Lueck, J. A. Ramos Pastrana, and G. Torrens. Interest groups, partisan politics, and environmental polarization. *Partisan Politics, and Environmental Polarization (June 16, 2022)*, 2022.

H. Mathurin. Pricing of green regulatory and technological risks. *Available at SSRN 5065450*, 2024.

E. Nakamura and J. Steinsson. High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330, 2018.

L. Pástor and P. Veronesi. Political cycles and stock returns. *Journal of Political Economy*, 128 (11):4011–4045, 2020.

L. Pástor, R. F. Stambaugh, and L. A. Taylor. Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2):550–571, 2021.

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

S. Ramelli, A. F. Wagner, R. J. Zeckhauser, and A. Ziegler. Investor rewards to climate responsibility: Stock-price responses to the opposite shocks of the 2016 and 2020 US elections. *The Review of Corporate Finance Studies*, 10(4):748–787, 2021.

P. Santa-Clara and R. Valkanov. The presidential puzzle: Political cycles and the stock market. *The Journal of Finance*, 58(5):1841–1872, 2003.

Z. Sautner, L. Van Lent, G. Vilkov, and R. Zhang. Firm-level climate change exposure. *The Journal of Finance*, 78(3):1449–1498, 2023.

Internet Appendix to

Legislator Tweets about the Green Transition and the Returns of Green versus Brown Stocks

This Internet Appendix reports the following robustness checks and additional results:

1. validating the ChatGPT-generated measures of sentiment, and
2. identifying regulatory mentions combining a dictionary approach and ChatGPT.

IA.1 Validating the ChatGPT-generated sentiment

In Section 2.3 of the main body of the manuscript, we obtain numerical and categorical tweet sentiment using ChatGPT-3.5. In this section, we validate these measures by comparing them to those manually generated by five researchers for a random sample of 100 tweets.

IA.1.1 Numerical sentiment

We ask five researchers to independently assign each tweet a continuous rating between −10 (strong *opposition* to the green transition) and 10 (strong *support* to the green transition); a value of 0 implies that the tweet is neutral, ambiguous, or unrelated to climate change. Panels (a) through (e) of Figure IA.3 depict the distribution of the sentiment measure assigned by the five researchers. Panel (f) of Figure IA.3 depicts the numerical sentiment measure (between −100 to 100) obtained from ChatGPT-3.5 in Section 2.3; the values are divided by ten to facilitate comparison with the manually generated measures. We observe that the sentiment scores assigned by researchers I, III, and IV are more polarized than those by ChatGPT, while the sentiment scores assigned by researchers II and V are less polarized. Thus, there is variation in how the support of a legislator tweet toward green transition is perceived.

Table IA.3 shows the matrix of correlations between the sentiment scores generatd by the five researchers and ChatGPT. Our main observation is that the sentiment scores generated by the five researchers are highly correlated with those generated by ChatGPT, with the correlation between every manually generated measure and the ChatGPT measure above 50%. Next, we compute the *average* sentiment produced by the five researchers for each tweet and compare it with the sentiment produced by ChatGPT. The first row of Table IA.4 shows that the correlation between the average sentiment produced by the five researchers and that by ChatGPT is 70%. Next, we compute the average sentiment produced by the research team after dropping one researcher at a time and report its correlation with the sentiment produced by ChatGPT. These correlation values are reported in the subsequent rows of Table IA.4 and range between 68% and 72%.

IA.1.2 Categorical sentiment

We replicate the analysis performed for the numerical sentiment by asking five researchers to independently assign one of the following categorical labels for each randomly selected tweet: Strongly ProTransition, Weakly ProTransition, Neutral, Weakly AntiTransition, or Strongly AntiTransition. To compare the categorical labels assigned by the researchers to those of ChatGPT, we aggregate labels Strongly ProTransition and Weakly ProTransition as 'favorable' and labels Weakly AntiTransition and Strongly AntiTransition as 'Unfavorable.' Figure IA.4 depicts the distribution of the categorical sentiment measures.

Next, we assign numerical values of -1, 0, and 1 to tweets labeled favorable, Neutral, and Unfavorable respectively. We then produce a correlation matrix between the categorical sentiment obtained from different researchers and ChatGPT-3.5. The values are reported in Table IA.5. We observe that the correlation between manually produced sentiment and ChatGPT sentiment ranges between 56% and 68%.

Finally, we compute the *average* sentiment produced by the research team for each tweet and compare it with the sentiment produced by ChatGPT-3.5. The first row of Table IA.6 shows that the correlation between the average sentiment produced by the five researchers and that by ChatGPT-3.5 is 68%, which is very close to that obtained using the numerical sentiment measure (70%). We also compute the average sentiment produced by the research team after dropping one researcher at a time and report its correlation with the sentiment produced by ChatGPT-3.5. These correlation values are reported in the subsequent rows of Table IA.6 and they range between 67% and 69%.

IA.2 Details of regulatory-channel analysis

This appendix gives a detailed description of the methodology we use to identify legislator tweets that mention environmental regulation and how we measure the support of a particular legislation toward the green transition.

IA.2.1 Identifying regulatory tweets

We now describe how we identify the legislator tweets with regulatory mentions by using ChatGPT combined with a dictionary approach.

For the 38,055 legislator tweets about climate change considered in Section 2.4, we ask ChatGPT-4o the following question: "*The following tweet was posted by a U.S. Congress member on* {`postdate`}*. Which regulation, policy, or bill does the tweet refer to? Provide the official name of the regulation and its bill number from congress.gov in the following format: [Regulation Name, Bill Number]. If the tweet does not refer to a regulation, policy, or bill, please respond with [None, None]. If the tweet refers to multiple regulations, policies, or bills, please separate the outputs with a comma*", where {`postdate`} refers to the date on which the tweet was posted.

We drop regulations that are tweeted about less than twenty times. We remove spurious entries and collate instances where the same regulation is referred to by different names.[25] We then cross check each regulation against `congress.gov` to document their bill number, the dates on which they were introduced or voted on, and the sequential number of the two-year term of the legislative body. Doing so yields a set of 19 regulations listed in Table 5. For example, the 'Climate Action Now Act' with bill number 'H.R.9' was introduced in the House of Representatives on March 27, 2019 by the 116th U.S. Congress. However, following its vote on May 2, 2019, it did not progress to the Senate.

Next, we manually inspect the content of the tweets to produce a dictionary in which each shortlisted regulation is associated with a set of keywords. We identify three categories of keywords related to the 19 regulations, which we report in Columns (2), (3), and (4) of Table IA.7, respectively. The first category is the bill number itself, as regulators often mention a regulation using its bill number directly. For example, consider the following tweet by Representative Claudia Tenney on March 30, 2023 which mentions the Lower Energy Costs Act: "*The average price for a gallon of gas has increased by over 40% since Biden took office. Today, House Republicans fulfilled another promise to the American people by passing H.R. 1, legislation to unleash our energy independence and drive down prices.*" We parse tweets against four permutations of the bill number with and without spaces and dots. For example, to check whether the tweet mentions the Lower Energy Costs Act, it is parsed against 'H.R.1', 'H.R. 1', 'HR1' and 'HR 1'.

The second category of keywords is *case-sensitive* abbreviations or acronyms associated with the regulation. The following tweet by Senator Jeanne Shaheen on December 21, 2022 uses the abbreviation 'BIL' for the Bipartisan Infrastructure Law, which is another name for the Infrastructure Investment and Jobs Act: "*Thrilled to see $45M now available to implement updated building energy codes in BIL following implementation of my bill w/*

[25] For instance, the 'Infrastructure Investment and Jobs Act' is also commonly referred to as the 'Bipartisan Infrastructure Law'.

@senrobportman. These grants will improve energy efficiency in buildings, lower costs & support good jobs." Because these keywords are case-sensitive, we search the tweet for 'BIL' but not 'bil'. Furthermore, we assert that 'BIL' is a *standalone word* and not part of another word (e.g., BILLION, MOBILE, or BILL).

The third category of keywords contains combinations of multiple words that are *not* case sensitive. For example, when parsing against the Infrastructure Investment and Jobs Act, searching for 'infrastructure bill' detects the tweet by Senator Patty Murray on 20 January, 2023: "*It's been two years since @POTUS & @VP took office. Since then, we've taken the boldest action on climate ever, created millions of jobs, lowered health care costs, passed a historic infrastructure bill, protected marriage equality, & so much more. And we're just getting started!*" We also search for these keywords without spaces to detect hashtags associated with the regulation. For instance, this allows us to identify the following tweet by Representative Shri Thanedar on 30 March, 2023: "*Do you know when would be a good time to act on climate change? That's right: NOW. Yesterday, I met with a group of inspiring students who are in D.C. to demand climate action. We owe it to future generations to act. #ClimateChangeIsReal #ClimateActionNow*".

Finally, we leverage the dictionary in Table IA.7 to parse the content of all 2.8 million legislator tweets downloaded from Twitter as explained in Section 2.1. We identify 40,821 tweets posted during NYSE trading hours that mention one of the 19 regulations. Of these, 4,092 tweets were classified by FinBERT as climate change related and were part of our baseline regression sample in Section 3.2. The remaining 36,729 tweets were not categorized by FinBERT as climate change related. Because congressional acts are comprehensive and contain clauses both related and unrelated to climate change, legislators are likely to discuss these regulations in alternative contexts. However, legislator outlook toward a clause unrelated to climate change could also have a differential impact on green versus brown stocks because the enactment of the regulation will also enforce clauses related to climate change.[26]

IA.2.2 Validating the measure of pro-transition legislative support

To validate the ChatGPT measure of pro-transition legislative support (PTLSupport) developed in Section 4.2.1, we compare it with that implied by the vote classification at the League of Conservation Voters (LCV), which was used by Lueck, Ramos Pastrana, and Tor-

[26]For example, Representative Zoe Lofgren tweeted about the Inflation Reduction Act on October, 12, 2022 in the context of healthcare: "*House Republicans want to repeal prescription drug provisions included in the Inflation Reduction Act & increase costs for seniors by banning Medicare from negotiating drug prices. The difference is clear: - GOP (handshake) Big Pharma - @HouseDemocrats (handshake) #PeopleOverPolitics*"

rens (2022) to study polarization in the U.S. Congress on environmental legislation. For votes related to environmental legislation, the LCV provides guidance on whether a vote in support of a bill or an amendment is a pro- or anti-environment position. We obtain the roll call numbers of votes associated with the 19 regulations listed in Table 5 from `congress.gov` and cross-reference them against the LCV database. Regulations can be associated with multiple votes during a given legislative session. We consider a vote determined as pro-environment by the LCV as a signal that the underlying regulation is pro-transition, and vice versa.

Table IA.8 reports the stance of the regulation toward the green transition obtained from ChatGPT and LCV. The third column shows that ChatGPT labels all regulations as pro-transition except for the 'Lower Energy Costs Act'. The fourth column reports the roll call number of the votes corresponding to the regulation, provided that the vote is documented in the LCV database. The fifth column reports the LCV classification of the vote. We observe that for regulations which have a single vote recorded in the LCV, classification by the LCV is aligned with that of ChatGPT. However, for regulations with multiple votes recorded in the LCV, the comparison is less straightforward because individual votes can be either pro- or anti-environment. However, we observe that for regulations classified as pro-transition by ChatGPT, the number of pro-environment votes recorded in LCV is greater or equal to that of anti-environment votes. Thus, we retain the categorization obtained from ChatGPT.

Figures and Tables for Internet Appendix

Figure IA.1: Average monthly numerical sentiment toward the green transition by party

This figure plots the average monthly numerical sentiment toward the green transition of tweets posted by Republicans (red) and Democrats (blue) between January 2014 and March 2023. The vertical gray dashed lines demarcate changes in U.S. presidency.

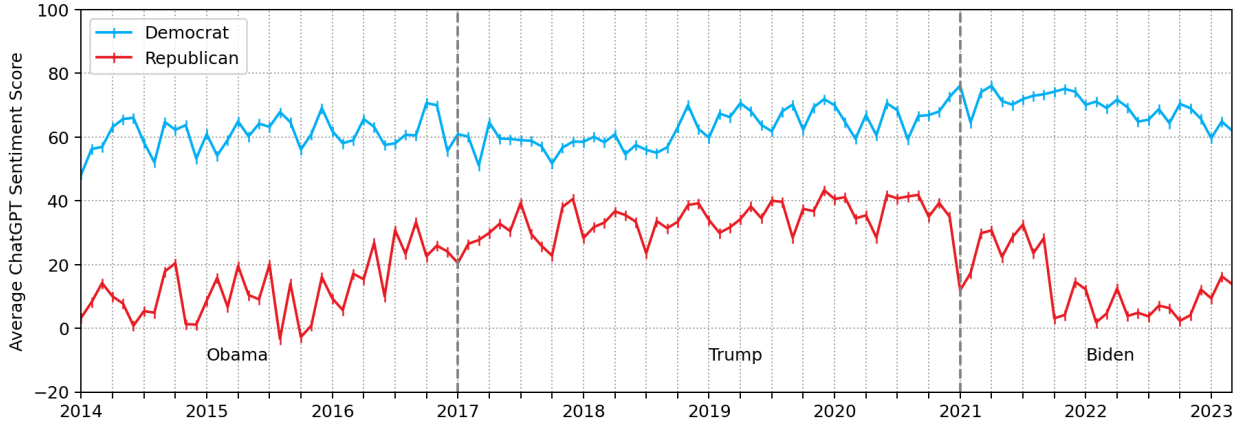


Figure IA.2: Distribution of pro-transition legislative support

This histogram depicts the empirical distribution of pro-transition legislative support (PTLSupport) obtained as described in Section 4.2.1, which measures the stance of a tweet toward pro-transition regulation between zero (strong opposition) and one (strong support).

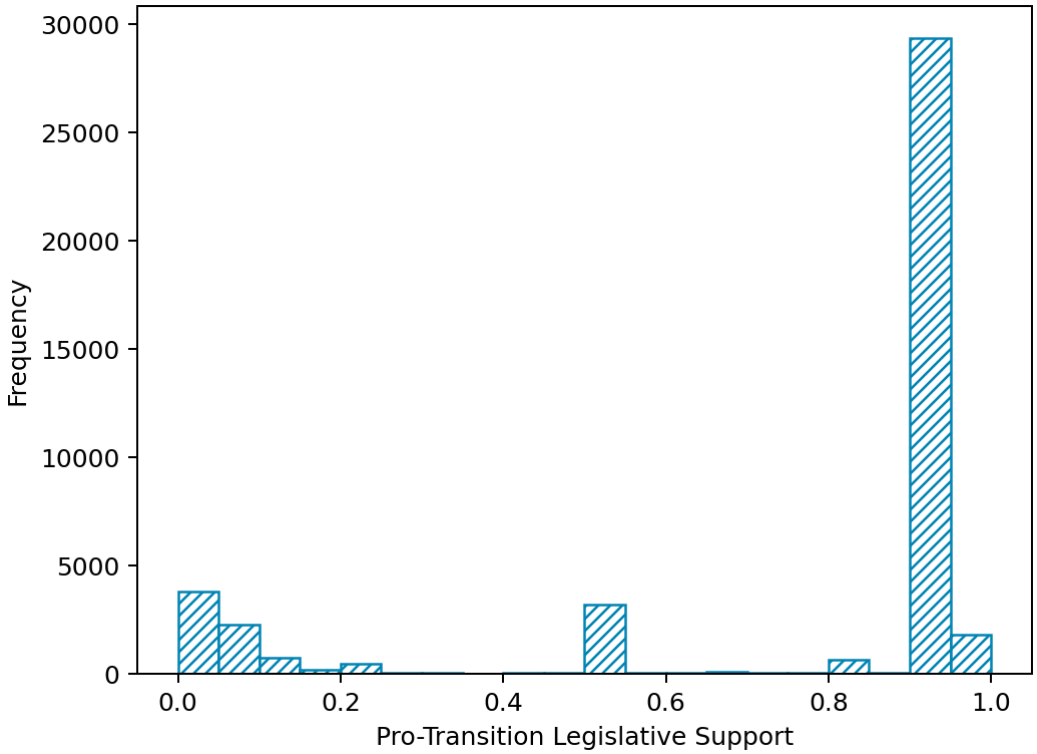


Figure IA.3: Distribution of manually assigned sentiment measures

Panels (a) through (e) depict the empirical distribution of numerical sentiment toward green transition assigned manually and independently by five researchers. Panel (f) depicts the numerical sentiment measure (between −100 to 100) obtained from ChatGPT in Section 2.3; the values are scaled by ten to facilitate comparison with the manually generated sentiment.

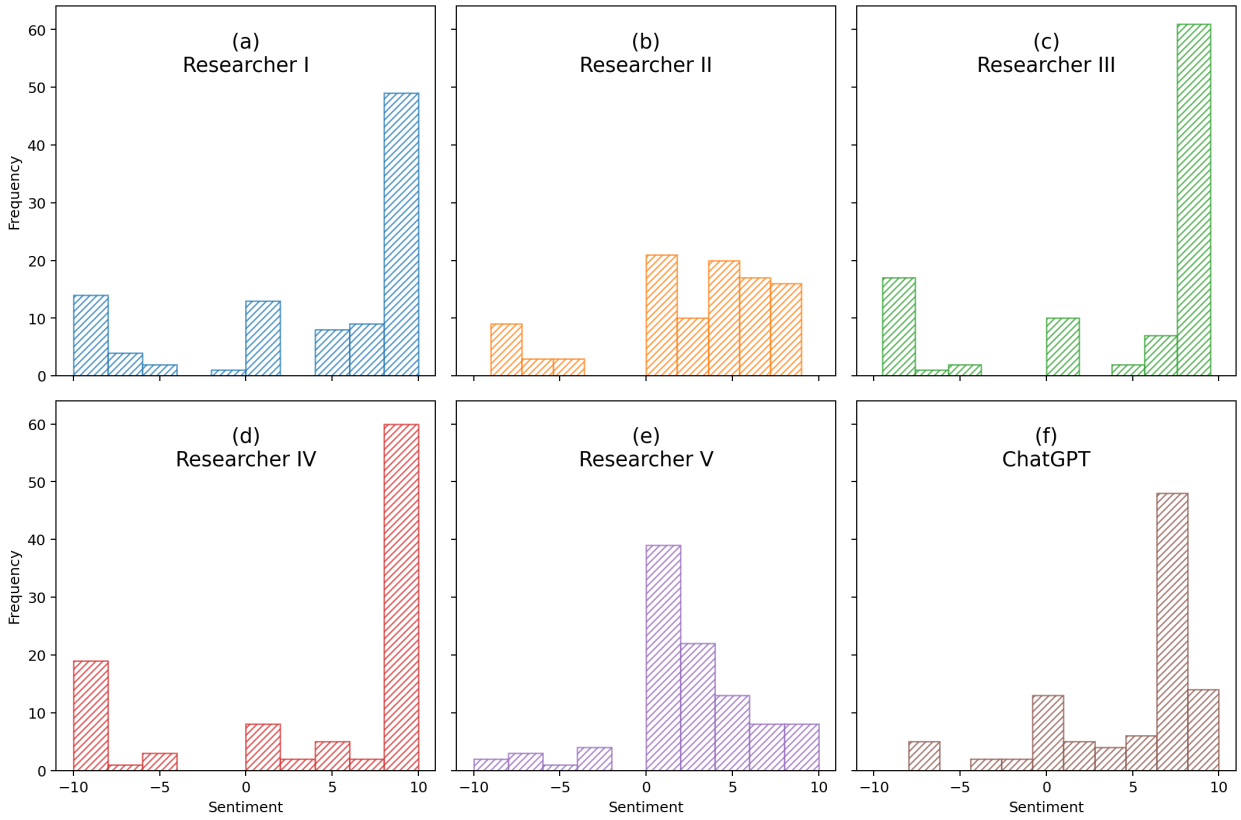


Figure IA.4: Distribution of categorical sentiment

This figure depicts the distribution of the categorical sentiment toward the green transition generated by five researchers and ChatGPT.

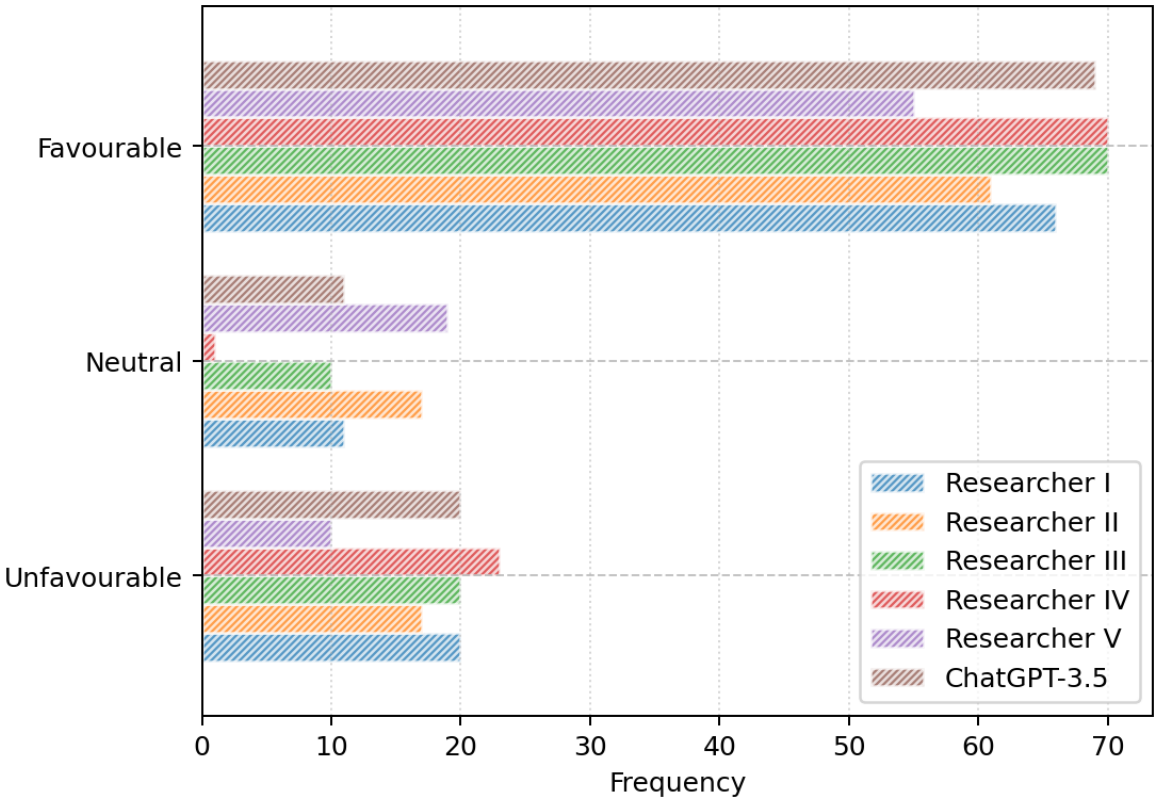


Table IA.1: ChatGPT sentiment for tweets classified "Negative" by FinBERT

This table reports the sentiment toward the green transition obtained from ChatGPT for a random sample of tweets classified as 'Negative' by the Tone FinBERT model.

	Tweets classified at Negative by FinBERT	ChatGPT Score
1	The Russian war on Ukraine makes clear that we cannot depend on foreign oil forever. To keep energy prices affordable and create a livable future for our children, we had to invest in #AmericanMade clean energy. The Inflation Reduction Act does just that.	70
2	New Jersey is a top #solar state & we must make sure we support our clean-energy workforce. This #solartariff move is bad for our state & I oppose it –> "President Trump Slaps Tariffs on Solar Panels in Major Blow to Renewable Energy" via @TIME:	70
3	Along with our thoughts & prayers to Pakistan, the U.S. must send tangible aid for flood victims. What we're witnessing are the deadly effects of climate change extending beyond the borders of top CO2 emission-producing nations. Together, we must do more to save our planet.	70
4	RT @RepDonBeyer: Climate change is one of the gravest dangers our society has faced in our lifetimes. The bill we are about to pass is mos...	80
5	To reach our ambitious climate goals, we must unleash the full power of American clean energy – and that includes offshore wind energy development. The former administration's arbitrary offshore wind leasing moratorium would stifle the growth of our clean energy economy in NC.	70
6	There is no single barrier to the transition to electric vehicles, but rather an entire ecosystem of challenges to address. We must do more to accelerate our transition to an electric vehicle future. My op-ed in @thehill	80
7	Solar tariffs could cost us 100,000 jobs in the solar industry and hurt our transition to a cleaner energy future. I've joined my colleagues, including @RepScottPeters, in calling on the Biden admin to protect domestic solar energy capacity & halt consideration of solar tariffs.	60
8	Banning Russian oil imports is a necessary step to continue to put economic pressure on Putin and support our Ukrainian allies. This situation underscores why we must invest in clean energy long-term to reduce our dependence on foreign oil.	60
9	In North Carolina, chemical companies have polluted the Cape Fear River with PFAS chemicals for years, resulting in serious health consequences for residents. This legislation invests $55 billion in clean drinking water, the largest investment in American history.	80
10	We advanced investments in resilient infrastructure against the dire effects of climate change. North Carolina knows the detrimental effects of extreme weather too well. From highways to roads to farms, hurricanes cost our state billions of dollars.	70
11	Hurricane Ida devastated communities up and down the East Coast while wildfires persist on the West Coast. Extreme flooding, lethal storms, dangerous wildfires. This is our future if we don't take serious steps to address the climate crisis now.	90
12	RT @CCLCharlotteNC: We appreciate your concern for the #climate and action towards #cleanenergy @RepAdams @RepRichHudson!	80
13	Republicans are delivering on their Commitment to America. The Lower Energy Costs Act is just one example. We know Americans are suffering in President Biden's America, but we are hard at work to end the multitude of crises created by this administration's failed policies!	0
14	WATCH: By closing the Keystone XL Pipeline, President Biden has killed American jobs and limited our ability to produce clean energy here at home. I introduced a resolution in @HouseCommerce to condemn this terrible move.	0
15	Rush-to-green energy policies created a widespread energy crisis in Europe as we head into winter. This should serve as a warning that enacting rushed Green New Deal energy policies in the U.S. would be nothing short of a catastrophe.	-65
16	FACT CHECK: FALSE. Gas in California cost $3.36 per gallon on average in January 2021, when President Biden took office. The sky-high price of gas is unmistakably a result of President Biden and Democrats' disastrous energy policies.	0
17	Democrats' rush to green energy agenda has created an energy crisis in California. I was proud to join @boblatta, @RepGusBilirakis, and @RepJohnJoyce's letter to President Biden sounding the alarm on the state's decision to ban gas-powered cars by 2035.	-65
18	(2/2) Farmers could have used this water in light of the record droughts and Arizona's reduction in its water allocation.	60
19	(1/2) I led the Arizona Delegation in requesting answers from the Bureau of Indian Affairs regarding the Bureau's failure to notify local farmers about a recent excess water release from the San Carlos Reservoir.	0
20	California's rush-to-green energy policies have sent energy prices soaring, including at the gas pump. House Democrats' Green New Deal climate mandates would force the rest of our nation to face the same disastrous fate as California.	-65

Table IA.2: Carbon emissions for Fama-French 48 industries

This table reports the (absolute) Scope 1, 2, and 3 carbon emissions and carbon emissions intensity averaged across firms in each of the Fama-French 48 industries, obtained from Refinitiv. Absolute carbon emissions are reported in metric tonnes of carbon-dioxide equivalent (tCO2e). The carbon emissions intensity of a firm in a given fiscal year is obtained by dividing the sum of its absolute Scope 1 and Scope 2 GHG emissions by total revenue (expressed in millions USD). Industries are sorted in an ascending order by their carbon emissions intensity. To categorize firms into industries, we link tickers in emissions data to CUSIPs obtained from CRSP. We then link CUSIPs to SIC codes obtained from Compustat. We then map these SIC codes to the corresponding Fama-French industry classifications using the definitions on Ken French's website.

	Industry	Scope 1	Scope 2	Scope 3	Intensity
1	Insurance	11,119	58,541	256,153	4.96
2	Banking	85,407	507,727	497,034	9.29
3	Recreation	9,312	52,725	725,761	11.99
4	Business Services	30,963	247,228	845,818	14.14
5	Apparel	29,288	91,703	2,860,287	15.46
6	Real Estate	17,943	17,974	4,521,554	17.14
7	Printing and Publishing	19,277	134,941	1,160,446	19.66
8	Wholesale	148,763	97,171	12,788,785	19.67
9	Tobacco Products	274,127	265,522	3,138,091	20.16
10	Computers	38,481	277,267	6,321,508	20.94
11	Measuring and Control Equipment	36,750	94,473	1,005,097	21.42
12	Defense	284,372	797,762	26,699,347	21.57
13	Construction	90,099	30,169	1,505,612	21.69
14	Medical Equipment	105,007	161,297	2,909,812	22.46
15	Healthcare	135,113	494,937	1,841,901	26.81
16	Retail	666,578	1,385,147	8,573,759	34.48
17	Communication	547,863	3,164,319	4,228,587	37.41
18	Machinery	118,630	221,445	131,741,114	41.75
19	Consumer Goods	396,655	460,887	23,587,101	43.26
20	Aircraft	330,942	522,363	6,638,688	43.28
21	Trading	21,883	197,339	190,512	45.04
22	Shipbuilding, Railroad Equipment	54,801	69,798	31,861	51.96
23	Automobiles and Trucks	454,429	1,015,292	119,519,381	58.47
24	Pharmaceutical Products	250,149	299,499	2,534,873	71.37
25	Rubber and Plastic Products	55,292	172,968	1,390,118	75.89
26	Beer & Liquor	3,800,645	1,896,167	27,105,002	86.74
27	Entertainment	169,944	475,641	645,072	91.76
28	Electronic Equipment	316,274	475,287	3,826,229	95.32
29	Candy & Soda	1,367,703	2,732,258	26,672,261	101.09
30	Fabricated Products	55,659	41,508	–	102.34
31	Food Products	1,157,164	663,036	7,423,648	102.96
32	Electrical Equipment	98,717	111,212	5,816,558	131.86
33	Restaurants, Hotels, Motels	263,139	1,076,852	9,382,590	164.91
34	Personal Services	1,538,768	84,728	3,886,528	206.50
35	Construction Materials	685,089	469,601	1,487,454	220.13
36	Textiles	995,130	817,932	546,690	231.74
37	Business Supplies	2,253,913	1,503,909	6,184,353	278.99
38	Shipping Containers	1,834,866	1,033,090	4,615,030	333.35
39	Precious Metals	1,591,431	770,847	3,036,569	412.14
40	Non-Metallic and Industrial Metal Mining	2,046,440	1,284,317	846,727	428.44
41	Agriculture	54,960	538	–	511.20
42	Transportation	11,989,562	280,562	4,045,709	576.07
43	Petroleum and Natural Gas	16,475,804	1,738,141	125,520,137	691.89
44	Chemicals	3,785,647	1,426,095	18,617,774	758.60
45	Steel Works Etc	8,255,203	2,131,831	13,103,009	1,119.81
46	Almost Nothing	19,845,106	241,514	5,251,020	1,756.65
47	Coal	6,976,498	666,579	–	1,914.86
48	Utilities	15,334,635	1,158,413	22,524,289	2,018.30

Table IA.3: Correlation between manually generated and ChatGPT sentiment

This table shows the correlations between the numerical sentiment generated by the five researchers and ChatGPT.

	Researcher I	Researcher II	Researcher III	Researcher IV	Researcher V	ChatGPT3.5
Researcher I	1.00	0.82	0.86	0.88	0.61	0.61
Researcher II		1.00	0.83	0.80	0.64	0.59
Researcher III			1.00	0.91	0.67	0.71
Researcher IV				1.00	0.62	0.71
Researcher V					1.00	0.53
ChatGPT-3.5						1.00

Table IA.4: Correlation between average researcher and ChatGPT sentiment

The first row reports the correlation between the average sentiment produced by the research team and that by ChatGPT. The subsequent rows report correlations with the sentiment produced by ChatGPT after dropping one researcher at a time.

	Correlation with ChatGPT
Overall Average	0.70
Dropped Researcher I	0.72
Dropped Researcher II	0.71
Dropped Researcher III	0.69
Dropped Researcher IV	0.68
Dropped Researcher V	0.70

Table IA.5: Correlations between manually generated and ChatGPT sentiment

This table shows the correlation matrix between categorical sentiment generated manually by the five researchers and that obtained from ChatGPT. The correlations are computed by assigning numerical values of −1, 0, and 1 to tweets labeled Favorable, Neutral, and Unfavorable, respectively.

	Researcher I	Researcher II	Researcher III	Researcher IV	Researcher V	ChatGPT3.5
Researcher I	1.00	0.75	0.86	0.86	0.69	0.56
Researcher II		1.00	0.82	0.82	0.62	0.62
Researcher III			1.00	0.91	0.77	0.67
Researcher IV				1.00	0.77	0.68
Researcher V					1.00	0.56
ChatGPT-3.5						1.00

Table IA.6: Average correlation between researcher and ChatGPT sentiment

The first row reports the correlation between the average categorical sentiment produced by the research team and that by ChatGPT-3.5. The subsequent rows report correlations with the sentiment produced by ChatGPT-3.5 after excluding one researcher at a time. The correlations are computed by assigning numerical values of −1, 0, and 1 to tweets labeled favorable, Neutral, and Unfavorable, respectively.

	Correlation with ChatGPT
Overall Average	0.68
Dropped Researcher I	0.69
Dropped Researcher II	0.67
Dropped Researcher III	0.68
Dropped Researcher IV	0.67
Dropped Researcher V	0.68

Table IA.7: Dictionary for regulations discussed in environmental tweets

This table reports a representative collection of keywords for each regulation used to identify tweets with direct regulatory mentions. The first column titled 'Regulations' reports the name of the regulation against which we would like to examine the content of tweets. The second column titled 'Code' reports the corresponding bill number. The tweet is parsed against permutations of the bill number with and without spaces and dots. For example, to check whether the tweet mentions the Lower Energy Costs Act, it is parsed against 'H.R.1', 'H.R. 1', 'HR1' and 'HR 1'. The third column titled 'Abbrevs' contains *case sensitive* keywords against which a given regulation is parsed. For example, the Inflation Reduction Act is parsed against 'IRA' but not 'ira'. Finally, the fourth column titled 'Tokens' contains keywords that are *not* case sensitive. Tweets are converted to lowercase and then parsed against these keywords with and without spaces. For example, the Climate Action Now Act is parsed against 'climate action now' and 'climateactionnow'. Keywords without spaces help identify hashtags (e.g., '#ClimateActionNow').

Regulation	**Code**	**Abbrevs**	**Tokens**
Inflation Reduction Act	H.R.5376	IRA	inflation reduction, build back better, buildback
Infrastructure Investment and Jobs Act	H.R.3684	IIJA, BIL	infrastructure investment and jobs, bipartisan infrastructure, infrastructure bill, infrastructure framework, infrastructure investment
Climate Action Now Act	H.R.9	CANAct	climate action now
Lower Energy Costs Act	H.R.1		lower energy costs
Green New Deal	H.Res.109	GND	green new, green deal
Growing Climate Solutions Act	S.1251	GCSA	growing climate solutions
Climate Risk Disclosure Act	H.R.2570		climate risk disclosure
America COMPETES Act	H.R.4521	COMPETES	america competes
Moving Forward Act	H.R.2		moving forward
INVEST in America Act	H.R.7095	INVEST	invest in america
American Innovation and Manufacturing Act	S.2754	AIM, HFC	american innovation and manufacturing, phase down, phasing down, hydrofluorocarbon
CLEAN Future Act	H.R.1512	CLEAN	clean future
Growing Renewable Energy and Efficiency Now Act	H.R.848	GREEN	growing renewable energy and efficiency now
Trillion Trees Act	H.R.2639		trillion trees, trillion new
Energy Policy Modernization Act	S.2012	EPMA	energy policy modernization, energy modernization, energybill
CHIPS and Science Act	H.R.4346	CHIPS	chips and science
100% Clean Economy Act	H.R.5221		100% clean economy
Energy Innovation and Carbon Dividend Act	H.R.763		energy innovation and carbon, carbon dividend
USE IT Act	S.383	USE	use it

Table IA.8: Regulation support to green transition by ChatGPT and LCV

This table reports the stance of regulations toward the green transition. The first column titled 'Regulations' reports the name of the regulation. The second column titled 'Congress' reports the sequential number of the two-year term of the legislative body. The third column titled 'ChatGPT' reports the stance of the regulation toward green transition obtained from ChatGPT—"Support" means that the regulation is pro-transition, while "Oppose" means that it is anti-transition. The fourth column subtitled 'Vote' documents the roll call number of the vote corresponding to the regulation. The fifth column subtitled 'Position' reports whether a vote in support of the regulation is considered a pro-environment or an anti-environment position according to the League of Conservation Voters (LCV). A regulation can be associated with several votes. The table reports all votes associated with a given regulation in `congress.gov` that are documented in the LCV scorecard. Entries are marked with a "—" if (i) the regulation was introduced but not voted on (e.g., it was deferred to a subcommittee) or (ii) the vote corresponding took place but was not documented in the LCV scorecard.

Regulation	**Congress**	**ChatGPT**	**LCV**	
			Vote	*Position*
Inflation Reduction Act	117th	Support	420	Pro-environment
Infrastructure Investment and Jobs Act	117th	Support	385	Pro-environment
Climate Action Now Act	116th	Support	181	Pro-environment
Lower Energy Costs Act	118th	Oppose	182	Anti-environment
Green New Deal	116th	Support	—	—
Growing Climate Solutions Act	117th	Support	—	—
Climate Risk Disclosure Act	117th	Support	—	—
America COMPETES Act	117th	Support	27	Pro-environment
			31	Anti-environment
Moving Forward Act	116th	Support	138	Pro-environment
			136	Pro-environment
			135	Anti-environment
INVEST in America Act	116th	Support	—	—
American Innovation and Manufacturing Act	116th	Support	—	—
CLEAN Future Act	117th	Support	—	—
Growing Renewable Energy and Efficiency Now Act	117th	Support	—	—
Trillion Trees Act	117th	Support	—	—
Energy Policy Modernization Act	114th	Support	9	Pro-environment
CHIPS and Science Act	117th	Support	—	—
100% Clean Economy Act	116th	Support	—	—
Energy Innovation and Carbon Dividend Act	116th	Support	—	—
USE IT Act	116th	Support	—	—

Table IA.9: Voting margin and legislator support toward pro-transition regulation

This table reports the impact of legislator support toward pro-transition regulation (PTLSupport) on green versus brown stock prices in the weeks leading up to congressional votes based on the tightness of the voting margin. Voting margin is computed by dividing the difference between the *Yay* (for) and *Nay* (against) votes by the total number of votes cast for a given regulation. A congressional vote is considered "close" when the voting margin is less than 5%, and "clear" otherwise. Columns (1) and (2) reproduce Columns (4) and (5) of Table 6 for tweets posted in the weeks leading up to congressional votes that were close. Columns (3) and (4) reproduce Columns (4) and (5) of Table 6 for tweets posted in the weeks leading up to congressional votes that were clear. The dependent variable is the one-to-ten minute return stock return around each tweet, $\Delta p(-01, +10)$. Stocks that fall into the top and bottom tercile of carbon emissions intensity—obtained by using the most recent Scope 1 and 2 emissions data available on Refinitiv Eikon at the time each tweet was posted—are classified as green and brown, with the rest classified as gray. Standard errors are robust to heteroscedasticity and clustered at the firm level. The t-statistics are reported in parentheses and p-values smaller than 0.10, 0.05, and 0.01 are indicated by one (*), two (**), and three (***) stars, respectively.

	(1) CLOSE (leading 2wks)	(2) CLOSE (leading 1wk)	(3) CLEAR (leading 2wks)	(4) CLEAR (leading 1wk)
Brown	0.566	0.370	-0.572	-0.0845
	(1.16)	(0.59)	(-1.05)	(-0.14)
Green	0.0649	-0.162	0.982	1.196
	(0.17)	(-0.32)	(1.19)	(1.16)
PTLSupport	-2.610***	-2.010***	1.623***	2.131***
	(-12.76)	(-7.19)	(3.17)	(3.21)
Brown × PTLSupport	-0.739***	-0.779***	-0.160	-0.0351
	(-3.62)	(-2.92)	(-0.33)	(-0.06)
Green × PTLSupport	-0.208	0.130	0.275	-0.968*
	(-0.97)	(0.47)	(0.55)	(-1.78)
Constant	1.834***	1.592***	-0.195	-0.0995
	(7.39)	(4.78)	(-0.41)	(-0.17)
Firm FE	Yes	Yes	Yes	Yes
Legislator FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Adjusted R2	0.0417	0.0600	0.139	0.189
Observations	1281044	783919	163675	112669