Climate Regulation and Emissions Abatement: Theory and Evidence from Firms' Disclosures

Finance Working Paper N° 730/2021
May 2023

Tarun Ramadorai
Imperial College London, ABFER, NCAER, CEPR and ECGI

Federica Zeni
Imperial College London

ECGI Working Paper Series in Finance

Climate Regulation and Emissions Abatement: Theory and Evidence from Firms' Disclosures

Working Paper N° 730/2021
May 2023

Tarun Ramadorai
Federica Zeni

We thank Pat Akey, Ian Appel, Harjoat Bhamra, Patrick Bolton, Charlie Donovan, Chris Hansman, Harrison Hong, Zacharias Sautner, and seminar participants at Imperial College Business School, Cornell University, the UZH workshop on Climate Finance, the HEC Paris Spring Finance conference, and the Western Finance Association conference for comments.

Climate Regulation and Emissions Abatement: Theory and Evidence from Firms' Disclosures

Tarun Ramadorai and Federica Zeni†

May 19, 2023

Abstract

We measure firms' beliefs about climate regulation, plans for future abatement, and current emissions mitigation from responses to the Carbon Disclosure Project. These measures vary strikingly around the Paris announcement. A dynamic model of a representative firm facing a future carbon levy, trading off abatement and capital growth, and facing convex adjustment costs cannot fit the data. A two-firm model with cross-firm reputational externalities, heterogeneous beliefs over climate regulation, and leader-follower interactions does. Out-of-sample, the model predicts firms' reactions when the US exits the Paris agreement. Firms' beliefs about climate regulation strongly affects abatement, and cross-firm interactions amplify regulatory impacts.

1 Introduction

Climate change poses a looming threat to economic and financial stability, lending urgency to calls for action on global warming.[1] Faced with such warnings, in December 2015, 196 nations signed an agreement at the United Nations Framework Convention on Climate Change (UNFCCC) in Paris, to limit greenhouse gas emissions to a level consistent with global temperatures rising less than

*We thank Pat Akey, Ian Appel, Harjoat Bhamra, Patrick Bolton, Charlie Donovan, Chris Hansman, Harrison Hong, Zacharias Sautner, and seminar participants at Imperial College Business School, Cornell University, the UZH workshop on Climate Finance, the HEC Paris Spring Finance conference, and the Western Finance Association conference for comments.

†Ramadorai (corresponding author): Imperial College Business School and CEPR. Email: t.ramadorai@imperial.ac.uk. Zeni: World Bank Research. Email: fzeni@worldbank.org.

[1] For example, see Carney (2015).

2°C. The agreement also determined a five-year window within which countries could meet and renew the so-called Nationally Determined Contributions (NDCs). However, between 2015 and 2020, most signatory countries fell far short of required targets,[2] and the world's second largest emitting country withdrew from the agreement.[3]

How important is such coordinated climate regulation in determining firms' carbon mitigation, and through what channels does such regulation affect firms? To answer these questions, we harness comprehensive micro data and interpret these data through the lens of structural models. The data capture firms' voluntary disclosures about their beliefs about regulation, their plans for future abatement, and their current carbon mitigation actions. In the years leading up to and following the Paris announcement, these measures exhibit significant and striking variation. To understand the economic drivers of these empirical observations, we build a set of dynamic models of firms' emissions abatement, and explore which combination of model ingredients can best match the patterns observed in the data.

Our analysis reveals that firms' reported beliefs about climate regulation strongly influence their planned and actual abatement activities. To match the patterns and magnitudes in firms' responses up to and just following the Paris announcement, we find that two key model ingredients are needed, namely i) the presence of cross-firm reputational externalities, and ii) heterogeneous beliefs across firms about the stringency of the regulatory policy. These ingredients amplify firms' reactions to climate regulatory announcements, leading us to conclude that climate regulation can have substantial effects on firms' abatement actions. Our approach allows us to recover estimates of key parameters such as firms' implied priors about the cost of climate regulation, and the extent to which firms respond to one anothers' actions. We validate the model and these conclusions in an out-of-sample exercise, in which we predict abatement actions from firms' revised beliefs following President Trump's announcement to pull back from the Paris agreement.

Our data track North American public firms that voluntarily disclose environmental information through the Carbon Disclosure Project (CDP) between 2011 and 2017.[4] The CDP data comprise

[2]As an example, see the WSJ article *https://www.wsj.com/articles/u-n-finds-nations-climate-plans-fall-short-of-paris-accord-11635259444*.

[3]In June 2017, U.S. President Donald Trump announced his intention to withdraw from the agreement, with the final decision made in November 2020. On February 2021, under the Biden administration, the United States officially rejoined the Paris agreement.

[4]We verify the accuracy of these data using third-party sources such as Bloomberg, Thomson Reuters, and MSCI,

three important dimensions, namely, firms' self-reported beliefs about the horizon and impact of future climate regulation; firms' plans for future carbon emissions abatement; and finally, data on firms' emissions abatement actions to date, which reflect the actual changes in their carbon footprints.

In our empirical work, we compare, at each reporting date, the dynamics of firms' reported beliefs about the intensity of future climate regulation with their carbon abatement actions. We document that there are important cross-sectional differences between two groups of firms in the data. One set comprises firms that publicly report their plans for future emissions reduction in addition to reporting their beliefs about the intensity of future climate regulation and their current actions on abatement. The other set comprises firms that report beliefs and current abatement actions, but do not report their plans for future abatement. The two sets of firms differ in several other ways. The plan-reporting firms are larger and more profitable, more emissions intensive, and also have a greater propensity than non-plan-reporting firms to i) engage with policymakers and climate regulators, ii) actively issue "green" debt securities, and iii) provide direct funding to climate regulatory activities.

Between 2011 and 2015, prior to the Paris agreement announcement, both plan-reporting and non-plan-reporting firms, on average, moderately downgrade their expectations over the impact of future regulation, while progressively increasing their actual carbon footprints. In 2015, the year of the Paris agreement announcement, plan-reporting firms upwardly revise their beliefs about future climate regulation. In contrast, firms without plans revise their beliefs downward, in a manner that seems to reflect the reduction in global crude oil prices in 2015. In 2016, the year of the enactment of the Paris agreement, firms without plans reverse the trend in their beliefs, upwardly revising their expected impact of climate regulation. Following the Paris agreement, all firms sharply increase carbon abatement over the year from 2016 to 2017, but strikingly, despite the fact that the belief revisions of plan-reporting firms are considerably *smaller* than those of non-plan reporting firms, plan-reporting firms react *far more* to the Paris agreement than non-plan-reporting firms. Put differently, plan-reporting firms have *more* extreme reactions to the climate regulation event, despite being seemingly *less* surprised by the agreement.[5]

who produce external audits and ratings of firms' ESG activities.

[5]Indeed, the belief revisions of these reporting firms seem to anticipate the regulatory announcement, intensifying

We view these three observations, namely that i) revisions in beliefs are far less pronounced than revisions in actions; ii) there is pronounced heterogeneity in beliefs around the Paris announcement; and iii) reported beliefs for different groups of firms don't map in the same way to their observed actions after the announcement, as important new targets for any model.

As a first step towards rationalizing these observations, we build a simple dynamic model of a representative firm's emissions reduction activities. In the model, the polluting firm produces output (and carbon emissions proportional to output) in each period using capital. The firm is exposed to a future climate regulation event, which takes the form of a terminal carbon levy of uncertain intensity. At any time period prior to the regulation event, depending on its belief over the levy, the firm can abate or increase emissions by reducing or increasing its level of polluting capital, but faces standard convex adjustment costs for such changes. The firm's optimal policy balances the tradeoff between output growth and emissions reduction. Since the carbon levy is only realized at the terminal date, the firm discounts the uncertain cost of regulation to the present, and sets an optimal profile ("plan") of abatement which begins in the current period—i.e., its current abatement action—and then in every period leading up to the date of the levy.

An important object of interest is firms' prior belief about the size of the regulator-imposed carbon levy. To recover this quantity, we calibrate the simple model to the cross-sectional average of plan-reporting firms. We feed the model with data on firms' self-reported beliefs about future climate regulation, and adjust the prior on the carbon levy to maximize the fit of model-implied abatement plans and actions with those seen in the data. We estimate an implied prior of roughly 80\$/mt CO_2e, which falls in the range estimated in the extensive literature on the social cost of carbon (see, e.g., Tol (2011)), but is far higher than the prevailing market price of carbon.[6] While these estimates could be reasonable, we show that the resulting dynamics of abatement actions implied by the simple model are excessively smooth, and fail to capture the substantial increases in abatement seen in the data around the Paris announcement.

To improve the performance of the model, we introduce a second firm into the economy, to

the puzzle about their more extreme reaction to the event.

[6]The price of carbon allowances traded in the European cap and trade market between 2011 and 2016 averaged around 10\$/mtCO_2e; and between 2010 and 2017, the US government's official estimate of the social cost of carbon, as provided by the inter-agency working group (IWG), averaged around 52\$/mtCO_2e. In 2017, the Trump administration disbanded the IWG, and since that time have used estimates of the social cost of carbon that range between 1\$ and 7\$/mtCO_2e.

capture the observed heterogeneity between plan- and non-plan-reporting firms. We allow the two firms in the extended model to have different beliefs about the terminal levy, and to emphasize their strategic interactions, we model one firm as a "leader',' which anticipates the beliefs of the other firm, while the "follower" firm simply optimizes with respect to its own beliefs. This assumption is motivated by the evidence discussed earlier on the observed differences between plan- and non-plan-reporting firms, which we map in the model to plan-reporting firms' propensities to act as leaders, and non-plan-reporting firms to act as followers in the carbon abatement market.

The final ingredient that we add to this extended model is a reputational externality which connects each firm's profits to the abatement actions of the other firm. More specifically, the reputational parameter controls how a firm's profits from polluting change when the other firm abates emissions at the same time. This model feature amplifies the impact of firms' beliefs about climate regulation on their physical actions on abatement.

Armed with these assumptions, we solve the enriched model for the equilibrium of a dynamic Stackelberg leadership game where the leader firm has a first-mover advantage over the follower firm. The leader maximizes profits, internalizing the follower's reaction to its actions. In keeping with standard Stackelberg intution, this leads to the leader reacting more to changes in information about the terminal levy despite it being less surprised by this information than the follower. The model delivers several other useful predictions. First, we find that the reputational externality generates an amplified reaction by firms to changes in the levy, with the leader (i.e., plan-reporting, in the data) firm reacting more than the follower (non-plan-reporting) firm to variations in its own beliefs about the levy because of its leadership position in the Stackelberg equilibrium. Second, we find that the leader's reaction to variation in the beliefs of the follower firm can be larger than the follower's own reaction to such variation if the reputation externality is large enough.

We next take this more complex model to the data, allowing for the size of the parameter governing the reputational externality to be structurally estimated. In the data, we link this parameter to a measure of current attention paid to all firms' sustainability performance using news sources, and find that it is strongly positive, i.e., firms' profits from polluting are significantly reduced when their competitors contemporaneously abate emissions. This result provides evidence of relative performance evaluation of firms along the dimension of their Environmental, Social, and

Governance (ESG) activities.

The more richly parametrized model, as we might expect, yields predictions that are closer to the observed data—although it is worth pointing out that this is not just mechanical, since the model now needs to fit the average and variance of disclosures of *both* plan-reporting and non-plan reporting firms. Overall, we find that the model is well-able to fit the observed dynamics of firms' abatement plans and actions before and after the announcement of the Paris agreement, helping to explain why firms' reactions to the Paris announcement are both high—the new model ingredients result in substantial amplification of the impact of climate regulation relative to the basic model—and different across the two groups of firms.

Our estimates of the reputation externality reveal that on average, in each time period, reputation benefits account for roughly 15% of the terminal costs introduced by the levy. We use these estimates to evaluate the optimal path of carbon emissions generated by the calibrated models for a time horizon of ten years and two policy scenarios corresponding, respectively, to distributed levies (i.e., applied at each time period) of 25\$/mtCO_2e and 125\$/mtCO_2e.[7] We postulate that reputational externalities are highest around the regulatory event and then decline over time, a condition that we show to generate a declining time-path of abatement, i.e., firms will optimally abate a large share of their polluting capital immediately. Comparing these cases, we show that a 125\$/mtCO_2e levy will be needed to meet reduction targets established at the Paris agreement.[8] Importantly, in the more stringent case, the augmented model with reputation predicts a substantial amplification of the firm's baseline reaction to the policy in the short run.

In a final exercise, to validate these estimated parameters as well as our conclusions from the model, we acquire data to extend our sample through 2019, to evaluate the impacts of Trump's announcement in June 2017 to pull back from the Paris agreement. We show that firms' reported beliefs about the intensity of future climate regulation following Trump's announcement drop sharply, with larger reported belief updates seen for plan-reporting firms. Firms also report revisions to their expected horizons of emissions abatement, which are pushed further into the future. We feed these reported beliefs from the extended sample into the model with parameters fixed at their es-

[7]According to recent academic studies (see Carleton and Greenstone (2021)), a social cost of carbon updated to respond to the frontier of climate science would average around 125\$/mtCO_2e.

[8]That is, the extra emissions reduction introduced by reputation in the case of a 25\$/mtCO_2e is not enough to meet the Paris target. This because the amplification effect increases in the size of the carbon levy.

timated values in the pre-2017 period, and demonstrate that our model predicts the patterns seen in emissions abatement in the out-of-sample period.

The remainder of this paper is structured as follows. The remainder of this section contains a brief discussion of closely related literature. Section 2 introduces the CDP dataset, validates the disclosure data using external sources, and describes the construction and measurement of the empirical evidence. In Section 3, we describe and solve the simple dynamic abatement model with an atomistic firm, and calibrate it to the data. Section 4 introduces, solves, and calibrates the more complex two-firm model, and discusses the differences between this model and the simple model. Section 5 describes our out-of-sample exercise, and Section 6 concludes. An Online Appendix contains more detailed descriptions of the underlying data and our constructed measures, detailed model derivations, a more comprehensive review of related literature, and several other auxiliary exercises.

1.1 Related Literature

Our paper shows that firms' beliefs about future climate regulation influence their emissions reduction activities. A recent study by Biais and Landier (2022) takes our evidence as a starting point, and warns that in an equilibrium where regulators have limited commitment power, beliefs about weak future climate regulation can be self-fulfilling. We also document a strong link between firms' reported plans to reduce carbon emissions and their subsequent emissions reductions, which connects our work to Bolton and Kacperczyk (2021b). This finding parallels macro-evidence in Tenreyro and De Silva (2021) that relates countries' climate pledges to future emissions reductions.

A second important contribution of our paper is our use of a simple abatement model to infer firms' implied prior about the future social cost of carbon. The estimates that we acquire are higher than those revealed by market prices of pollution permits traded in the European cap and trade market, or those inferred by Meng (2017) (who considers the failure of the U.S. Waxman-Markey bill). These findings suggest that raising the cost of climate regulation may not come as a shock to firms. In a similar spirit to Barnett et al. (2020), our implied estimates help to quantify the negative impact of regulation uncertainty,[9] although in contrast with their work, we find that the

[9] In our model, firms are risk neutral yet they act as if they are risk-averse, decreasing emissions abatement when future regulatory uncertainty is higher, because they face convex adjustment costs in physical capital (this is similar, for example, to Rampini et al. (2014)).

effects of uncertainty are more muted in comparison with those arising from shifts in firms' priors.

Third, we confirm that market interest in firms' sustainability practices (which we proxy using the dynamics of ESG ratings' related news embedded in our model through the reputational externality channel), is necessary to rationalize puzzling patterns in the data on firms' responses to climate regulation. Theoretically, the assumption that investors have preferences for sustainability has been the starting point of related work such as Pástor et al. (2021) and Hong et al. (2021), which study the equilibrium implications of such preferences on firm value and global welfare, and Barbalau and Zeni (2023), which studies the impact of such preferences on firms' financing choices. Empirically, the literature has attempted to document the existence (and quantify the impact) of ESG preferences, for example, Hartzmark and Sussman (2019) study announcements of mutual funds' ESG ratings to infer preferences from investors' capital reallocation, while Zerbib (2019) and Flammer (2021) attempt to infer pro-environmental preferences from green bond issuance. We show that the reputation externalities amount to roughly 15% of the (perceived) cost of the future levy, providing an additional quantitative assessment that is useful for this growing literature.

2 Data

2.1 Carbon Disclosure Project (CDP) Data

We employ detailed data on firms' voluntary disclosures from the Carbon Disclosure Project (CDP) (https://www.cdp.net/en), an international, not-for-profit organization providing a system for companies to measure, disclose, manage, and report environmental information. CDP sends out detailed questionnaires to a large set of firms each year, and we obtain the annual responses to these questionnaires from 2011 to 2017. These data provide information rarely available in SEC-mandated 10-K annual reports, and information that is only occasionally provided by voluntary firm CSR reports.

We focus on three sets of firm disclosures in these questionnaires, namely, (i) firms' self-reported measures of their current carbon emissions (henceforth referred to as their *actions*), (ii) firms' forecasts of the future impact of environmental regulation on their operations (henceforth referred to as their *beliefs*), and (iii) firms' self-reported targets for future emissions reductions (henceforth referred to as their *plans*). We describe how we convert the raw data from CDP into the specific measures that we use in our empirical analysis later in this section, but first describe the construction of our sample below.

While it does provide detailed information on firms' environmental activity, we should mention here that the CDP dataset does have several limitations. First, firms self-report to CDP, meaning that the data comprise a selected subsample of the CRSP COMPUSTAT universe (see, for example, Luo et al. (2012)). More specifically, firms in the dataset are substantially larger than the average firm in the universe. While this does introduce concerns about external validity, it is worth noting that these firms comprise a substantial fraction (25%) of the total emissions reported in the US. Second, since the information reported in CDP is voluntary and not subject to third party auditing, it is potentially subject to "greenwashing".[10] We are therefore careful to assess the validity of the disclosures in CDP on firms' carbon footprint, their beliefs about the expected impact of regulation, and their reported plans for future abatement using a range of internal and external data. This includes three different datasets (Bloomberg, Thomson Reuters, and MSCI) of third-party verified indicators of firms' sustainability collected from publicly available sources.

2.2 Sample Construction

We match the CDP data to the CRSP COMPUSTAT North America merged database, which is a panel of 5,691 public firms reporting data over the 2010–2016 accounting period.[11] To ensure that we can measure firms' changing actions and revisions of their beliefs about regulatory risks, we require that firms in CDP report *both* current carbon emissions and their forecasts of the future impacts of regulation for at least two consecutive years in the dataset. Firms also have the option of self-reporting their targets for future emissions reductions (i.e., their plans), resulting in firms that reported as well as those that did not report plans in the *previous year*, a distinction that we return to during our analysis of the data. When we match the CDP data to the CRSP COMPUSTAT sample after applying these filters, the sample comprises a total of 446 unique North American public firms, with between 229 and 375 firms reporting in any given year between 2011 and 2017.

The top panel of Figure 1 shows, in the last reporting year 2017 in our data, the fraction of

[10] Greenwashing is the use of marketing to portray an organization's products, activities or policies as environmentally friendly when they are not.

[11] We keep only firms in the CRSP/COMPUSTAT North America (Fundamental Annual) dataset with non-missing Tickers and Total Assets within the 2010–2016 accounting period. We lag the information from CRSP/COMPUSTAT by one year to account for the time window between the filling and final release of the CDP questionnaires.

Figure 1
Sector Composition and Market Capitalization

Summary statistics of the CRSP/COMPUSTAT North America universe and the CDP subsample, as of 2017. The top histogram summarizes the proportion of CDP firms in the CRSP/COMPUSTAT North America universe at the GICS two digit level, the bottom histogram summarizes the proportion of total market value (MKVALT from CRSP/COMPUSTAT as of 2016) represented by these firms. Black (red) bars refer to the total of CDP firms (subset of CDP firms that disclose plans for at least one previous reporting period) in our sample.

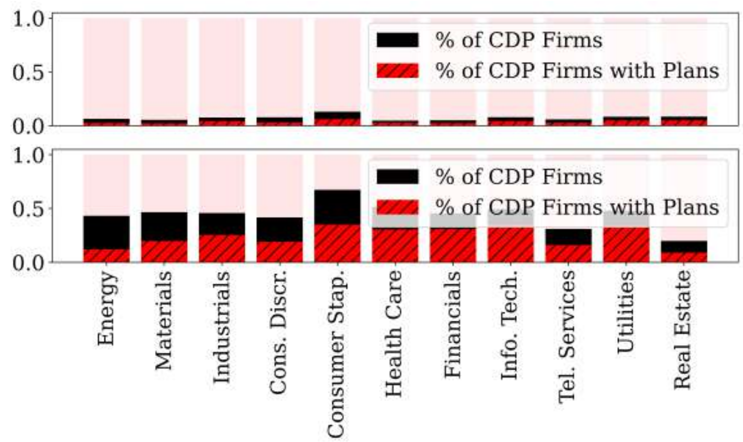

firms in the CRSP COMPUSTAT North America universe that are in our final merged sample of firms. Each bar represents a broad GICS industry. The fractions of firms reporting and not reporting future emissions reductions plans are represented in red and black respectively. Relative to the CRSP COMPUSTAT universe, there are more firms in the merged sample in Consumer Staples, Materials, Utilities, and Real Estate, and fewer Financial and Health Care firms, though these differences are not substantial. Firms that report plans for future emissions reduction are overrepresented in Utilities and Real Estate, though a roughly similar number of firms report and do not report plans in each industry.

The bottom panel of the figure shows that despite the number of firms in the left panel comprising less than 15% of the total *number* of firms, the firms in the merged sample account for 20% to 60% of the total *market capitalization* across all industries, meaning that firms that report to CDP are substantially larger than the average firm in the universe. It is also worth noting here that in 2017, the sample firms emit a total of 1,910 million metric tonnes CO_2e, which represents over 25% of the total emissions produced in the United States in 2017.[12]

[12]See *https://www.epa.gov/ghgemissions.*

Table 1
Financial and Sustainability Indicators: Summary Statistics

Summary statistics (mean and 95^{th} percentile) of the CRSP/COMPUSTAT North America universe compared with the CDP subsample over the 2010–2016 accounting period. The column Plan (No Plan) refers to the subset of CDP firms that disclose plans (do not disclose plans) in the previous reporting period. Market Value (MKVALT), Total Assets (AT), Total Liabilities (LT) and Income Before Extraordinary Items (IB) are provided by CRSP/COMPUSTAT. Return on Operating Assets (ROA) is computed as Income/(Total Assets - Total Liabilities), expressed in percentage terms. Weighted Average Cost of Capital (WACC) and Altman Z-Score are built-in functions provided by Bloomberg EquitY. Environmental, Social and Governance (ESG) disclosure scores are provided by Bloomberg ESG Data Service (1), Asset 4 ESG (2), and MSCI (3) respectively. Emissions are collected from CDP disclosures (as detailed in the data construction Appendix C). Emissions intensity is computed as Emissions/Total Assets, expressed in $\frac{\text{mtCO}_2\text{e ml}}{\$ \text{ bn}}$. All variables are collected at the annual level.* indicates that the variable has been winsorized between the 1^{st} and the 99^{th} percentiles of the pooled distribution. $^+$ indicates that the variable is available for a subset of the entire sample.

	CDP	Plan	No Plan	CRSP/COMPUSTAT	
Variable	**Mean**	**Mean**	**Mean**	**Mean**	**95^{th} perc.**
Market Value* ($ bn)	19.0	21.2	17.3	2.5	11.9
Total Assets* ($ bn)	43.5	51.6	37.6	8.9	36.6
Total Liabilities* ($ bn)	30.4	36.4	26.2	6.3	25.2
Income B. E. Items* ($ bn)	1.3	1.5	1.1	0.2	1.2
Liabilities to Assets Ratio*	0.7	0.7	0.7	0.8	1.3
ROA*	14.0	14.8	13.4	3.4	77.4
WACC*$^+$	8.4	8.0	8.7	8.4	11.2
Altman Z-Score*$^+$	3.8	3.9	3.8	3.7	13.6
ESG Score (1)$^+$	38.2	40.3	36.8	18.3	49.6
ESG Score (2)$^+$	66.4	68.9	64.7	51.3	82.3
ESG Score (3)$^+$	4.9	5.0	4.9	4.4	6.1
Emissions* (mtCO_2e ml)	6.0	7.3	5.2	-	-
Emissions Intensity*	1.8	2.3	1.5	-	-
Unique Firms	446	262	191	5,691	

Table 1 shows pooled means of a selected set of characteristics from CRSP COMPUSTAT, Bloomberg, Thomson Reuters, and MSCI. The average firm in the merged sample (i.e., reporting to CDP) is above the 95^{th} percentile firm in the size distribution of the CRSP COMPUSTAT universe. The firms in the merged sample also have substantially higher average income than the average firm in the CRSP COMPUSTAT universe, as well as a higher Return on Operating Assets (ROA), but a similar liabilities-to-assets ratio, and a slightly lower probability of bankruptcy.[13]

Firms which report plans for future emissions reduction are on average larger, have higher income, substantially lower cost of capital,[14] and lower probability of bankruptcy than firms which do not report such plans. Moreover, plan-reporting firms have greater emissions intensity (as measured by their higher emissions-to-capital ratio) than non-plan-reporting firms. The size, performance, and emissions intensity of firms can affect their incentives to disclose emissions reduction plans, as increases in these attributes can make firms more visible, resulting in greater scrutiny and pressure to disclose.[15] To verify CDP disclosures, we also acquire, for a subset of firms, their Environmental, Social, and Governance (ESG) rating scores from three separate sources, namely, Bloomberg ESG Data Service, Thomson Reuters Asset 4 ESG, and MSCI ESG, who independently assess firms' performance on carbon emissions and environmental-related activities.[16] Bloomberg, Thomson Reuters, and MSCI report ESG scores for 30%, 17%, and 32% of the pooled CRSP–COMPUSTAT sample respectively. Coverage of CDP-reporting firms in our sample, however, is substantially higher (69% in Bloomberg, 57% in Thomson Reuters, and 90% in MSCI respectively). Interestingly, across the three providers, the externally generated ESG rating scores are not substantially higher for firms in CDP than for the average firm in the universe—this raises the possibility that

[13] As implied by the Altman (1968) Z-score, an indicator of the probability of a company entering bankruptcy within the next two years, based on financial ratios obtained from 10-k reports.

[14] Weighted Average Cost of Capital (WACC) is provided by Bloomberg Equity.

[15] Size and performance can also be related to incentives to disclose through common determinants of these variables. For example, firms in CDP have substantially higher fractions of institutional ownership than firms in the universe (82% vs 64%), and firms with plans have slightly higher fractions of institutional ownership than firms without (82 vs 81%). We also find that firms with plans have a significantly greater fraction of ownership by ESG funds than firms without plans (Appendix A Table 5). Institutional ownership has been associated with higher firm value (e.g., McConnell and Servaes (1990)) and ESG funds ownership with pressures for firms to consider environmental issues (e.g., Hoepner et al. (2018) and Dyck et al. (2018)).

[16] Despite multiple controversies on ESG rating methodologies (see, for example, Christensen et al. (2022)), we find that these three different ESG metrics are positively correlated in our sample (74% correlation between Thomson Reuters and Bloomberg, 24% correlation between MSCI and Bloomberg, 34% correlation between MSCI and Thomson Reuters). These providers also make a range of environmental specific indicators available—such as the Emissions Reduction score and the total carbon footprint—which we later use in our analysis.

a certain degree of "greenwashing" might motivate firms to report. We are careful, therefore, to consider this factor, and to attempt to validate the CDP data along the dimensions in which we are interested, as we describe more fully below.

2.3 Firms' Actions, Beliefs, and Plans

In this section, we discuss how we use the CDP data to construct three measures that summarize important dimensions in the context of climate risk mitigation, namely, firms' climate mitigation *actions* to date, reflected in their actual changes in carbon footprints; their *beliefs* about the risk of climate-related regulation; and finally, their *plans* for future carbon footprint mitigation activities. We begin by describing the measures that we construct, and discuss how we validate these metrics using a range of internal and external data, including third-party verified indicators of firms' sustainability collected using publicly available sources. Then, we show that firms' plans help to predict their subsequent actions, and we uncover interesting variation along both belief and action dimensions, which we subsequently attempt to rationalize using a theoretical model.

2.3.1 Actions

We measure a firm's "abatement action" in each year as the annual change in its reported carbon emissions. Specifically, we define firm i's *abatement rate* at time t as the percentage difference between time t and time $t+1$'s emissions as:

$$x_{i,t} = -\left(\frac{Emissions_{i,t+1} - Emissions_{i,t}}{Emissions_{i,t}}\right), \tag{1}$$

where the variable $Emissions_{i,t}$ measures firm i's direct emissions from production (scope 1) as well as indirect emissions from consumption of purchased energy (scope 2),[17] as collected from CDP in reporting year t. We exclude from the study other self-reported indirect emissions from the production of purchased materials, product use etc. (scope 3) as the disclosure quality is low (see, for example, Bolton and Kacperczyk (2021a)). In Figure 1 in the Appendix A we plot carbon emissions disclosures in CDP against third-party estimates provided by Thomson Reuters—to summarize, we obtain consistent figures across the two datasets for the majority of firms in the

[17] Disclosures of carbon emissions in CDP follow the Greenhouse Gas Protocol Corporate Standard classification.

sample.[18]

2.3.2 Beliefs

In CDP, firms are queried about their exposures to three broad types of risks. The first is risk arising from likely changes in the physical climate, the second is risk arising from future environmental/greenhouse gas emissions regulation, and the third is risk arising from changes in consumer tastes and social/macroeconomic conditions. We focus on the second type of risk given our interest in the responses of firms to climate regulation events. In CDP, almost 90% of the reporting firms state that they associate climate regulation events with an increase in their operational costs, which in turn may lead to a reduced capacity to conduct "business as usual" operations.

In each reporting year t, firms provide the following pieces of information about the expected impact of a future climate regulation event:

1. An horizon H at which the environmental regulation event is expected to occur.
2. The likelihood of the event q occurring, ranging between *exceptionally unlikely, very unlikely, unlikely, about as likely as not, more likely than not, likely, very likely, virtually certain.*
3. The expected magnitude of the impact of the event M, which ranges between *low, low-medium, medium, medium-high,* and *high*, to which we assign values 1, 2, 3, 4, and 5 respectively.

To convert these reported data to a measure of beliefs, we define the expected discounted impact of the regulation event reported by firm i in year t as:

$$\Lambda_{i,t} = \beta_{i,t}^{(H_{i,t}-t)} M_{i,t} q_{i,t}. \tag{2}$$

In equation (2), $\beta_{i,t}$ is the firm's discount rate, which is the weighted average cost of capital of the firm.[19] In Figure 3 in Appendix A, we show the frequency of responses of Λ_{it} at each horizon

[18]For example, in 2017, we are able to match a total of 154 firms out of the 368 firms to the Asset 4 ESG dataset. These firms are spread across sectors. For roughly 85% of these matched firms, we find perfect matches between the two datasets, or discrepancies below 10% of the Asset 4 ESG value. For the remaining observations, CDP disclosures are lower than the Asset 4 ESG estimates, especially in pollution intensive sectors such as Energy and Utility.

[19]This is the Weighted Average Cost of Capital (WACC) from Bloomberg Equity. When not available for the single firm, we take the average WACC of the CDP sample in the same accounting year.

H_{it}, and the average expected impact (i.e., the $t-$pooled cross-sectional average of M_{it}) reported over the 2011 to 2017 period. The plot shows that the reported event horizon H_{it} ranges between zero years and over ten years from the date of reporting, and varies considerably across firms. Moreover, the expected impact of the event Λ_{it} increases, on average, with the time horizon of the event T_{it}. In Appendix A Table 2, we also regress Λ_{it} on firms' current carbon footprint and current market value, as well as a set of dummy variables to soak up industry, time, and firm headquarter-specific variation. We find that firms' self-reported beliefs about the future risks of climate regulation increase significantly with their current carbon footprint, though they decrease with firm size, controlling for the level of emissions.

In addition to these more structured quantitative assessments, firms also report unstructured text about the *specific form* of climate regulation that they expect. This text information varies with firms' location and industry, as well as varying across time. We show a word cloud of these unstructured text disclosures in Figure 2 in Appendix A. Firms' two most frequently stated types of anticipated climate regulation are, as one might expect, i) a fossil-fuel energy tax, and ii) a carbon tax/levy, generally associated with a cap and trade system. Firms also refer to mandatory emissions reporting programmes as a third category of potential climate regulation. These text disclosures partly motivate our modelling choice, described later, of regulation in the form of a carbon levy.

2.3.3 Plans

We use firms' self-reported emissions reduction targets to measure their plans for future emissions abatement. As noted earlier, some firms report these targets, while others do not. Firms that do report provide the following information in each year t:

1. A maturity T by or before which the target is planned to be achieved.

2. The total percentage of carbon emissions in year t that the firm plans to reduce between year t and the target year T, which we denote as $\hat{x}$.

We assume a constant emissions reduction rate between each reporting year t and the stated target year T, which gives us a present discounted abatement stream (i.e., a *plan* for abatement) for each firm i:

$$plan_{i,t} = \sum_{\tau=t+1}^{T_{it}} \beta_{i,t}^{\tau-t}(\frac{\hat{x}_{i,t}}{T_{i,t}-t}), \tag{3}$$

where the first timing of abatement $\tau = t+1$ refers to one year after the year of reporting.[20] In Figure 4 in the Appendix A, we plot the various reported components of the abatement plan in equation (3). The most frequently reported target horizon is between 1 and 5 years, though some firms report far longer horizons, up to 25 years ahead. As before, the longer the stated horizon, the greater the reported $\hat{x}$, on average across firms and reporting years.

In Appendix A Figure 5 we externally validate these estimates. We do so by once again relying on the subset of reporting firms that are also tracked by Thomson Reuters in their Asset 4 ESG dataset, as well as by MSCI in their MSCI ESG dataset. We plot the environmental score that feeds into the ESG rating (a measure of firms' environmental commitment) in Thomson Reuters and MSCI against our measured $plan_{i,t}$, and find a strong positive relationship between our measure and these two ratings.

2.4 Patterns in Firms' Actions, Beliefs, and Plans

Figure 2 plots the beliefs and actions of firms across our sample period, as well as 95 percentile confidence intervals across reporting years in the dataset. The left-hand panel of the figure plots conditional averages of beliefs in each reporting period, i.e.:

$$\bar{\Lambda}_t^p = \frac{1}{N_t^p}\sum_{i=1}^{N_t^p}\Lambda_{i,t}, \quad \bar{\Lambda}_t^{np} = \frac{1}{N_t - N_t^p}\sum_{i=1}^{N_t-N_t^p}\Lambda_{i,t}, \tag{4}$$

where N_t^p is the number of plan-reporting firms, and N_t is the total number of firms in reporting year t.[21] The right-hand panel plots firms' actions (with notation as above):

$$x_t^p = \frac{1}{N_t^p}\sum_{i=1}^{N_t^p}x_{i,t}, \quad x_t^{np} = \frac{1}{N_t - N_t^p}\sum_{i=1}^{N_t-N_t^p}x_{i,t}. \tag{5}$$

[20] It is worth noting that CDP questionnaires are released in October of each reporting year, while firms' responses are submitted in June or July of the same year, with exceptions of later submissions. Planned emissions reduction, as reported from firms in the second-half of the year, refer to the year ahead onwards.

[21] In Figure 9 in the Appendix A, we repeat this exercise using a weighted average of firms' beliefs, plans, and actions, where the weights vary with the emission intensity of the firms in CDP, as described in Table 1.

Figure 2

Beliefs, Plans, and Actions

The left plot shows the belief metric as in (2) against reporting years in the CDP questionnaires. The red (black) line refers to firms that disclose (do not disclose) plans in the previous reporting year (initial beliefs are set equal to the average belief observed in the next available year, e.g. in 2012). The right plot shows abatement rates and plans as in (1) and (3) respectively against reporting years in the CDP questionnaires. The red (black) line refers to actions (i.e., abatement rates) for firms that disclose (do not disclose respectively) plans in the previous reporting year. Actions are normalized between the 5^{th} and 95^{th} percentiles of the pooled distribution. The red thin line at the top of the right-hand panel shows previous year plans for emissions abatement.

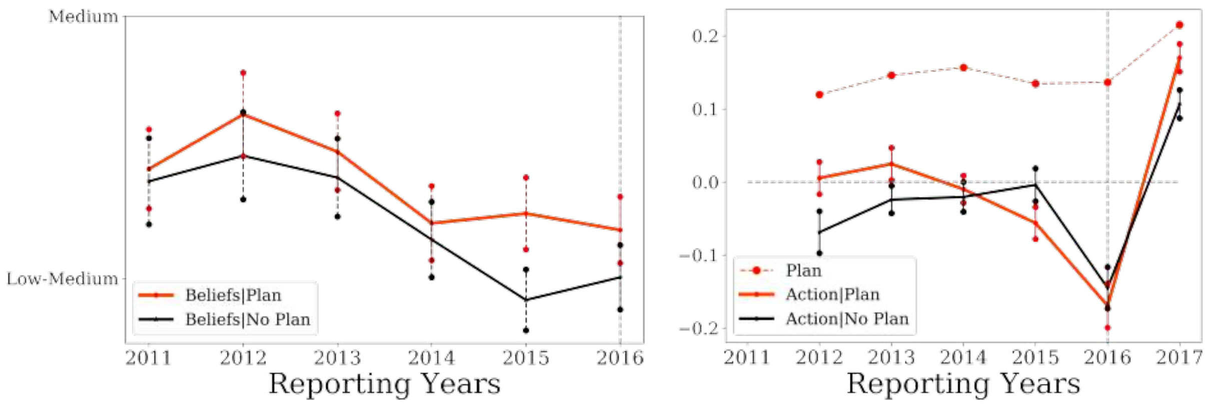

In each plot, the firms that report plans are displayed in red, and those that do not report plans are displayed in black. In the right-hand plot, we also show a thin red line, which plots the average planned abatement rate, i.e., $plan_t = \frac{1}{N_t}\sum_{i=1}^{N_t} plan_{i,t}$ for those firms that report plans.[22]

As described earlier, we construct our measure of beliefs using firms' qualitative responses about the expected impact of future climate regulation. Observed beliefs exhibit a similar and moderately decreasing trend up to 2015 for all firms (from medium (numerical value 3) to low-medium (2) expected impact across reporting years).[23] However in 2015, the year ahead of the Paris agreement, revisions in beliefs differ statistically across firms with and without plans for future abatement. In this year, plan-reporting firms revise their beliefs upwards, whereas firms without plans downward-update their beliefs about future climate regulation. This difference reverses following the Paris agreement, when firms without plans revise their beliefs up about the expected impact of climate

[22]At this stage, we ignore the distinction between the size of the emissions reduction that firms plan, and the horizon over which they choose to implement this emissions reduction. We conflate the two into the planned abatement rate in what follows.

[23]In Figure 10 in the Appendix A, we also show that such a decreasing trend is common to both components of the beliefs measure (e.g., likelihood and magnitude of future regulation), meaning that firms are downward-adjusting both the likelihood of the regulation as well as its expected impact across reporting years.

regulation, whereas firms with plans moderately revise their beliefs down.

In Appendix A Figure 6, we compare the reported belief revisions with the dynamics of global crude oil prices (West Texas Intermediate (WTI) Cushing). Interestingly, the belief revisions of non-plan-reporting firms seem to follow the dynamics of crude-oil prices across the entire observation period. In contrast, plan-reporting firms' belief revisions diverge from crude oil prices from the year prior to the Paris agreement. This evidence, although merely descriptive, seems to suggest that firms with plans anticipate the international agreement, in contrast with non-plan reporting firms.

The right-hand panel of Figure 2 shows how the current actions of firms on emissions reduction vary over time, once again splitting firms into two groups based on whether they do or do not report plans for future emissions reduction.[24] The plots show patterns similar to the dynamics of beliefs—both groups of firms increased their emissions between 2012 and 2016, leading up to the Paris climate change agreement. Perhaps surprisingly given their reported beliefs, firms with plans reduced their abatement activities *more* than firms without plans over this period. Once the Paris agreement is ratified, however, both groups sharply reduce their emissions, i.e., increase abatement activities, in 2017. And again, firms with plans increase abatement activities *more* than firms without reported plans for future emissions reduction.

The plans themselves are plotted as a thin red line in the right-hand panel of the figure. The expected future abatement rate remained steady until 2015, but rose significantly in 2016, predicting the realized spike in emissions reduction in 2017. Figure 11 in Appendix A shows that predicted and realized emissions reductions persist in disaggregated firm disclosures at the sector level, though there is variation across sectors. This bolsters the case that the spike observed in the data is a reaction to a global shock of the Paris agreement announcement, rather than a sector-specific regulatory shock.

A note on robustness is in order here. As with many ESG-related datasets, CDP is an unbalanced and expanding panel with few firms at the beginning of the sample and more as time goes on. An issue with such datasets is that results could potentially be driven by composition effects—as from one year to another, firms that are potentially very different from the average sample firm

[24]Note that actions, which are percentage changes in total reported carbon emissions, are winsorized between the 5^{th} and the 95^{th} percentiles of the pooled distribution.

join the panel. For completeness, Figure 8 in Appendix A shows the average dynamics of beliefs, actions, and plans for a balanced panel of CDP firms that reported consistently in each year beginning with 2011. The sample shrinks as a result of this filter, so the magnitudes are different. Nevertheless, all of the major patterns identified in Figure 2, i.e., the decreasing trend in beliefs, the reaction to the Paris announcement, and the more pronounced reaction of firms with plans, are preserved in the restricted sample. To better understand the underlying source of these intriguing patterns, we build a dynamic model of firms' carbon emissions reduction, as we describe below.

3 A Baseline Dynamic Model of Carbon Emissions Reduction

Our modelling strategy proceeds in two steps. We first begin with a dynamic model of a single representative firm, considering its optimal abatement strategy. In a second step, to better model the heterogeneity in responses that we observe across firms with and without plans, we extend the model to a two-firm version adding strategic considerations.

3.1 Setup: Single-Firm Model

The economy exists for $t = 0, \dots, T$ time periods, and we model a single firm operating in this economy. At the beginning of each time period t, the firm operates with a stock of polluting capital k_{t-1}, producing a proportional amount of carbon emissions $\xi_{t-1} = \eta k_{t-1}$ (measured at the end of time period $t-1$). The firm can reduce or increase its emissions at a rate x_t.[25] If the firm decides to abate, the capital stock has the following law of motion:

$$k_t = k_{t-1}(1 - x_t), \tag{6}$$

with corresponding carbon emissions (measured at the end of time period t) of:

$$\xi_t = \eta k_t = \eta k_{t-1}(1 - x_t) = \xi_{t-1}(1 - x_t). \tag{7}$$

[25]Note that the polluting firm can reduce emissions by reducing production from dirty capital. An alternative option would be to replace the polluting capital with clean capital to increase energy efficiency (i.e., decrease the emission intensity η). Such an option is generally irreversible and takes time to build. We discuss the implications of introducing this possibility in Section 4.3.

At the end of time period $t = 0, \ldots T-1$, the firm makes profits π_t from its operations:

$$\pi_t = \omega k_t - \frac{1}{2}\phi x_t^2 k_{t-1}, \tag{8}$$

where ωk_t is the firm's output from a linear production function (ω is a productivity constant), and ϕ is a parameter for the quadratic adjustment cost of capital (we simply normalize the cost of incremental investment to zero).[26]

At the end of time period $t = T$, a regulation event occurs, and the firm pays a carbon levy λ for each unit of emissions it produces at that time.[27] As a result, the firm's terminal profits can be expressed as:

$$\pi_T - \lambda \xi_T. \tag{9}$$

At the beginning of the current time period $t = 0, \ldots, T-1$, the true intensity of the terminal levy λ is unknown. The firm has a belief over the levy which is normally distributed as $\lambda_t \sim \mathcal{N}(\bar{\lambda}_t, \sigma_t)$. Conditional on this belief, the firm determines the current abatement rate x_t which maximizes the Bellman equation:

$$V_t = \max_{x_t} \pi_t + \beta \mathbb{E}_t[V_{t+1}] \tag{10}$$

where β is the discount rate of the firm and the terminal value given by:

$$V_T = \max_{x_T} \pi_T - \lambda \xi_T. \tag{11}$$

3.1.1 Solving the Model

In the theory Appendix B.1, we show from the first order condition of the Bellman equation in (10), that the optimal current abatement rate x_t^* for each $t = 0, \ldots, T-1$ satisfies

$$x_t^* = \beta \mathbb{E}_t[x_{t+1}^* - \frac{1}{2}(x_{t+1}^*)^2] - \frac{\omega}{\phi}, \tag{12}$$

[26]It is simple to show that introducing a linear cost for incremental investment does not alter any of the main predictions of the model in that it enters the expression for the optimal abatement rate x_t^* as a modified (lower) productivity constant ω.

[27]In the interests of parsimony, we choose to model the carbon pricing mechanism as a tax applied to each unit of emissions produced by the firm. As mentioned in the data section, the carbon tax is one of the most frequent types of regulation explicitly mentioned by reporting firms in the data.

with the terminal abatement rate:

$$x_T^* = \frac{(\eta\lambda - \omega)}{\phi}. \tag{13}$$

It is worth reiterating here that when determining the current abatement rate by backward induction, we assume that the conditional distribution of future beliefs is constant (i.e., $\lambda_\tau | t \sim \lambda_t$ for $t < \tau < T$). This amounts to assuming that while the firm internalizes a certain degree of uncertainty about the terminal levy at time t, it does not expect to learn more about the levy up until the terminal date.[28] This implies that observed revisions in beliefs are the result of unanticipated shocks about the levy. This assumption is not fully consistent with the observed patterns (firms seem to revise beliefs in each time period and it is therefore likely that they anticipate further revisions), but we make it to reduce the number of free parameters in the model given our limited data sample (different assumptions about the anticipated dynamics of beliefs would strongly affect the abatement dynamics). Concerns about this choice are alleviated by the fact that the main mechanisms in the model are geared towards explaining observed reactions around and after the Paris agreement announcement, which can be intended as an exogenous shock to firms' beliefs about future regulation. In Appendix D.1 we discuss the implications of an alternative specification of the conditional distribution of beliefs, and show that it does not materially alter our conclusions.

3.1.2 Comparative Statics

The comparative statics of the terminal abatement rate x_T^* in (13) are intuitive. The abatement rate increases with the levy (λ) and with the parameter η, which captures the pollution intensity of the firm. On the other hand, the abatement rate decreases with the productivity of polluting capital, ω. Finally, regardless of whether the model predicts an abatement or an increase in polluting capital (i.e., regardless of whether $x_T^* > 0$ or $x_T^* < 0$), the magnitude of any abatement decreases as the adjustment cost parameter ϕ rises.

We now outline the key comparative statics of the abatement profile $\{\mathbb{E}_t[x_\tau]\}_{\tau=t,\dots T}$. For a given belief over the levy λ_t, we show in the Appendix B.1 that, for realistic values of the model

[28] Equivalently, the firm expects future signals to be very noisy compared to the precision of its current belief.

parameters, the optimal abatement profile satisfies

$$x_t^* < \mathbb{E}_t[x_{t+1}^*] < \cdots < \mathbb{E}_t[x_{T-1}^*] < \mathbb{E}_t[x_T^*], \tag{14}$$

that is, an *upward-sloping* term structure of abatement, as seen in Figure 3.

Figure 3
Optimal Abatement Profile

The plot shows the optimal abatement profile $\mathbb{E}_t[x_\tau^*]$ as a function of the maturity $\tau = t, \dots T$ for two values of the parameter $\sigma_t = 20$ (black line) and $\sigma_t = 0.0$ (blue line) respectively. Other model parameters are: $\phi = 10$, $\omega = 0.2$, $\beta = 0.9$, $T = 10$, $\eta = 0.1$, $\bar{\lambda}_t = 15$.

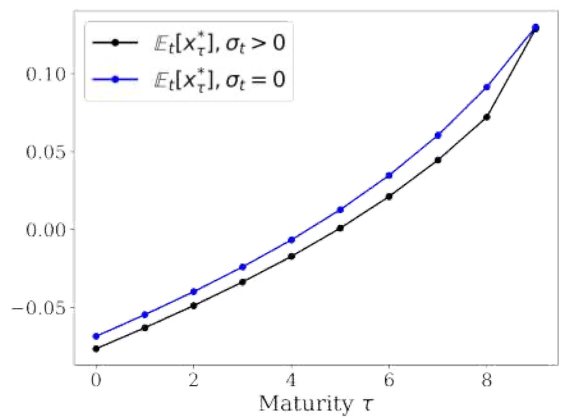

This result is intuitive: the benefits to the firm from an additional unit of polluting capital (the productivity constant ω) accrue at the time at which the capital is in place (i.e., any time t before and including the terminal date), while the social cost (the levy λ) is always incurred at the terminal date, and hence always discounted more heavily than the benefits. This gap between the present value of costs and benefits shrinks as we approach the terminal date, resulting in the upward-sloping abatement term structure.

For a given maturity $\tau = t, \dots T-1$, we show in the Appendix B.1 that the expected abatement rate satisfies

$$\frac{\partial \mathbb{E}_t[x_\tau^*]}{\partial \sigma_t^2} < 0, \tag{15}$$

that is, the expected abatement rate decreases for higher values of the levy's uncertainty σ_t^2. This because emissions abatement costs are convex,[29] hence the value function is concave in emissions abatement, and therefore also in the terminal levy.

To summarize, this simple first model predicts an upward-sloping term structure of planned abatement for any realization of the levy (i.e., abatement rates increase up to levy imposition), as the costs of abatement are incurred in the present, but the levy is only incurred at the terminal date, meaning that its impact is diminished by discounting at any intermediate date. Moreover, the model predicts that everything else equal, abatement should be lower on average for higher values of uncertainty about the terminal levy (Figure 3). This latest result relates to recent work by Barnett et al. (2020) which quantifies the impact of uncertainty on the social cost of carbon (SCC), showing that risk-adjusted SCCs are significantly lower than risk-neutral estimates.

3.1.3 Single Firm Model Calibration

We calibrate the single firm model on the representative firm that reports a plan (denoted as firm p) in the selected CDP dataset. We summarize the measurement choices below, and report the corresponding parameters' values in Panel A, Table 2. Further details on the measurement of the parameters can be found in Appendix D.

The discount rate, $\hat{\beta}^p$, is measured using the weighted average cost of capital from Bloomberg Equity; the pollution intensity, $\hat{\eta}^p$, is measured using the ratio of total emissions to total assets from CDP and CRSP/Compustat respectively; the productivity constant, $\hat{\omega}^p$, is measured using the net income to total assets ratio from CRSP/Compustat; the adjustment cost of capital $\hat{\phi}^p$, is borrowed from Liu et al. (2009);[30] the regulation event, T, is either set equal to $\hat{T} = 2020$, with 2020 is the most frequent target year reported by firm p since the Paris agreement announcement, or rolls over reporting years $\hat{T}_t = t + \hat{h}$ for $t = 2011, \ldots 2016$, with $\hat{h}$ the average target horizon reported by firm p in the selected dataset. For each reporting year $t = \{2011, \ldots 2016\}$, firm p's

[29] This result holds true if two conditions are satisfied. First, the firm must abate at least some capital in order to control its emissions, and second, abatement of capital must involve convex adjustment costs—these conditions together imply that emissions abatement has convex costs.

[30] The values refer to the first estimate obtained in Liu et al. (2009), who use the q-theory of investment to derive (and test) moments on the cross-section of stock returns. Previous estimates of the adjustment cost of capital obtained by directly matching the observed and model-implied moments of investment range from over 20 to as low as 3.

belief about the levy, expressed in $/mtCO$_2$e, are specified as a linear function of reported beliefs

$$\lambda_t^p = \bar{\lambda}^p(1 + \hat{\alpha}(\hat{\Lambda}_t^p - \bar{\Lambda}^p)) \tag{16}$$

where $\hat{\Lambda}_t^p$ is the belief of the representative firm with plans, which we construct from the distribution of reported beliefs assuming $\hat{\Lambda}_t^p \approx \mathcal{N}(\bar{\Lambda}_t^p, \Sigma_t^p)$ with $\bar{\Lambda}_t^p$ the average belief in (4) and Σ_t^p the standard deviation of the beliefs reported by the group of firms with plans.[31] The constant $\bar{\Lambda}^p$ is the average belief reported by the firm with plans across the entire observation period in CDP, whereas $\alpha = 5$ is a scale parameter which accounts for the fact that beliefs in (2) are extracted from categorical disclosures which range between 0 and 5, whereas the levy intensity is continuous.

Denoting the set of input parameters $\Theta^p = \{\hat{\beta}^p, \hat{\omega}^p, \hat{\eta}^p, \hat{\phi}^p, \hat{\Lambda}_t^p, \hat{\alpha}\}$, we estimate the average belief of firm p, $\bar{\lambda}^p$, to minimize the squared distance between the empirical and model-implied abatement actions and abatement plans:

$$\min_{\bar{\lambda}^p} \sum_{t=2011}^{2016} \left(x_t^p - x_t^*(\Theta^p, \bar{\lambda}^p)\right)^2 + \left(plan_t - plan_t^*(\Theta^p, \bar{\lambda}^p)\right)^2 \tag{17}$$

where the optimal plan is the sum of expected future discounted optimal abatement rates, i.e., $plan_t^* = \sum_{\tau=t+1}^{T_t}(\hat{\beta}^p)^{\tau-t}\mathbb{E}_t[x_\tau^*]$, and x_t^* is the optimal current abatement rate.[32] It is worth recalling that, from the specification of the firm's emissions in (7) and the capital stock dynamics in (6), we have that $x_t^* = -\left(\frac{\xi_{t+1}-\xi_t}{\xi_t}\right)$, which allows for a direct comparison with the relative change in realized emissions, x_t^p, as measured in (5). In the same way, the model-implied abatement plan ($plan_t^*$) also allows for a direct comparison with the measured abatement plan $plan_t$ in (3), reported by the representative firm with plans at year t and anticipating relative changes in emissions from year $t+1$ onwards.

Calibration results for the two specifications of the regulation event are reported in the first and second column of Table 2, Panel B. Looking at the case of the fixed maturity (first column), the

[31] This in turn means that the average belief at each time t is $\bar{\lambda}_t^p = \bar{\lambda}^p + \alpha\bar{\lambda}^p(\bar{\Lambda}_t^p - \bar{\Lambda}^p)$ whereas the standard deviation is $\sigma_t^p = \alpha\bar{\lambda}^p\Sigma_t^p$.

[32] We impose full consistency between reported plans and actions in CDP, simply excluding the possibility of cheap talk equilibria in our setting (see, for example, Hämäläinen and Leppänen (2017). More specifically, we assume that the firm can only truthfully report its abatement plan in CDP. One way to justify this choice is to assume that, as in reality, the informational quality of the announcement is subject to a high degree of third-party scrutiny.

plan-disclosing firms' actions are consistent with an average belief over the levy of roughly $\bar{\lambda}^p = 80$ $/mtCO$_2$e. Despite the growing literature dedicated to the topic, there is still large uncertainty about the social cost of carbon (SCC), and the economic implications of carbon policies (see, for example, Nordhaus (2014)). In a review by Tol (2011), the average of SCC estimates across over 300 published articles is over 150$/mtCO$_2$e, while the mode of the distribution is below 50$/mtCO$_2$e. Our levy-implied estimate of the SCC, i.e, assuming a Pigouvian levy that equates the average marginal cost of one additional unit of carbon emissions, falls into the range provided by Tol (2011), while is well above the price of carbon permits traded in the EU ETS during the period of observation (roughly 5$/mtCO$_2$e).[33] To the extent that our exercise is credible, the policy implication is that a substantial rise in the global price of carbon might not come as a shock to firms given these implied priors.

Figure 4
Model-Implied and Observed Moments - Single Firm

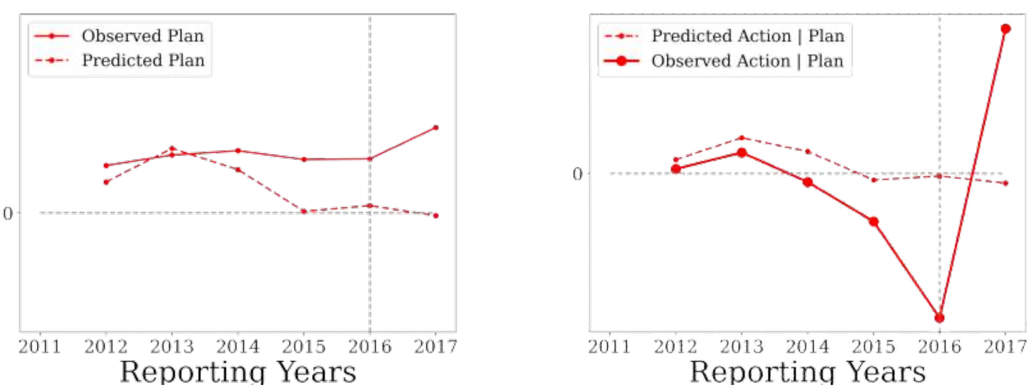

The left plot compares the model-implied and observed abatement plan of the firm with plans (firm p) against reporting years in CDP. The right plot compares the model-implied and observed abatement rate of the firm with plans (in red) and firm without plans (in black) against reporting years in CDP. Thick (dashed) lines refer to observed (model-implied) moments. Model parameters in input are reported in Table 2, Panel A. Estimated parameters in Table 2, Panel B, first column.

In Figure 4 we report empirical and model-implied abatement plans, as well as model-implied actions and observed abatement actions on emissions reduction.[34] In both plots, the model-implied moments are dashed lines, while the solid lines show the patterns in the data. The figure shows

[33]Information on the current pricing of carbon in the EU emissions trading scheme (EU ETS) can be found at https://climate.ec.europa.eu/eu-action/eu-emissions-trading-system-eu-ets_en.

[34]Parameters in input refer to the case of fixed maturity, i.e., first column of Panel B, Table 2. The case of rolling maturity is reported in Figure 18 in Appendix A.

that the model captures the dynamics of plans and actions reasonably well up to the Paris agreement announcement, although there is an issue of magnitude, which might be expected given the simplicity of the model. However, the figure shows that model performs poorly around and after the Paris agreement announcement. To attempt to better capture the relationship between beliefs and actions in the dataset, as well as to capture differences in reactions across the firms with and without plans, we therefore move to a model with two firms, which we describe in the next section.

4 A Leader-Follower Model of Carbon Emissions Reduction

We introduce a second firm into the model, and denote one firm by l (for *leader*) and the other by f for (*follower*). We assume that firm l and firm f have heterogeneous primitives $\{\omega^l, \phi^l, \eta^l, \beta^l\}$ and $\{\omega^f, \phi^f, \eta^f, \beta^f\}$, as well as heterogeneous beliefs over the levy $\{\lambda_t^l, \lambda_t^f\}$ respectively.

In each time period $t = 0, \dots T$, we augment the baseline profit function with a payoff externality that makes firm l and firm f's profits depend symmetrically on the other firm's actions

$$\begin{aligned} \pi_t^l(x_t^f) &= \omega^l k_t^l - \frac{1}{2}\phi(x_t^l)^2 k_{t-1}^l - \gamma_t^l x_t^f(\xi_t^l - \xi_{t-1}^l), \\ \pi_t^l(x_t^f) &= \omega^f k_t^f - \frac{1}{2}\phi(x_t^f)^2 k_{t-1}^f - \gamma_t^f x_t^l(\xi_t^f - \xi_{t-1}^f), \end{aligned} \tag{18}$$

when γ_t^l and γ_t^f are positive, this can be interpreted as a reputation externality, in that a firm's profits are reduced in any period t in which the *other* firm abates emissions, and vice-versa. The term can also be thought of the degree of attention paid by society to firms' abatement activity, manifested in relative performance evaluations along this dimension. In the years leading to the Paris agreement event, the dynamics of media attention to climate change issues has been increasing, along with attention paid to firms' ESG scores. Some evidence to support this assumption can be seen in Figure 5, which documents the number of articles in Dow Jones newswire on selected keywords, but also directly from firms' unstructured disclosures about climate change risks.[35]

In what follows, we derive l and f's optimal abatement profiles in a Stackelberg leadership equilibrium with heterogeneous beliefs. In each time period t, firm l (the firm reporting its plans) moves first, rationally anticipating actions and beliefs of firm f (the firm not reporting its plans),

[35] For example: "Failure to meet investors' expectations [..] could result in a risk to corporate reputation, with incremental financial impact given the expanding role of Environmental, Social and Governance (ESG) issues in evaluations."

Figure 5
Historical Environmental Media Coverage

The figure shows the time-series of the number of Dow Jones articles containing the words *"Climate Change"* (black dotted line) and *"ESG"* (red dotted line) in headlines or lead paragraphs as recorded from the Factiva database between 2000 and 2020.

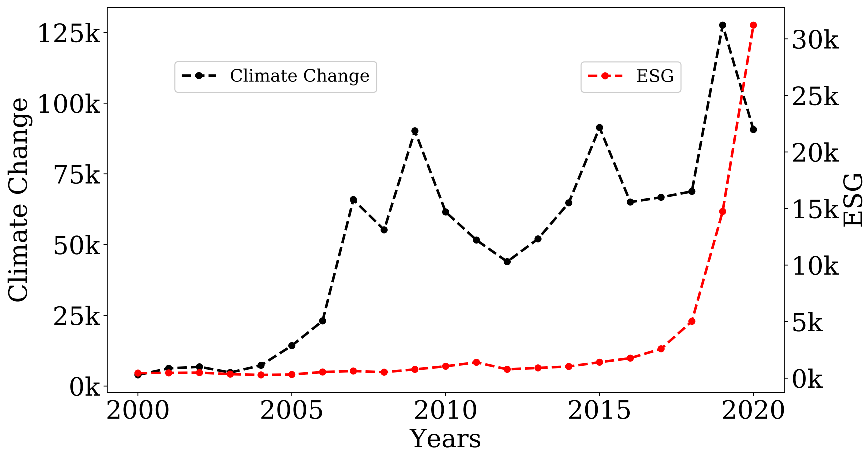

while the latter takes the abatement choices of the competitor as given, without updating its prior beliefs about climate regulation. We summarize a set of arguments in support of this setup (i.e., 1) heterogeneous beliefs and 2) a Stackelberg leadership equilibrium) below.

The leader and follower firms might be exposed to different terminal levies, i.e., λ^f and λ^l, implying that it is rational for the firm without plans to mantain an independent belief over the levy after observing that of the leader. This assumption of heterogeneous beliefs helps to rationalize observed differences in belief dynamics across the two types of firms discussed in Section 2. More specifically, while firms without plans report beliefs that broadly track global crude oil prices, plan-reporting firms' beliefs diverge around the Paris agreement announcement. Finally, we also later demonstrate that moving away from a setting with heterogeneous beliefs to a Stackelberg leadership equilibrium where firm f (the follower) learns about the levy from the leader's abatement choices does a worse job in matching firms' reactions around and after the Paris agreement announcement.

As far as the leadership equilibrium is concerned, we show in Table 1 that firms with plans in our dataset are larger and more profitable, on average, than firms with no plans. Furthermore, Appendix A Table 4 shows that plan-reporting firms have a greater propensity to a) engage with policymakers and b) provide direct funding to regulatory activities. This greater proximity to climate regulation (which is also consistent with firms' own disclosures of climate regulation risk) supports the assumption that firms that report plans act as leaders in the "carbon abatement market"(see Ovtchinnikov et al. (2019), Zhang et al. (2019) and Heitz et al. (2021) respectively). Additional important evidence supporting the assumption of firms with plans acting as "carbon abatement leaders"comes from the fixed income market. Figure 7 in Appendix A shows that firms with plans anticipate firms without plans in the issuance of green debt securities,[36] which are securities that embed incentives to innovate the production plan by switching to green technologies.[37] Finally, for completeness, we study an alternative Cournot equilibrium where firms with heterogeneous beliefs move simultaneously. Appendix B.2 shows that this alternative setting is unable to capture firm l and firm f's reactions to changes in beliefs around and after the Paris agreement announcement.

4.1 Equilibrium Abatement Profiles.

Holding fixed the model parameters $\{\phi^l, \phi^f, \beta^l, \beta^f, \omega^l, \omega^f, \eta^l, \eta^f\}$ and the maturity of the regulation event T, for any time $t = 0, \dots T$ for which firm l and firm f beliefs are $\{\lambda_t^l, \lambda_t^f\}$ and the payoff externalities $2\eta^l\gamma_t^l\eta^f\gamma_t^f < \phi^l\phi^f$, the optimal abatement profiles $x_t^{*,l}$ and $x_t^{*,f}$ satisfy:[38]

- Firm f (follower):

$$x_t^{*,f} = w_t^f x_t^{*,l} + \beta^f \mathbb{E}_t\big[x_{t+1}^{*,f} - w_{t+1}^f x_{t+1}^{*,l} - \frac{1}{2}(x_{t+1}^{*,f})^2\big] - \frac{\omega^f}{\phi^f} \tag{19}$$

with $w_t^f = \frac{\eta^f\gamma_t^f}{\phi^f}$, and

$$x_T^{*,f} = w_T^f x_T^{*,l} + \frac{\eta^f\lambda^f}{\phi^f} - \frac{\omega^f}{\phi^f} \tag{20}$$

[36]Green debt securities as classified from Bloomberg fixed income include green bonds, green loans, sustainable loans as well as sustainability-linked bonds and loans. For more information, see Barbalau and Zeni (2023).

[37]In our model, carbon abatement is uniquely achieved by reducing production from dirty capital, whereas the option to replace the dirty technology (i.e. reduce the firm's emission intensity) is not accounted for. In Section 4.3, we discuss this alternative abatement option and its implications on the model results.

[38]The upper bound on the magnitude of the strategic parameters is a requirement that we impose to get well-defined abatement plans and actions. This can be though of as a bound on the size of the reputation externality.

- Firm l (leader):

$$x_t^{*,l} = \frac{\beta^l}{1-2w_t^l w_t^f}\mathbb{E}\big[x_{t+1}^{*,l}(1-w_{t+1}^l w_{t+1}^f - \frac{\beta^f}{\beta^l}w_t^l w_{t+1}^f) + x_{t+1}^{*,f}(\frac{\beta^f}{\beta^l}w_t^l - w_{t+1}^l)\dots$$
$$\dots - \frac{1}{2}((1-2w_{t+1}^l w_{t+1}^f)(x_{t+1}^{*,l})^2\big] - \frac{\omega^l}{\phi^l(1-2w_t^l w_t^f)} - w_t^l \frac{\omega^f}{\phi^f(1-2w_t^l w_t^f)} \tag{21}$$

with $w_t^l = \frac{\eta^l \gamma_t^l}{\phi^l}$, and

$$x_T^{*,l} = \frac{1}{1-2w_T^l w_T^f}\Big(\frac{\eta^l\lambda^l}{\phi^l} - \frac{\omega^l}{\phi^l}\Big) + \frac{w_T^l}{1-2w_T^l w_T^f}\Big(\frac{\eta^f\lambda^f}{\phi^f} - \frac{\omega^f}{\phi^f}\Big) \tag{22}$$

The derivations of these expressions are in Appendix B.2.

4.2 Comparing the Single-Firm and Two-Firm Models

We now compare the equilibrium abatement rates in the previous subsection with the baseline solution established in (12) and (13). We first state the following proposition:

- **Proposition 1.** *Provided* $0 < w_T^l w_T^f < 1/2$, *the sensitivity of the leader* l*'s expected terminal abatement rate* $\mathbb{E}_t[x_T^{*,l}]$ *to changes in beliefs* $\bar{\lambda}_t^l$ *is greater than the sensitivity of the follower* $\mathbb{E}_t[x_T^{*,f}]$ *to changes in beliefs* $\bar{\lambda}_t^f$. *Moreover, both are greater than the respective sensitivities in the baseline (i.e. single-firm) model with no cross-firm payoff externalities.*

  - **Corollary 1**. *If* $w_T^l > (1+w_T^f)^{-1}$, *then the sensitivity of the leader* l*'s rate* $\mathbb{E}_t[x_T^{*,l}]$ *to changes in beliefs of the follower* $\bar{\lambda}_t^f$ *is also greater than that of the follower.*

The proof of this proposition can be found in Appendix B.2. There, we also outline a sufficient condition under which the corollary extends to current rates $x_t^{*,l}, x_t^{f,*}$.[39]

To develop intuition, we begin by discussing the second statement in proposition 1, which is easy to verify—starting from the explicit expressions in (20) and (22), one can easily derive that the belief parameter $\bar{\lambda}_t^l$ ($\bar{\lambda}_t^f$ respectively) has a higher marginal effect on $\mathbb{E}_t[x_T^{*,l}]$ ($\mathbb{E}_t[x_T^{*,f}]$ respectively) than on the baseline solution in (13). The intuition is that the cross-firm externalities make firms

[39] Due to the presence of convex adjustment costs, the corollary does not necessarily hold for all for each maturity $t = 0 \dots T-1$. However, as we show in the Appendix B.2, the corollary holds when the model parameters generate expected negative abatement in the next period, i.e. $\mathbb{E}_t[x_{t+1}^{f,*}], \mathbb{E}_t[x_{t+1}^{l*}] < 0$ in equilibrium. Importantly, such condition is almost always verified in the data.

endogenously increase their reaction to changes in the policy, because the way the model is set up in equations (18), firms have incentives *to act alike* provided that the weights w_T^f and w_T^l are positive. More specifically, when the weights are positive, firms find more costly to act such that $\mathbb{E}_t[x_T^{*,f}]\mathbb{E}_t[x_T^{*,l}] < 0$. This tendency towards similarity amplifies their actions relative to the "atomistic" optimum which is unencumbered by such externalities.

The proposition then states that this amplification mechanism is greater for the leader firm than for the follower firm. Inspecting equations (18), we can see that they bear a resemblance to the expressions that one might get from a traditional Stackelberg duopoly, with a modified "demand function of abatement".[40] Essentially, since firm profits respond to (own and other firm) abatement negatively in a similar way that price responds to demand in the traditional Stackelberg model, the leader firm has an incentive to grab "abatement market share" in a similar way to the traditional Stackelberg model, since it has a first-mover advantage.

Finally, the corollary says that if the reputation weight of the leader is sufficiently high, then it reacts *more* than the follower to variations in the beliefs of the follower, i.e. $\partial\mathbb{E}_t[x_T^{*,l}]/\partial\bar{\lambda}_t^f > \partial\mathbb{E}_t[x_T^{*,f}]/\partial\bar{\lambda}_t^f$. To derive intuition, imagine a negative shock to the follower's belief $\partial\bar{\lambda}_t^f < 0$ such that the leader's belief $\partial\bar{\lambda}_t^l = 0$ remains unchanged. *Ceterus paribus*, the model predicts that leader firm should react *more* than the follower by decreasing its abatement (i.e., increasing its emissions) at a rate that is larger than that of the follower, despite the fact that its information over the levy remains unchanged. This condition, which is essential to predict how firms' beliefs and actions are related around and after the Paris agreement announcement, cannot be generated by a simultaneous equilibrium setting unless more restrictive conditions on the weight parameters are imposed (see Appendix B.2).[41] Similarly, the condition would not hold in an equilibrium where the follower firm would incorporate information about the levy from the leader's actions (since in this case, $\partial\bar{\lambda}_t^f|\partial x_t^{*,l} = \partial\bar{\lambda}_t^l = 0$).

Finally, we introduce in the proposition an additional feature of the two-firm model which

[40]To see this, note that we can rewrite the firms' terminal profits as:

$$\pi_T^i(x_t^{-i}) \approx (\eta\lambda^i - \frac{\phi^i}{2}(x_T^i - 2w_T^i x_T^{-i}))x_T^i - \omega^i x_T^i \tag{23}$$

with $i = l, f$ and $-i = f, l$ respectively.

[41]In particular, the condition is never verified if the follower and the leader's reputation weights are equal, whereas it can still be verified in the leader-follower model.

relates the abatement term structure to changes in the time-path of the reputation externality:

Proposition 2. *For any maturity $t < \tau < T$, assume that the leader l's and the follower f's next period weights follow $w^l_{\tau+1} = \rho w^l_\tau$ and $w^f_{\tau+1} = \rho w^f_\tau$ respectively. Then, provided ρ is sufficiently small, the leader l's abatement profile satisfies $\mathbb{E}_t[x^l_\tau] > \mathbb{E}_t[x^l_{\tau+1}] > 0$. Furthermore, if the follower's weight w^f_τ is sufficiently large, then the follower's abatement profile also satisfies $\mathbb{E}_t[x^f_\tau] > \mathbb{E}_t[x^f_{\tau+1}] > 0$.*

Figure 6
Optimal Abatement Profile

The figure shows the optimal abatement profile for the leader firm (red line) in (21), the follower firm (black line) in (19) and the baseline single-firm (blue line) in (12). Model parameters are $\omega^f = \omega^l = \omega = 0.2$, $\eta^l = \eta^f = \eta = 0.3$, $\beta^l = \beta^f = \beta = 0.9$, $\phi^l = \phi^f = \phi = 10$, $\bar{\lambda}^f_t = \bar{\lambda}^f_t = \bar{\lambda}_t = 20$, $\sigma^l = \sigma^f = \sigma = 0$, $T = 8$, $t = 0$, $\gamma^l_\tau = \gamma^f_\tau = 15e^\tau 1_{\tau>3}$.

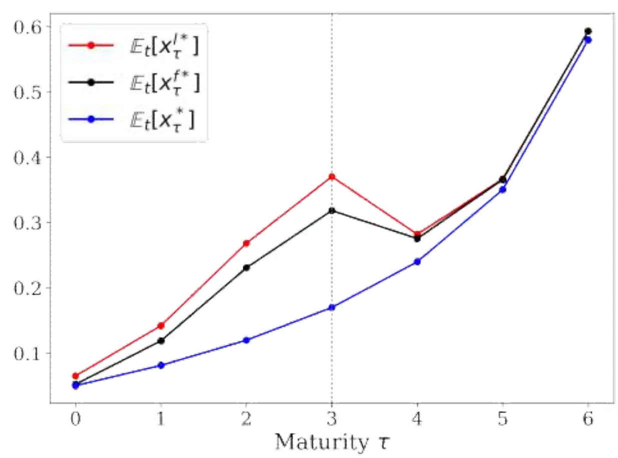

While we leave the details to Appendix B.2, Proposition 2 states that if there exists an interim maturity τ such that the reputation weights at τ are sufficiently large and decrease quickly after τ,[42] then the equilibrium plans can support a non-monotonic term-structure of abatement, i.e., abatement can peak at the interim maturity τ rather than increase until the terminal date T, as in the baseline model (Figure 6).[43] This is because a decreasing time-path of the reputational externality introduces an additional cost associated with carbon emissions that accrues more aggressively at the (current) time at which the capital is in place.

[42]For example, a sudden increase in attention to sustainability practices which gradually reverts back to the mean.
[43]Importantly, as we show in the Appendix B.2, the abatement profile of the follower can only be inverted if the one of the leader is inverted (and if the reputation weight of the follower is large enough).

4.2.1 Two-Firm Model Calibration

We conclude this section by calibrating the two-firm model to the data. We use the same set of parameters Θ^p and repeat the same exercise on the representative firm without plans (firm np) to obtain Θ^{np} (reported in the calibration Appendix D). In each reporting year $t = 2011, \dots 2016$, we then specify the sign and magnitude of the payoff externality assuming a functional form

$$\gamma_t^p = \gamma^p(ESG_t), \quad \gamma_t^{np} = \gamma^{np}(ESG_t) \tag{24}$$

where the term ESG_t measures the market attention to firms' sustainability practices as captured by the normalized number of Dow Jones articles containing the words *ESG* (Figure 5). We estimate the average beliefs $\bar{\lambda}^p, \bar{\lambda}^{np}$ as well as the strategic scale parameters γ^p, γ^{np} to match average abatement plans and actions of the leader as well as actions of the follower firm.[44]

The estimated parameters are reported in the third and fourth column of Table 2, Panel B. The matched moments are reported in Panel C. Looking at the case of fixed maturity (third column in Table 2, Panel B), the average belief consistent with the plan-reporting firm's actions and plans is $\bar{\lambda}^p \approx 120\$/\text{mtCO}_2\text{e}$, while that of the firm without plans is $\bar{\lambda}^{np} \approx 126\$/\text{mtCO}_2\text{e}$. The strategic scale parameters $\gamma^p \approx 1,356\$/\text{mtCO}_2\text{e} > \gamma^{np} \approx 986\$/\text{mtCO}_2\text{e}$ verify the condition in Corollary 1 around and after the Paris announcement.[45] Taking as reference the model-implied average abatement rates for the firm with and without plans in Panel C (i.e., 0.015 and 0.011 respectively), the model estimates that the reputation externality introduces a benefit from abating one unit of emissions that averages around $1,356\$/\text{mtCO}_2\text{e} * 0.011 \approx 15\$/\text{mtCO}_2\text{e}$ for firm p and $986\$/\text{mtCO}_2\text{e} * 0.015 \approx 15\$/\text{mtCO}_2\text{e}$ for firm np respectively (so equal across the two firms). One notes that these benefits account, on average, for roughly 15% of the costs introduced by the terminal levy, but are incurred at each time period t instead of only at the terminal date T.

Figure 7 plots the results of the calibration with fixed maturity (i.e., third column):[46] the left and right-hand panels show that the more complicated two-firm model with cross-firm externalities

[44] The moment-matching exercise is reported in more details in the calibration Appendix D.

[45] The condition is plausible, as these firms are larger, more emission intensive, and with a higher percentage of institutional investors than firms with no plans.

[46] The case of rolling maturity (i.e. fourth column) is reported in Figure 19 in Appendix A.

Table 2
Calibration Results

Panel A reports the input parameters as calibrated on the firm with plans (firm p) and firm with no plans (firm np) in the dataset. Panel B reports the calibration results (with t-values in the parenthesis) for the single-firm model and two-firm models respectively. The column "T constant" refers to a fixed regulation event $T = 2020$. The column "T rolling" refers to a regulation event T_t that rolls in each reporting year t so that the maturity $T_t - t$ remains fixed. Details on the measurement of the parameters are reported in the Appendix D.

Panel A: Input Parameters

$\hat{\beta}^p$	0.93
$\hat{\beta}^{np}$	0.92
$\hat{\eta}^p$	0.0037
$\hat{\eta}^{np}$	0.0029
$\hat{\omega}^p$	0.036
$\hat{\omega}^{np}$	0.023
$\hat{\phi}^p$	8
$\hat{\phi}^{np}$	8

Panel B: Estimates	**Single-Firm**		**Two-Firms (Stackelberg)**	
	T constant	T rolling	T constant	T rolling
$\bar{\lambda}^p$	79.2***	75.3***	120.0***	121.8***
	(10.3)	(10.6)	(2.96)	(2.99)
$\bar{\lambda}^{np}$			126.0***	128.2***
			(2.84)	(2.77)
γ^p			1,356***	1,315***
			(7.81)	(7.30)
γ^{np}			985.2***	971.5***
			(7.38)	(7.06)

Panel C: Moments	**Single-Firm**		**Two-Firms (Stackelberg)**	
	T constant	T rolling	T constant	T rolling
Average Action (firm p)	-0.006	-0.006	-0.006	-0.006
Model Implied	0.009	-0.005	0.015	-0.005
Variance Action (firm p)	0.010	0.010	0.010	0.010
Model Implied	0.000	0.000	0.003	0.002
Average Plan (firm p)	0.151	0.151	0.151	0.151
Model Implied	0.059	0.063	0.108	0.141
Variance Plan (firm p)	0.001	0.001	0.001	0.001
Model Implied	0.004	0.007	0.018	0.026
Average Action (firm np)			-0.026	-0.026
Model Implied			0.011	-0.004
Variance Action (firm np)			0.006	0.006
Model Implied			0.002	0.001

Figure 7
Model-Implied and Observed Moments - Leader-Follower Model

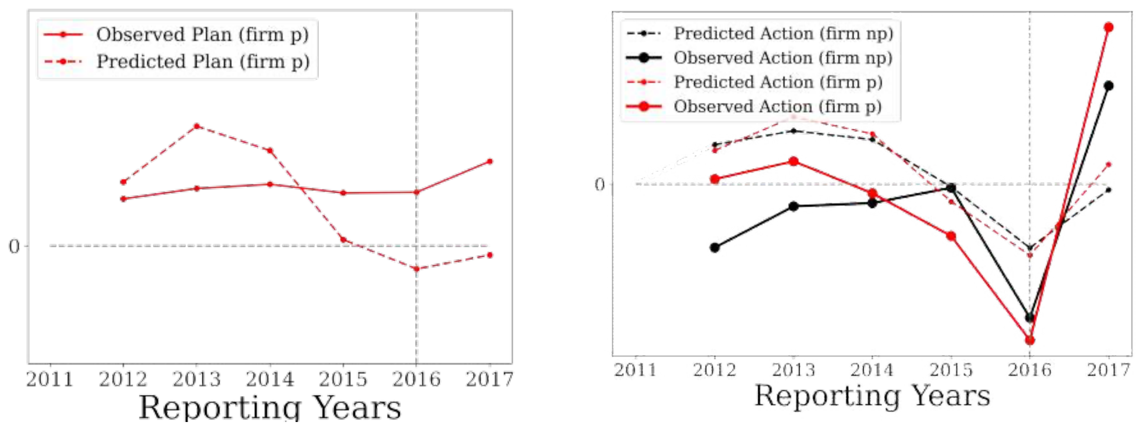

The left plot compares the model-implied and observed abatement actions for the leader firm against reporting years in CDP. The right plot compares the model-implied and observed actions for the follower firm against reporting years in CDP. Thick (dashed) lines refer to observed (model-implied) moments. Model parameters in input are reported in Table 2, Panel A. Estimated parameters in Table 2, Panel B, third column.

and leader-follower dynamics does result in a better ability to capture the observed dynamics of abatement in the data. In Figures 13, 14, 15, and 16 in Appendix A, we also show that alternative calibration exercises with different specifications of the strategic setting do a worse job in fitting reported plans and action in the dataset. In particular, we find that a leader-follower interaction with heterogeneous beliefs, and a time-varying reputation externality indexed to the ESG news is essential to capture the relationship between firms' beliefs and plans around and after the Paris agreement announcement.

In Figure 8, we plot the counterfactual optimal emissions paths of the firm with plans, $\{\xi^{*,p}_{t,\tau=t,\ldots T}\}$, generated in the single-firm model (red lines), and in the two-firm model (blue lines) respectively, assuming $t = 2022$ and $T = 2030$. From time $t = 2022$ onwards, we impose a true levy of $\lambda = 125\$/mtCO_2e$ (thick lines) or $\lambda = 25\$/mtCO_2e$ (dashed lines), setting it equal for both the leader's and follower's firms. Finally, we assume that the reputation externality decreases quickly after $t = 2022$ when the regulatory event occurs.

As the figure shows, a levy of $25\$/mtCO_2e$ in 2022 generates roughly 20% cumulative decrease in emissions by 2030 in the single-firm model, and almost 25% cumulative decrease in emissions by 2030 in the two-firm model. Both these reductions are insufficient to meet the 45% target established at the Paris agreement. In contrast, raising the levy to 125\$/mtCO_2e generates over 75% cumulative decrease in emissions by 2030 in the single-firm model, and over 90% cumulative decrease in emissions by 2030 in the two-firm model. Consistent with Proposition 2, abatement in the strategic setting is strongest at the earliest maturities: incorporating the decreasing time-path of the reputation externality, the leader firm anticipates that future profits from reputation will quickly vanish, and as a result abates the most at the earliest maturities.[47]

Figure 8
Optimal Emissions Path - Firm with Plans

The plots show the optimal carbon emissions path of the firm with plans $\{\xi^{*,p}_{t,\tau}\}_{\tau=t,\ldots T}$ as generated by the single-firm model (red lines) and the two-firm model (blue lines) respectively. Dashed and thick refer to the emissions generated by the models when the levy parameter $\bar{\lambda}^p = \bar{\lambda}^{np} = 25\$/\text{mtCO}_2\text{e}$ and $\bar{\lambda}^p = \bar{\lambda}^{np} = 125\$/\text{mtCO}_2\text{e}$ respectively and $\sigma^p = \sigma^{np} = 0$. The strategic parameters are reported in Table 2, Panel B, third column. The remainder of model parameters in input are reported in Table 2, Panel A.

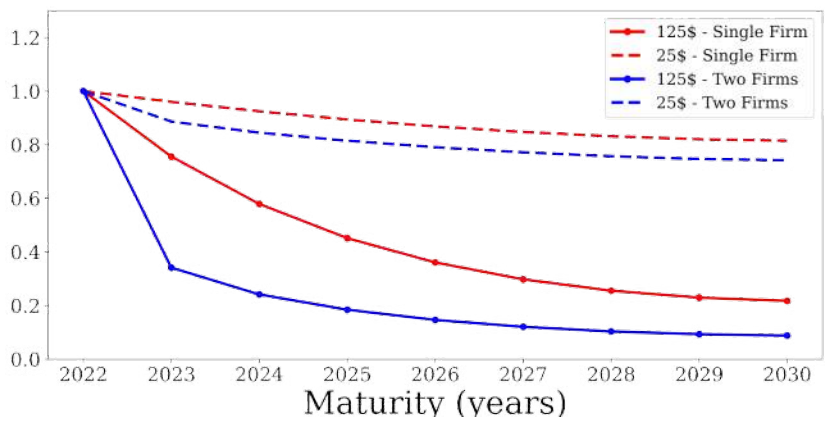

4.3 Discussion of alternative mechanisms

Although the leader-follower model of reputational externalities is able to match the main patterns identified in the data, we cannot rule out the possibility that an alternative mechanism might also match these patterns as well. Two such plausible alternative mechanisms include: i) managerial incentive schemes set by shareholders (e.g. so-called extrinsic motivations as in Ariely et al. (2009)

[47]In Figure 17 in the Appendix, we attempt to identify the effect of uncertainty by repeating the exercise in Figure 8, but keeping uncertainty at the levels estimated from CDP disclosures in $t = 2016$. In both single-firm and two-firm models, the effect of uncertainty seems to be smaller relative to that of the mean and the reputation externality, i.e., introducing uncertainty does not change the ordering of abatement reported in Figure 8. A caveat here is that the firm does not internalize the possibility of learning, for example, from future announcements.

such as monetary rewards/penalizations in case of overcompliance/undercompliance with climate regulations; and ii) clean technology adoption induced by the policy announcement (e.g., redirected technical change as in Acemoglu et al. (2016)), which would eventually make firms more efficient thereby reducing the sensitivity of output to carbon emissions.

We note here that for the mechanism in i) to generate a similar amplification effect as the one predicted in our model, it must be that the firm's shareholders expect undercompliance with the regulatory policy in absence of the monetary incentive, and that such beliefs are consistent with the cross-firm patterns observed in the data. Put differently, such a mechanism requires an additional free parameter which captures time-variation in shareholders' beliefs.

On the other hand, a more complex technological structure of the model as in ii), but also involving a time-varying emission intensity, could serve as a substitute for the reputation externality in matching some of the time-series dynamics in the model. The Paris Agreement shock might have induced acceleration in firms' optimal timing of adoption of clean technologies, as Figure 7 in the Appendix A shows (around and after the Paris agreement announcement, an increasing percentage of firms across the two groups issued green or sustainable debt, which is often taken to finance such technology adoption). That said, the market for such technologies experienced a broad slowdown in the years preceding the Paris agreement announcement,[48] furthermore, such preemption effects would still incur well-documented time lags[49] between the decision to adopt the technology and the effective reduction in emission intensities.[50]

Nevertheless, it is useful to understand how introducing such a technology option would help fitting the patterns in the data. Suppose that a shock in the firm's belief about the levy (such as the one that firms with plans experience in 2015), triggers the decision to adopt clean technology. Then firms with plans would know to be better hedged to the future regulatory event, which should reflect in a *lower* sensitivity of their current emissions to revisions in beliefs about the levy. Hence, while the clean technology adoption could substitute for the reputation externality in explaining why firms with plans do not react extensively (in terms of current emissions) to their belief revisions

[48]See, for example, the study in León et al. (2018) and the green growth statistics from the OECD database. A summary of the slowdown in US green innovation can be found in the OECD report *Green Growth Indicators, 2017*.

[49]See, for example, the review in Hoppe (2002)

[50]It is also worth mentioning that, during the observation period in CDP, average emissions intensities remain constant for firms with and without plans.

in 2015, it would make it even harder to explain (without reputation externalities) the greater reaction following the Paris agreement announcement. More specifically, in the single-firm setting, firms with plans should mantain a reduced sensitivity to changes in their own beliefs over the levy. Therefore, the reputation externality would still play a critical role in generating the reactions which follow the Paris agreement announcement.

As a final exercise, to develop greater confidence in our model, we consider how well it performs out-of-sample.

5 Out of Sample Predictions

To assess the quality of the model's out-of-sample predictions, we estimate the model using data up to 2017, and then use it to predict U.S. public firms' responses for the years 2018 and 2019.[51] This period is particularly interesting, as it allows us to evaluate the impact on firms' responses of a regulatory shock that goes in the opposite direction to those used to fit the model, namely, U.S. President Donald Trump's announcement to pull back from the Paris agreement, which occurred in June 2017.

The left-hand panel of Figure 9 shows beliefs computed from the extended CDP dataset that we use as input variables to our out-of-sample evaluation exercise. As can be seen in the figure, in the year following Trump's pull-back announcement, all firms significantly downgrade their expectations of the impact of climate policy regulation. In contrast with the patterns previously observed, however, firms reporting plans for future emissions abatement now appear to revise their beliefs about the intensity of future climate regulation more extensively than those not reporting plans in response to the announcement.

The right-hand panel of Figure 9 shows another interesting observation from the new disclosure data. Following Trump's pull-back announcement, in addition to their changing beliefs about the intensity of regulation, firms increase their expected time horizon over which climate regulation is expected to come into effect, by a median value of 2 years.

[51]Over these two years, CDP implemented a set of changes to make the questionnaires more aligned with the recommendations of the Task Force on Climate-Related Financial Disclosures (TCFD), which was established in 2016. In Appendix C we report the major changes to the responses and format arising from these changes and how they affected the measures that we compute. The Appendix also describes a few adjustments to the data that were needed to conduct the out-of-sample exercise.

Figure 9
Extended Beliefs

The left plot shows the belief metric as in (2) against reporting years in the extended CDP questionnaires. The red line (black line) refers to firms that disclose (do not disclose) plans in the previous reporting year. The right plot shows the distribution of time horizons for planned emissions reduction as reported by firms with plans in the dataset. The red bars refer to time horizons as reported in 2016, the year following the Paris agreement announcement, while grey bars refer to time horizons as reported in 2017, the year following President Trump's pullback announcement.

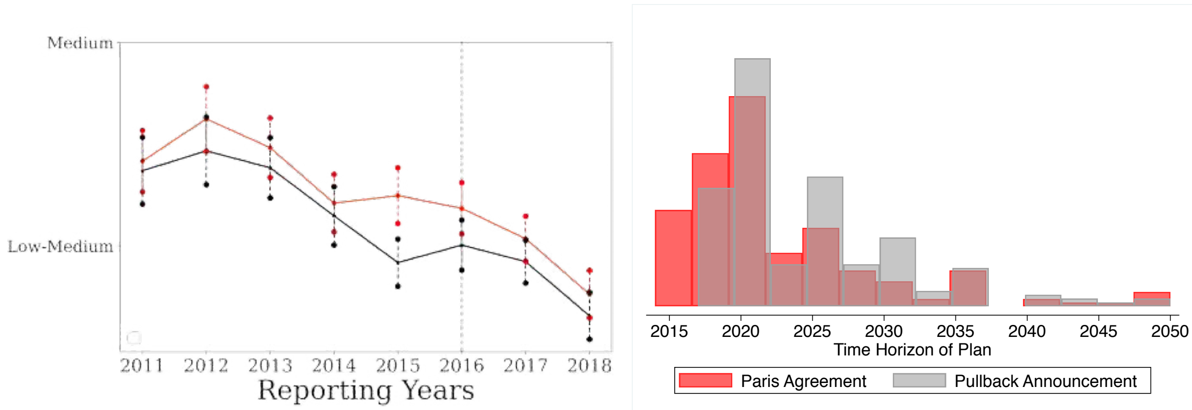

To conduct the out-of-sample exercise, we use the beliefs reported in the right panel of Figure 9 to generate emissions abatement plans and actions from the two-firm model. We fix the strategic parameters at the levels reported in the fourth column of Table 2, Panel A, estimated over the period from 2011 to 2017, and we use all other parameter values reported in Table 2, Panel A, except for the time horizon T, which we extend by 2 years to account for the evidence seen in Panel B of Figure 9.

The left- and right-hand plots in Figure 10 show, respectively, predicted and realized actions for the leader (plan-reporting) and follower (non-plan reporting) firms for each reporting year in the dataset. The vertical dashed line in the figures indicates the beginning of the out-of-sample forecasting period. The model captures the realized drop in emissions reduction predicted by the downward revision in beliefs following the pull-back announcement for both leader and follower firms, and correctly predicts a larger response for the leader firm. The model also captures the increase in emissions reduction observed in the final year of reported data, which once again is more pronounced for leader than follower firms in the data. Overall, the out-of-sample exercise helps to increase confidence in the augmented two-firm model's ability to predict the dynamics of reported

Figure 10
Out-of-sample Prediction

The left plot compares the model-implied and observed abatement actions for the leader firm against reporting years in CDP. The right plot compares the model-implied and observed actions for the follower firm against reporting years in CDP. Thick (dashed) lines refer to observed (model-implied) moments. Estimated parameters are reported in the third column of Table 2, Panel B. Other parameters in input are reported in Table 2, Panel A.

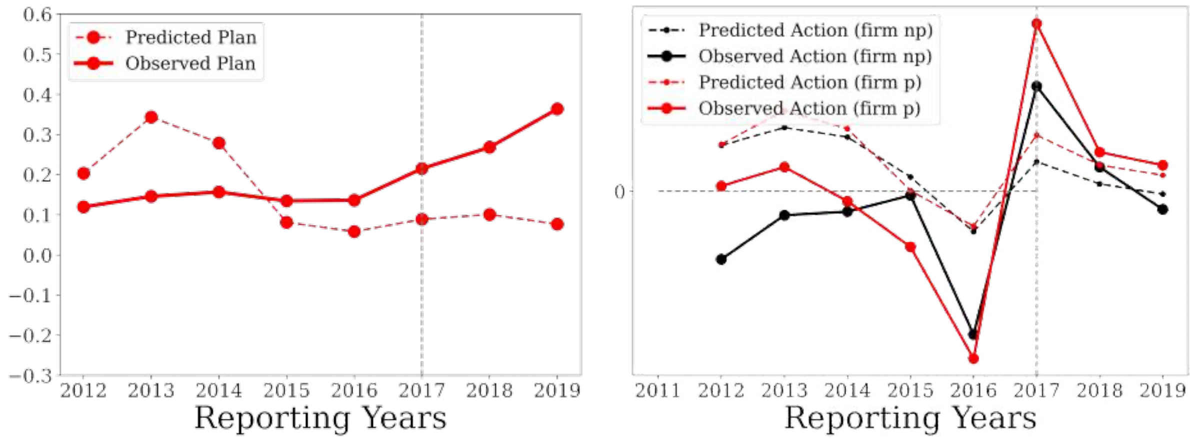

emissions abatement.

5.1 Conclusions

In this paper, we pursue a structural approach to identify the determinants of firms' decision-making when they face climate regulation risk. We begin by bringing new empirical observations to the table, using firms' disclosures to the Carbon Disclosure Project (CDP), which we verify using third-party sources (Bloomberg, Thomson Reuters, and MSCI) who produce ESG ratings of firms. We document patterns in firms' beliefs about the climate regulation risks that they face, their plans for future abatement, and their actions to date on mitigating carbon emissions. We find that in the five years prior to the Paris announcement, firms' actions on carbon abatement as well as their beliefs about climate regulation gradually reduce. However, firms' actions and beliefs both adjust sharply around the announcement of the Paris climate change agreement in 2016, with the size of these responses depending on whether or not a given firm pre-announces its plans for carbon emissions reduction.

To learn more about the underlying structure that can jointly rationalize these findings, we build two dynamic models of emissions abatement. The first model features an atomistic firm operating with polluting capital, which is exposed to a future uncertain carbon levy of known maturity, updates its belief over the levy in each time period prior to the regulation, and incurs convex capital adjustment costs when abating emissions. We calibrate the model to the data, feeding it with the dynamics of reported beliefs, and comparing the predicted plans and actions from the model with those in the data. While the model can fit the dynamics of abatement prior to the Paris agreement, the reactions to the Paris agreement predicted by this atomistic firm model cannot match the sharp variations observed in the data.

Our more complex model introduces a second firm into the economy, with the goal of understanding whether the amplification we observe in the data can be rationalized by firms strategic responses to one another. Specifically, we introduce a reputation externality in the firms' payoffs, which reduces the profits of a given firm when the other firm abates, and vice-versa. Moreover, we allow the firms to have different beliefs, with the "leader" firm given knowledge of the "follower" firm's beliefs to rationalize the observed differences between plan- and non-plan-reporting firms. We verify that the Stackelberg equilibrium of the model predicts abatement dynamics that closely match the patterns that we observe in the data, and that the ingredients that we add are needed to deliver the fit between the model predictions and the data. We validate the model, showing that it is well-able to predict firms' abatement actions out-of-sample in the period following the announcement of the U.S. pullback from the Paris agreement.

There is much work to be done on the economics of climate change and carbon emissions. Our paper contributes to this important agenda by demonstrating that i) climate regulation matters greatly to firms, and ii) to better understand firms' responses to regulation events, it is important to take strategic interactions, differences in beliefs, and information asymmetries between firms into account. We believe that further work to learn more about the specific microeconomic mechanisms at work along these lines will pay rich dividends.

References

Daron Acemoglu, Ufuk Akcigit, Douglas Hanley, and William Kerr. Transition to clean technology. *Journal of Political Economy*, 124(1):52–104, 2016.

Dan Ariely, Anat Bracha, and Stephan Meier. Doing good or doing well? image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1):544–55, 2009.

Adelina Barbalau and Federica Zeni. Designing green securities: Action vs. outcome- based incentives. *Available at SSRN 4155721*, 2023.

Michael Barnett, William Brock, and Lars Peter Hansen. Pricing uncertainty induced by climate change. *The Review of Financial Studies*, 33(3):1024–1066, 2020.

Bruno Biais and Augustin Landier. Emission caps and investment in green technologies. *Available at SSRN 4100087*, 2022.

Patrick Bolton and Marcin Kacperczyk. Do investors care about carbon risk? *Journal of financial economics*, 142(2):517–549, 2021a.

Patrick Bolton and Marcin T Kacperczyk. Firm commitments. *Columbia Business School Research Paper*, 2021b.

Tamma Carleton and Michael Greenstone. Updating the united states government's social cost of carbon. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2021-04), 2021.

M. Carney. Breaking the tragedy of the horizon–climate change and financial stability. *Speech given at Lloyd's of London, September*, 2015.

Dane M Christensen, George Serafeim, and Anywhere Sikochi. Why is corporate virtue in the eye of the beholder? the case of esg ratings. *The Accounting Review*, 97(1):147–175, 2022.

Alexander Dyck, Karl V Lins, Lukas Roth, and Hannes F Wagner. Do institutional investors drive corporate social responsibility? international evidence. *Journal of Financial Economics*, 2018.

Caroline Flammer. Corporate green bonds. *Journal of Financial Economics*, 142(2):499–516, 2021.

Raimo P Hämäläinen and Ilkka Leppänen. Cheap talk and cooperation in stackelberg games. *Central European Journal of Operations Research*, 25(2):261–285, 2017.

Samuel M Hartzmark and Abigail B Sussman. Do investors value sustainability? a natural experiment examining ranking and fund flows. *The Journal of Finance*, 74(6):2789–2837, 2019.

Amanda Heitz, Youan Wang, and Zigan Wang. Corporate political connections and favorable environmental regulatory enforcement. *Management Science*, 2021.

Andreas GF Hoepner, Ioannis Oikonomou, Zacharias Sautner, Laura T Starks, and Xiaoyan Zhou. Esg shareholder engagement and downside risk. *Available at SSRN 2874252*, 2018.

Harrison Hong, Neng Wang, and Jinqiang Yang. Welfare consequences of sustainable finance. Technical report, National Bureau of Economic Research, 2021.

Heidrun C Hoppe. The timing of new technology adoption: theoretical models and empirical evidence. *The Manchester School*, 70(1):56–76, 2002.

Lorena Rivera León, Kyle Bergquist, Sacha Wunsch-Vincent, Ning Xu, Kunihiko Fushimi, et al. *Measuring innovation in energy technologies: green patents as captured by WIPO's IPC green inventory*, volume 44. WIPO, 2018.

Laura Xiaolei Liu, Toni M Whited, and Lu Zhang. Investment-based expected stock returns. *Journal of Political Economy*, 117(6):1105–1139, 2009.

L. Luo, Y.-C. Lan, and Q. Tang. Corporate incentives to disclose carbon information: Evidence from the cdp global 500 report. *Journal of International Financial Management & Accounting*, 23(2):93–120, 2012.

John J McConnell and Henri Servaes. Additional evidence on equity ownership and corporate value. *Journal of Financial economics*, 27(2):595–612, 1990.

Kyle C Meng. Using a free permit rule to forecast the marginal abatement cost of proposed climate policy. *American Economic Review*, 107(3):748–84, 2017.

William Nordhaus. Estimates of the social cost of carbon: concepts and results from the dice-2013r model and alternative approaches. *Journal of the Association of Environmental and Resource Economists*, 1(1/2):273–312, 2014.

Alexei V Ovtchinnikov, Syed Walid Reza, and Yanhui Wu. Political activism and firm innovation. *Journal of Financial and Quantitative Analysis*, pages 1–36, 2019.

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

Adriano A Rampini, Amir Sufi, and S Viswanathan. Dynamic risk management. *Journal of Financial Economics*, 111(2):271–296, 2014.

Silvana Tenreyro and Tiloka De Silva. Climate-change pledges, actions and outcomes. *Journal of the European Economic Association*, 2021.

Richard SJ Tol. The social cost of carbon. *Annu. Rev. Resour. Econ.*, 3(1):419–443, 2011.

Olivier David Zerbib. The effect of pro-environmental preferences on bond prices: Evidence from green bonds. *Journal of Banking & Finance*, 98:39–60, 2019.

Liming Zhang, Fei Ye, Li Yang, and Guichuan Zhou. Impact of political connections on corporate environmental performance: From a green development perspective. *Sustainability*, 11(5):1317, 2019.

Climate Regulation and Emissions Abatement: Theory and Evidence from Firms' Disclosures

Online Appendix

Tarun Ramadorai and Federica Zeni[†]

May 19, 2023

*We thank Pat Akey, Ian Appel, Harjoat Bhamra, Patrick Bolton, Charlie Donovan, Chris Hansman, Harrison Hong, Zacharias Sautner, and seminar participants at Imperial College Business School, Cornell University, the UZH workshop on Climate Finance, the HEC Paris Spring Finance conference, and the Western Finance Association conference for comments.

[†]Ramadorai (corresponding author): Imperial College Business School and CEPR. Email: t.ramadorai@imperial.ac.uk. Zeni: World Bank Research. Email: fz@worldbank.org.

A Tables and Figures

Table 1
Selected Disclosures

Reporting Year	Risk (1)	Footprint (2)	Risk & Footprint (1)+(2)	Plans (3)	To Plans (4)	To No Plans (5)
2011	249	381	**229***			
2012	302	412	**229**	**106**		
2013	342	454	**274**	**123**	16	12
2014	350	460	**293**	**142**	22	6
2015	383	471	**331**	**152**	18	20
2016	410	498	**342**	**163**	25	12
2017	439	503	**375**	**186**	25	14
Total Firms	526	611	446	262	100	61

Number of firms in the CRSP/COMPUSTAT North America universe reporting selected disclosures in the CDP questionnaires between 2011 and 2017. Column (1) is the subset of firms that disclose climate regulation risk; column (2) is the number of firms that disclose total carbon footprint; column (1)+(2) is the selected dataset: firms that disclose regulation risk, carbon footprint, and report in the dataset for at least two consecutive years. Column (3) is the subset of firms in the selected sample that also disclose emissions reduction plans in the previous reporting year (*for 2011 only, the column (3) is the group of firms that report plans in the same reporting year). Column (4) is the subset of firms in the plan-reporting group (i.e. column (3)) that were in the no plan-reporting group in the previous year. Column (5) is the subset of firms in the no plan-reporting group (i.e. column (1)+(2) - (3)) that were in the plan-reporting group in the previous year.

Table 2

Beliefs - Linear Regressions

Regressor	**Beliefs**			
Emissions	0.16***	0.19***	0.20***	0.20***
	[0.04]	[0.04]	[0.05]	[0.05]
Market Value		-0.16***	-0.15***	-0.15***
		[0.04]	[0.04]	[0.04]
Intercept	0.09	1.07*	0.61	1.67**
	[0.51]	[0.55]	[0.71]	[0.78]
Sector dummy?	No	No	Yes	Yes
Year dummy?	No	No	Yes	Yes
HQ Country dummy?	No	No	No	Yes
$\mathcal{R}^2$	0.04	0.07	0.09	0.12
Firms	446	408	408	408

Beliefs and emissions are collected from CDP as detailed in appendix C. Market value is provided by CRSP/COMPUSTAT. Sector dummies are the Global Industry Classification (GIC) sectors, country dummies refer to the country of the company Head Quarter (HQ), both variables are provided by CRSP/COMPUSTAT. Standard errors in square brackets are clustered at the firm-level. *,**,*** indicates statistical significance at the 10%, 5% and 1% level respectively.

Table 3

Belifs and Actions – Summary Statistics

Panel A: Beliefs	**Plan**		**No Plan**		**Plan - No Plan**	
Reporting Year	Mean	(Std Err.)	Mean	(Std. Err)	Diff	(Std. Err)
2011^{+}	2.41***	(0.15)	2.36***	(0.16)	0.04	(0.23)
2012	2.62***	(0.16)	2.46***	(0.17)	0.16	(0.23)
2013	2.48***	(0.15)	2.38***	(0.15)	0.10	(0.21)
2014	2.20***	(0.14)	2.14***	(0.14)	0.06	(0.20)
2015	2.24***	(0.14)	1.92***	(0.12)	0.32*	(0.17)
2016	2.18***	(0.13)	2.00***	(0.12)	0.18	(0.17)

Panel B: Actions	**Plan**		**No Plan**		**Plan - No Plan**	
Reporting Year	Mean	(Std Err.)	Mean	(Std. Err)	Diff	(Std. Err)
2012	0.00	(0.02)	-0.07**	(0.03)	0.07**	(0.03)
2013	0.02	(0.02)	-0.02	(0.02)	0.05*	(0.03)
2014	0.00	(0.02)	-0.02	(0.02)	-0.01	(0.02)
2015	-0.05*	(0.02)	0.00	(0.02)	-0.05*	(0.03)
2016	-0.17***	(0.02)	-0.15***	(0.02)	-0.02	(0.04)
2017	0.17***	(0.02)	0.11***	(0.02)	0.06**	(0.03)

Summary statistics of beliefs in Panel A and actions in Panel B as reported in Figure 2 in the main text. The column Plan (No Plan) is the group of firms in column (3) (columns (1)+(2)-(3) respectively) summarized in Table 1. The column Plan - No Plan is the average difference in reported beliefs and actions between the group of firms with and without plans in the dataset. Standard errors are reported in parenthesis. *,**,*** indicates statistical significance at 10%, 5% and 1% respectively.

Table 4

Plans vs No Plan Reporting Firms - Selected Statistics

	Plan	**No Plan**	**Plan - No Plan**
Market Share	0.07	0.05	0.02***
Engagement with Climate Policymakers	0.94	0.78	0.15***
Direct Funding to Regulatory Activities	0.77	0.64	0.14**
Supplier Propensity to Report Plan	0.40	0.36	0.04***

The first row is the average fraction of the firm's annual sales (Sales Turnover Net from CRSP/Compustat) in the respective GIC Industry. The second and third rows are the fraction of firms that report to engage positively with climate policymakers and provide fundings to regulatory activities respectively, collected from the latest CDP questionnaire (2017). The fourth row is the propensity of the firm's representative supplier to report plans, measured as the sum of the propensity to report plans in each industry weighted by the share of the firm's total inputs from that industry. Input tables are provided by the Bureau of Economic Analysis and are available at the NAICS industry level. We only report firm's suppliers for which a straightforward conversion from NAICS codes to GIC codes is available. $*, **, ***$ indicate statistical significance at the 1%, 5%, and 10% respectively.

Table 5

ESG Funds Ownership - Linear Regressions

Variable	I	II	III
Has Plan	108.6**	94.6**	50.2
	(48.3)	(44.9)	(24.2)
Size		86.6***	123.3***
		(10.7)	(11.4)
Industry Dummy	No	No	Yes
$\mathcal{R}^2$	0.02	0.14	0.29
N	422	422	422

Linear regression of ESG funds ownership on firm size. The number of ESG funds is taken from Bloomberg ESG equity. The term ESG encompasses key words Clean Energy, Climate Change, ESG, Environmentally Friendly, and Socially Responsible. The word fund encompasses Fund of Fund, Hedge Fund, Private Equity Fund, Private Equity Fund of Fund, Separately Managed Acct (SMA). The variable size is the logarithm of the firm total assets as averaged across the reporting period in CDP.

Table 6
Calibration Results - Alternative Models

Externalities ($/mtCO$_2$e)	**Model I**	**Model II**	**Model III**	**Model IV**
$\bar{\lambda}^p$	119.1***	119.1***	130.3***	94.0***
	(2.89)	(2.88)	(3.01)	(21.37)
$\bar{\lambda}^{np}$	125.6***	124.2***	128.3***	50.0***
	(2.81)	(2.81)	(3.08)	(3.08)
γ^p	1,383***	1,394***	1,361***	100.0
	(9.61)	(10.27)	(12.3)	(0.51)
γ^{np}	998.9***	1014***	978.4***	100.0
	(8.05)	(10.3)	(8.46)	(0.14)
g			0.47*	
			(1.60)	

The table reports the calibration results for alternative specifications of the two-firm model. Model I refers to a leader-follower model where the firm without plans (firm np) learns about the levy from firm p's action in each reporting year t (Figure 13). Model II refers to a leader-follower model where the reputation externalities γ^p and γ^{np} are the same in each reporting year t instead of tracking the dynamics of ESG news (Figure 14). Model III refers to a leader-follower model where the reputation externalities have a monotonic term-structure $\gamma^p_{t,\tau} = e^{-g(\tau-t)}(\gamma^p ESG_t)$ and $\gamma^{np}_{t,\tau} = e^{-g(\tau-t)}(\gamma^{np} ESG_t)$ (Figure 19). Model IV refers to an otherwise equivalent two-firm model as the one described in the main text where firm p and firm np act simultaneously (Figure 16). Parameters in input are reported in Table 2, Panel A in the main text. We report only estimates for the rolling maturity of the regulation event.

Figure 1
Emissions - External Disclosures

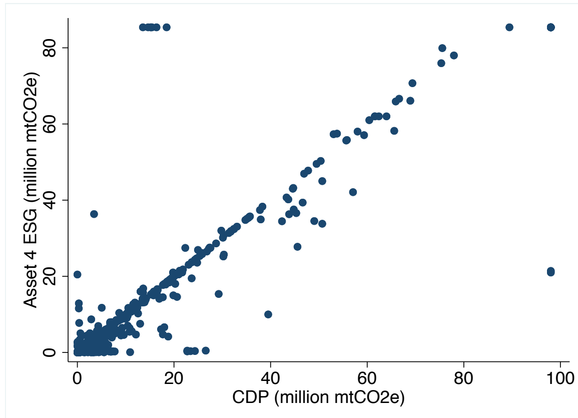

The scatter plot shows firm-level values of CO_2e emissions from Asset 4ESG (y-axis) against self-reported CO_2e emissions from CDP (x-axis), winsorized between the 1^{st} and the 99^{th} percentile of the pooled distribution. Asset 4 ESG emissions refer to the variable ENERDP023 (see the Asset 4 ESG Dada Glossary for details). The matched sample refers to the entire reporting period in the dataset.

Figure 2
Regulation Risk - Description by Firms

The word cloud highlights the most frequent words that appear in the unstructured text field in the pooled CDP dataset in which firms describe the specific regulation risks.

Figure 3
Beliefs - Constituents

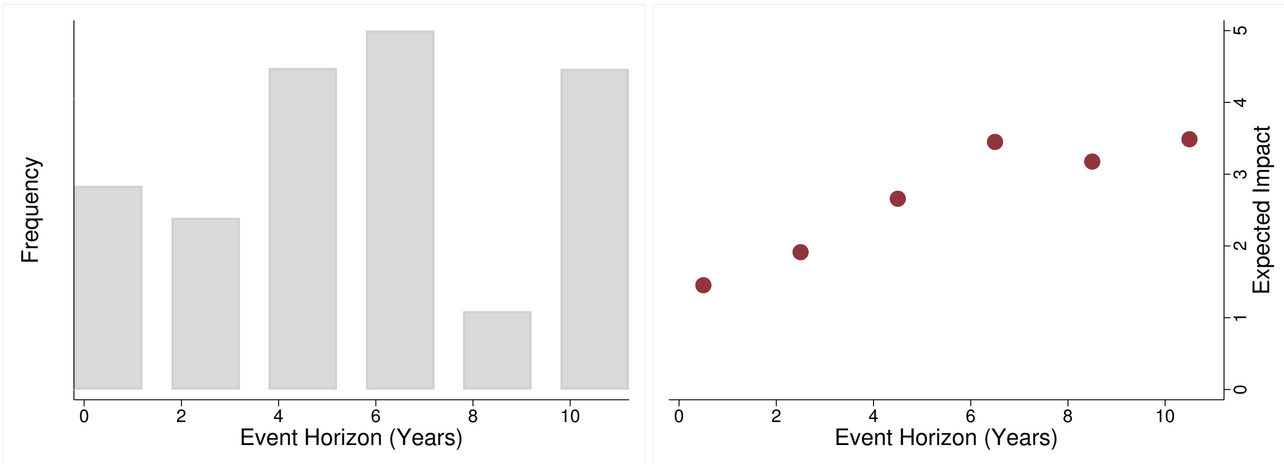

The right plot shows the average expected impact of the regulation event across different maturities of the regulation event. The left plot indicates the frequency of disclosures across each time horizon as collected from the selected CDP sample between 2011 and 2017.

Figure 4
Plans - Constituents

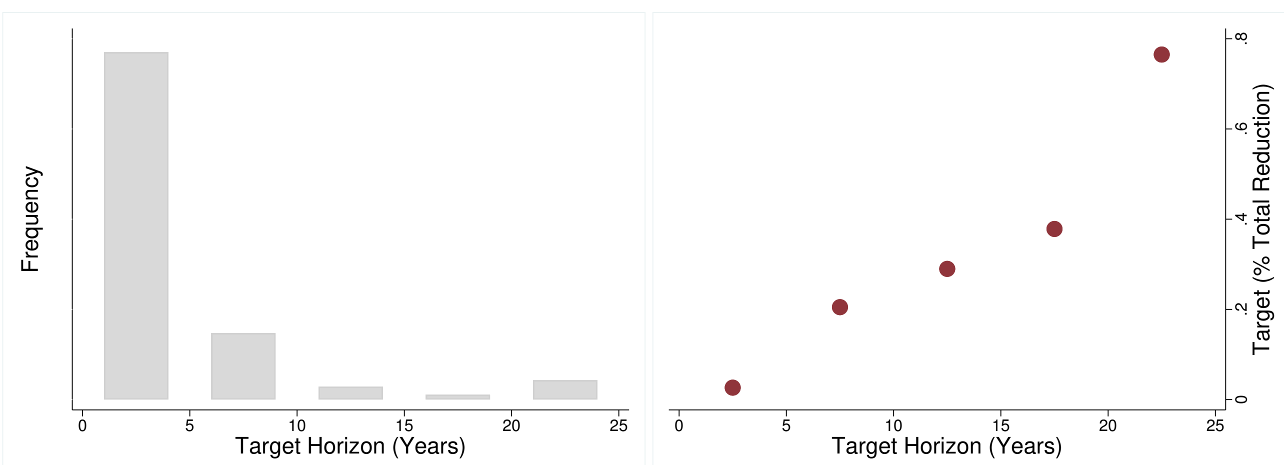

The right plot shows the average emissions reduction target across different target maturities. The left plot indicates the frequency of disclosures across each time horizon as collected from the selected CDP sample between 2011 and 2017.

Figure 5
Plans - External Environmental Ratings

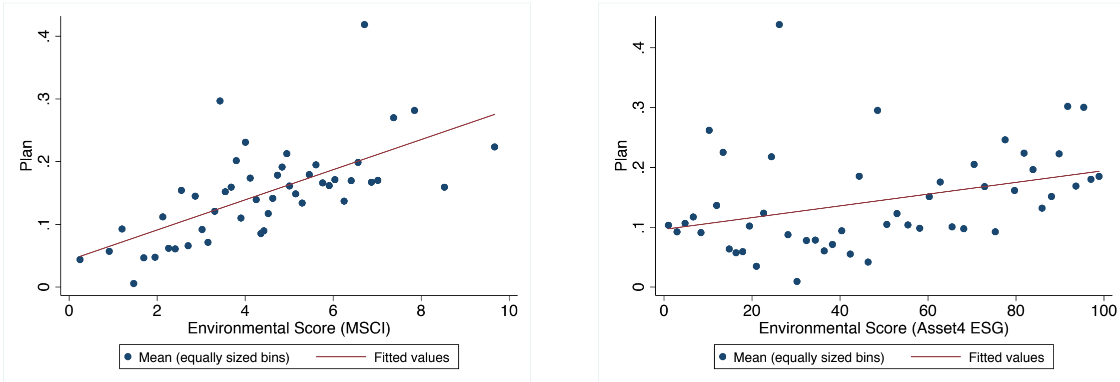

The top and bottom plots show abatement plans (y-axis) averaged across equally sized bins of the Environmental score (x-axis) provided by MSCI and Thomson Reuters Asset 4ESG respectively. Environmental scores are constituents of the ESG scores (see the Asset 4 ESG and the MSCI Dada Glossary for details).

Figure 6
Beliefs vs Crude Oil Spot Prices

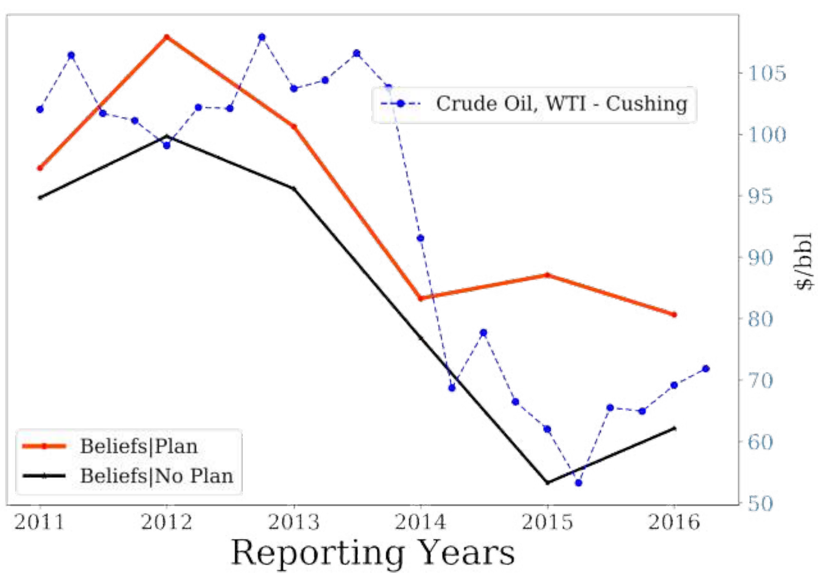

The plot compares average time-series dynamics of firms' self-reported beliefs about regulation risk (red and black thick lines) and quarterly spot prices of the West Texas Intermediate (WTI) Crude Oil - Cushing, Oklahoma as collected from the FRED database (blue dashed line). Oil prices are lagged of three quarters to match the fact that CDP questionnaires are submitted in June or July each year (with some later submissions).

Figure 7
Green Debt Issuance

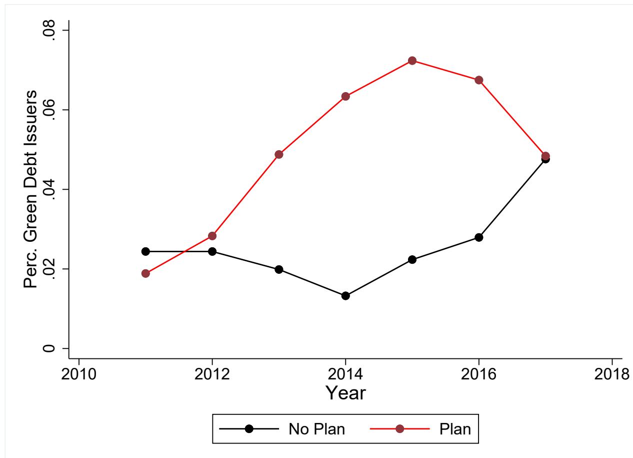

The plot shows the fraction of firms with outstanding green debt issuances across reporting years in the dataset. Green debt is identified in the Bloomberg fixed income database as green loans, green bonds, sustainable bonds, and sustainability-linked bonds and loans. The red (black) line refers to firms with and without plans in the dataset.

Figure 8
Beliefs, Plans, Actions - Balanced Panel

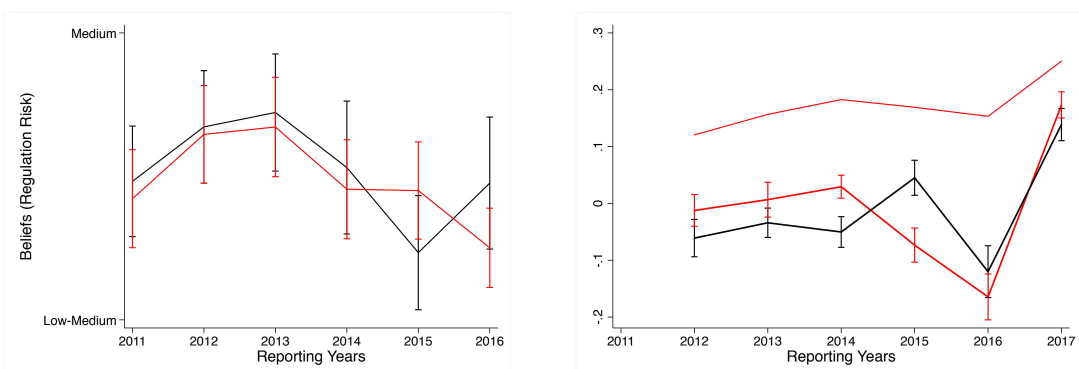

The left and right-hand plots report beliefs, actions, and plans respectively as constructed from a restricted CDP sample which includes only firms reporting to CDP since 2011. Variables are otherwise constructed following the same procedures as the ones reported in the main text.

Figure 9

Beliefs, Plans, Actions - Weighted Averages

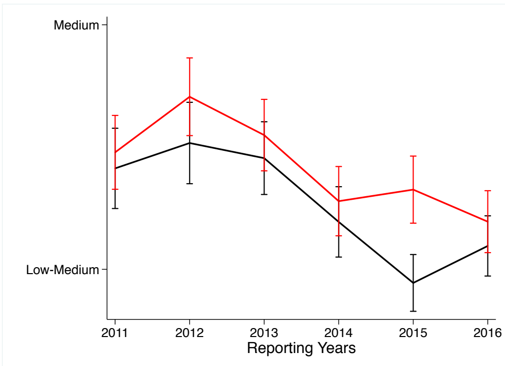

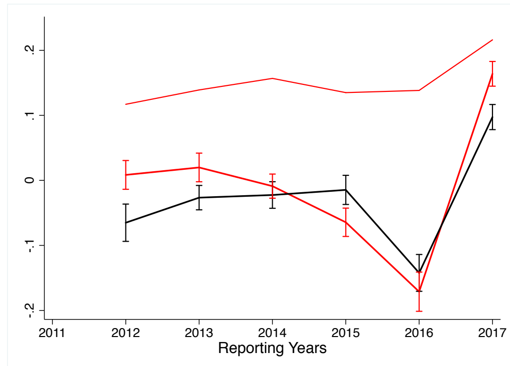

The left and right-hand plots report beliefs, actions, and plans respectively as constructed from weighted averages of firms' beliefs, plans, and actions across reporting years. Weights $w_{i,t}$ equal the emission intensity of each firm i in reporting year t, divided by the total emission intensity of firm i's group in reporting year t. Variables are otherwise constructed following the same procedures as the ones reported in the main text.

Figure 10

Beliefs - Magnitude/Likelihood Components

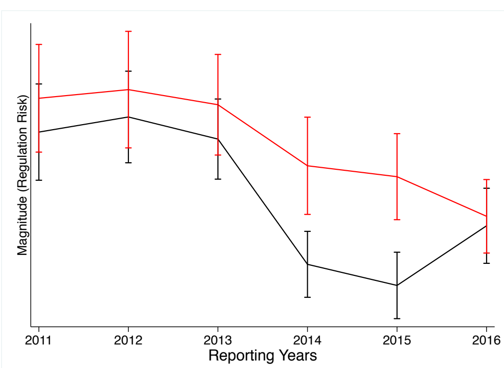

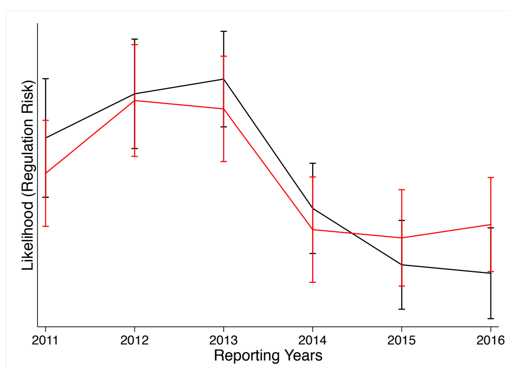

The left-hand plot reports the average magnitude of the regulation risk across reporting years in CDP. The right-hand plot reports the average likelihood of the regulation risk across reporting years in CDP. The red (black) line refers to firms that disclose (do not disclose) plans in the previous reporting year.

Figure 11
Revisions in Plans and Actions across Industries

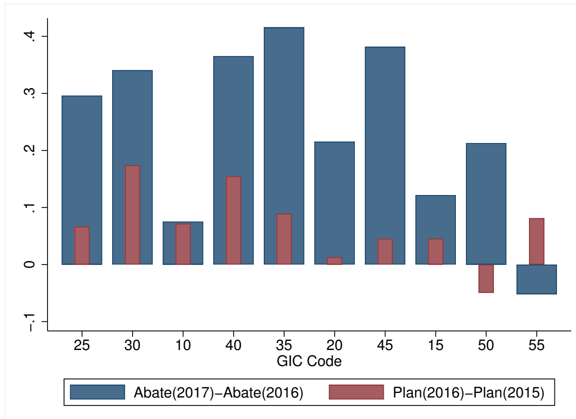

The bar plot shows average changes in reported plans and actions across GIC sectors. The red bars refer to changes in plans between the years surrounding the Paris agreement announcement. The blue bars refer to changes in actions between the years surrounding the Paris agreement.

Figure 12
Reported Beliefs - Standard Deviation

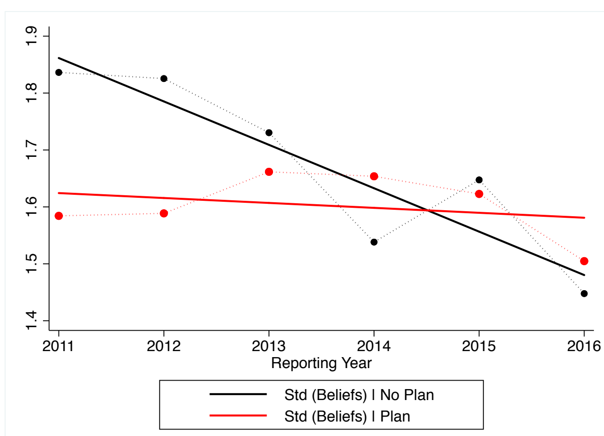

The plot shows the standard deviation of beliefs reported by the group of firms with plans (in red dots) and the standard deviation of beliefs reported the group of firms without plans (in black dots) for each reporting year $t = 2011, \ldots 2016$. Red and black lines show linear interpolations for the group of firms with and without plans respectively.

Figure 13
Alternative I - Stackelberg with learning

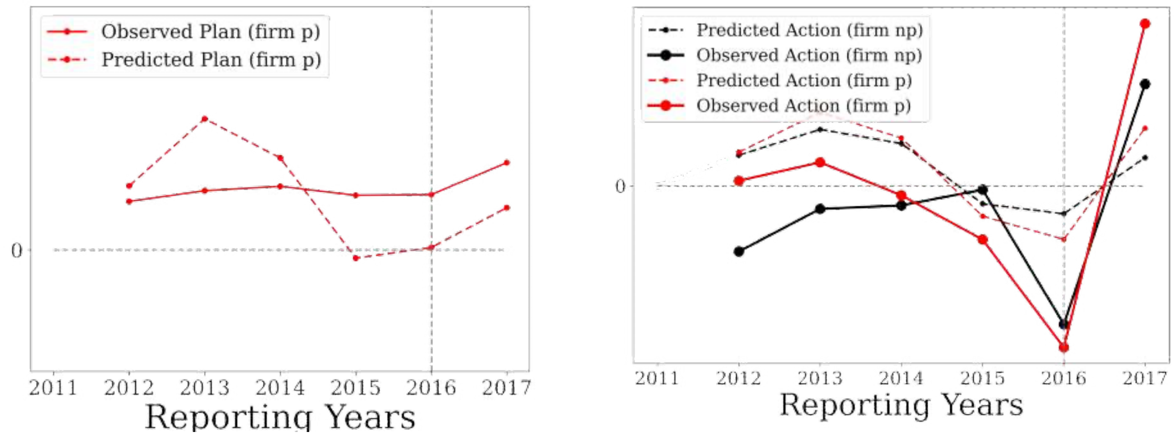

The left plot compares predicted and observed abatement plans and actions across reporting years in CDP. Predicted moments are generated from a leader-follower model where the firm without plans (firm np) learns about the levy from firm p's actions in each reporting year t. Results are reported in the first column of Table 6. Thick (dashed) lines refer to observed (model-implied) moments, red-circle (black-star) lines refer to the subset of firms with (without) abatement plans respectively.

Figure 14
Alternative II - Stackelberg with constant reputation

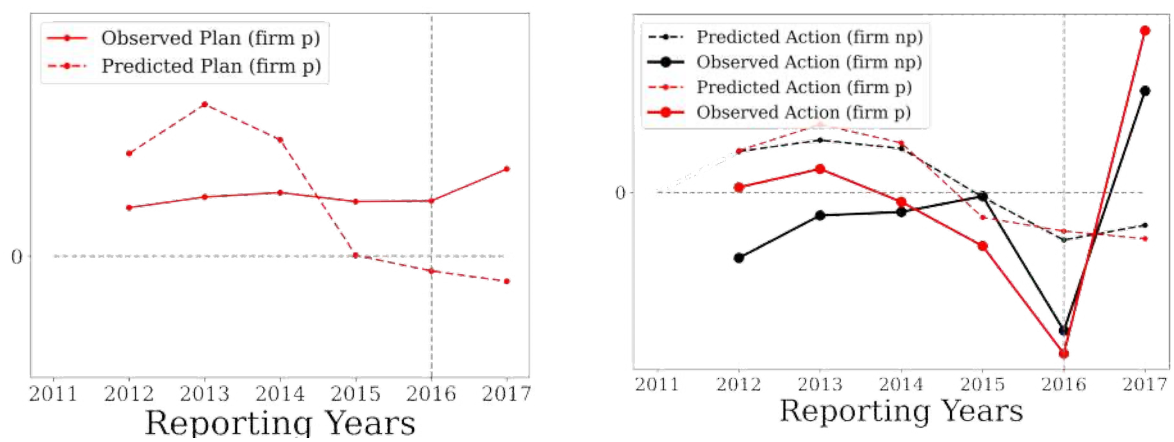

The left plot compares predicted and observed abatement plans and actions across reporting years in CDP. Predicted moments are generated from a leader-follower model where the reputation externalities γ^p and γ^{np} are kept constant across the entire reporting period instead of tracking the dynamics of ESG news. Results are reported in the second column of Table 6. Thick (dashed) lines refer to observed (model-implied) moments, red-circle (black-star) lines refer to the subset of firms with (without) abatement plans respectively.

Figure 15

Alternative III - Stackelberg with reputation term-structure

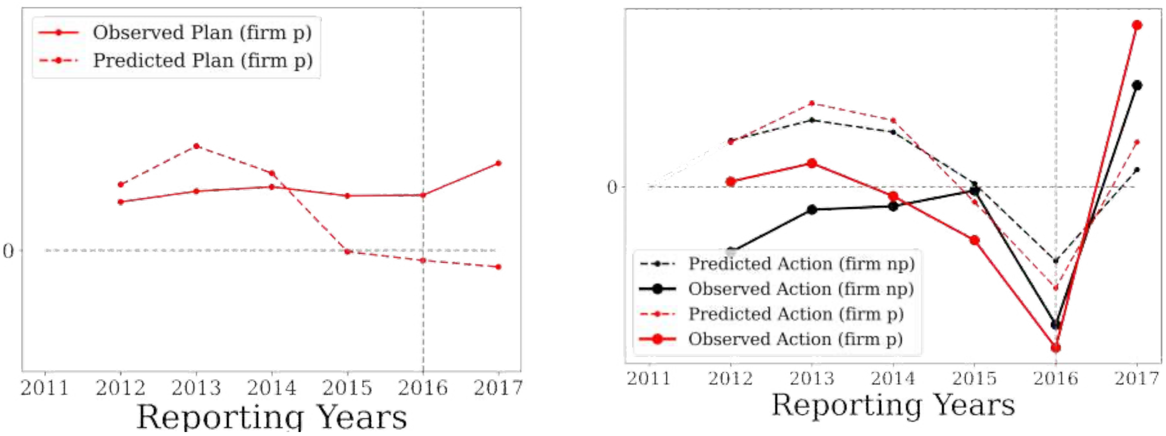

The left plot compares predicted and observed abatement plans and actions across reporting years in CDP. Predicted moments are generated from a leader-follower model where the reputation externalities have a monotonic term-structure $\gamma^{p}_{t,\tau} = e^{-g(\tau-t)}(\gamma^{p}ESG_t)$ and $\gamma^{np}_{t,\tau} = e^{-g(\tau-t)}(\gamma^{np}ESG_t)$, with g an additional parameter to be estimated. Results are reported in the third column of Table 6. Thick (dashed) lines refer to observed (model-implied) moments, red-circle (black-star) lines refer to the subset of firms with (without) abatement plans respectively.

Figure 16

Alternative IV - Simultaneous game

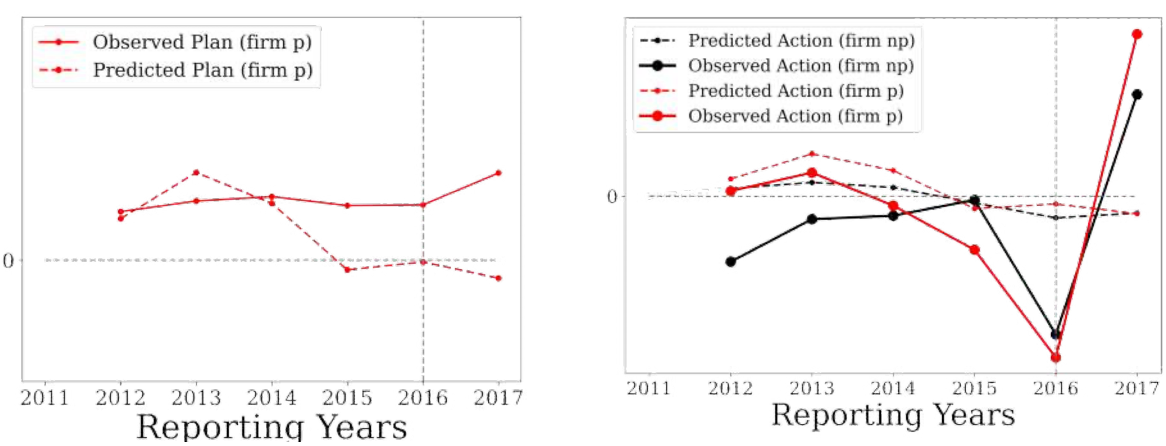

The left plot compares predicted and observed abatement plans and actions across reporting years in CDP. Predicted moments are generated from a simultaneous equilibrium model as the one described in section B.3. Results are reported in the fourth column of Table 6. Thick (dashed) lines refer to observed (model-implied) moments, red-circle (black-star) lines refer to the subset of firms with (without) abatement plans respectively.

Figure 17

Counterfactuals - Effect of Uncertainty

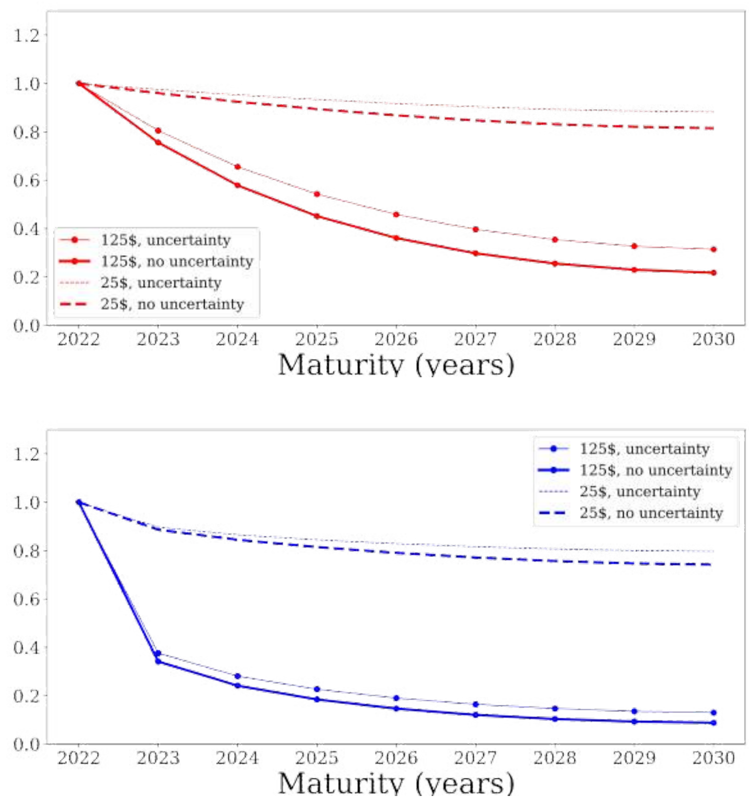

The plots show the optimal carbon emissions path of the firm with plans $\{\xi^{*,p}_{t,\tau}\}_{\tau=t,...T}$ as generated by the single-firm model (top plot, red lines) and the two-firm model (bottom plot, blue lines) respectively. Dashed and dotted lines refer to the emissions path generated by the models when the average levy $\bar{\lambda}^p = \bar{\lambda}_{np} = 25\$/\text{mtCO}_2\text{e}$ and $\bar{\lambda}^p = \bar{\lambda}^{np} = 125\$/\text{mtCO}_2\text{e}$ respectively. Thin and thick lines refer to the emissions path generated by the models when the uncertainty parameters are inferred from the CDP data and when $\sigma^p_t = \sigma^{np}_t = 0\$/\text{mtCO}_2\text{e}$ respectively. The strategic parameters are reported in Table 2, Panel B, third column in the main text. The remainder of model parameters in input are reported in Table 2, Panel A in the main text.

Figure 18

Single firm with rolling maturity

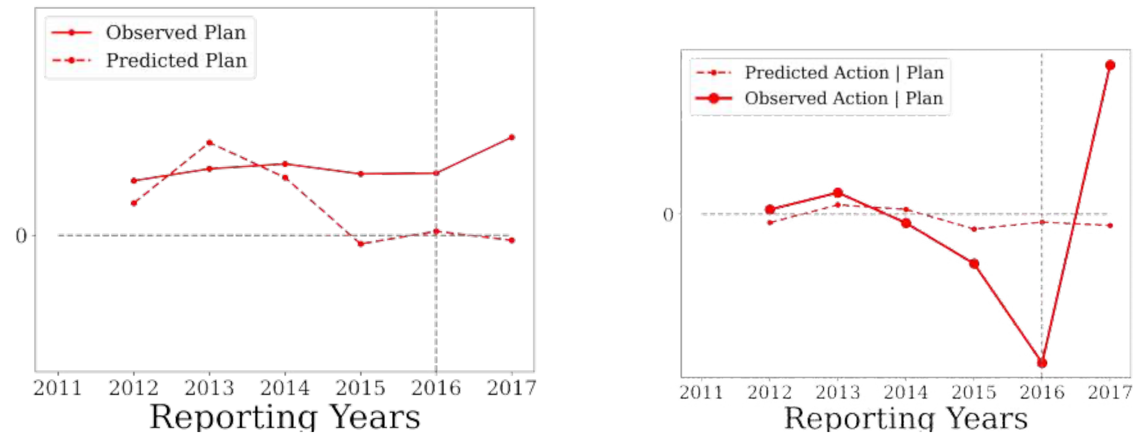

The left plot compares predicted and observed abatement plans and actions across reporting years in CDP. Calibrated parameters are reported in the second column of Panel B, Table 2, main text. Thick (dashed) lines refer to observed (model-implied) moments, red-circle (black-star) lines refer to the subset of firms with (without) abatement plans respectively.

Figure 19

Stackelberg with rolling maturity

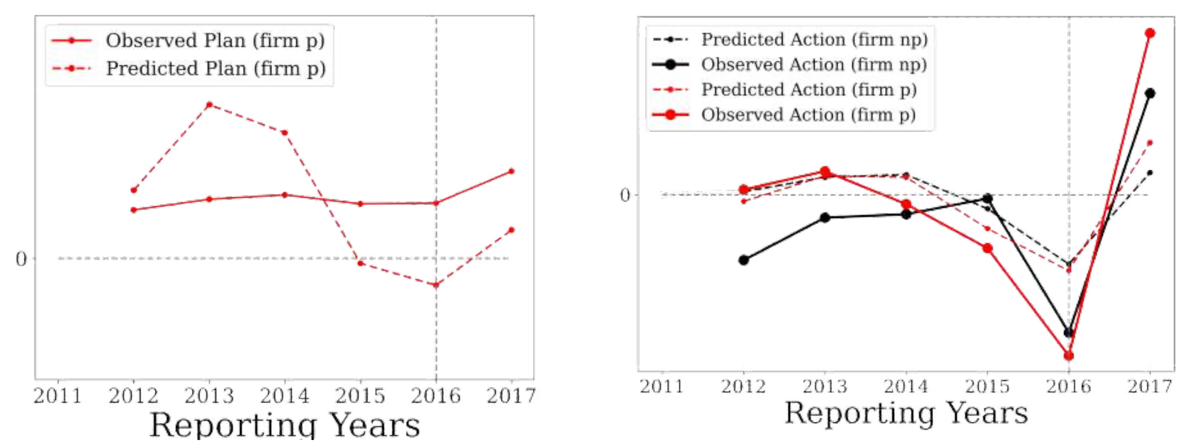

The left plot compares predicted and observed abatement plans and actions across reporting years in CDP. Calibrated parameters are reported in the fourth column of Panel B, Table 2, main text. Thick (dashed) lines refer to observed (model-implied) moments, red-circle (black-star) lines refer to the subset of firms with (without) abatement plans respectively.

B Theory Appendix

B.1 Single - Firm Model

Solving the single-firm model. We solve for the current optimal abatement rate, which we express as x_t^*, for each time $t = 0, \dots T$, starting from the terminal date.
At time $t = T$, firm value reads

$$\begin{aligned} V_T &= \max_{x_T} \pi_T - \lambda \xi_T \\ &= \max_{x_T} \omega k_{T-1}(1 - x_T) - \frac{1}{2}\phi x_T^2 k_{T-1} - \lambda \eta k_{T-1}(1 - x_T). \end{aligned} \tag{1}$$

From the first order condition (FOC) with respect to x_T, one gets

$$x_T^* = \frac{\eta\lambda}{\phi} - \frac{\omega}{\phi}. \tag{2}$$

For time $t < T$, the optimal abatement rate x_t^* maximizes the Bellman equation

$$V_t = \max_{x_t} \pi_t + \beta \mathbb{E}_t[V_{t+1}], \tag{3}$$

the FOC in (3) with respect to x_t gives, after rearrangement,[1]

$$\begin{aligned} x_t^* &= -\frac{1}{\phi}\beta \frac{\partial \mathbb{E}_t[V_{t+1}]}{\partial k_t} - \frac{\omega}{\phi} \\ &= -\frac{1}{\phi}\beta \mathbb{E}_t[\frac{\partial V_{t+1}}{\partial k_t}] - \frac{\omega}{\phi} \end{aligned} \tag{4}$$

Iterating (3) of one time period, one has

$$\frac{\partial V_{t+1}}{\partial k_t} = \omega(1 - x_{t+1}^*) - \frac{1}{2}\phi(x_{t+1}^*)^2 + \beta \mathbb{E}_{t+1}[\frac{\partial V_{t+2}}{\partial k_{t+1}}](1 - x_{t+1}^*) \tag{5}$$

iterating (4) to get $\mathbb{E}_{t+1}[\frac{\partial V_{t+2}}{\partial k_{t+1}}]$ and rearranging

$$\beta \mathbb{E}_{t+1}[\frac{\partial V_{t+2}}{\partial k_{t+1}}] = -\phi x_{t+1}^* - \omega \tag{6}$$

which substituting back into (5) gives

$$\frac{\partial V_{t+1}}{\partial k_t} = -\phi x_{t+1}^* + \frac{1}{2}\phi(x_{t+1}^*)^2 \tag{7}$$

[1] We interchanged the derivative with the mean as the capital k_t is known at time t and therefore independent of the argument in the expectation operator.

hence substituting back into (4)

$$\begin{aligned} x_t^* &= -\frac{1}{\phi}\beta\mathbb{E}_t[-\phi x_{t+1}^* + \frac{1}{2}\phi(x_{t+1}^*)^2] - \frac{\omega}{\phi} \\ &= \beta\mathbb{E}_t[x_{t+1}^* - \frac{1}{2}(x_{t+1}^*)^2] - \frac{\omega}{\phi} \end{aligned} \tag{8}$$

which proves the result. Solving explicitly by backward induction:

- For $t = T-1$, substituting (2) into (8) we have

$$\begin{aligned} x_{T-1}^* &= \beta\mathbb{E}_{T-1}[(\frac{\eta\lambda}{\phi} - \frac{\omega}{\phi}) - \frac{1}{2}(\frac{\eta\lambda}{\phi} - \frac{\omega}{\phi})^2] - \frac{\omega}{\phi} \\ &= \beta\mathbb{E}[(\frac{\eta\lambda_{T-1}}{\phi} - \frac{\omega}{\phi}) - \frac{1}{2}(\frac{\eta\lambda_{T-1}}{\phi} - \frac{\omega}{\phi})^2] - \frac{\omega}{\phi} \end{aligned} \tag{9}$$

  Since the belief is normally distributed as $\lambda_{T-1} \sim \mathcal{N}(\bar{\lambda}_{T-1}, \sigma_{T-1})$, computing the expectation we get the explicit expression

$$x_{T-1}^* = \beta\Big((\frac{\bar{\lambda}_{T-1}\eta}{\phi} - \frac{\omega}{\phi}) - \frac{1}{2}(\frac{\bar{\lambda}_{T-1}\eta - \omega}{\phi})^2 - \frac{1}{2}\frac{\eta^2}{\phi^2}\sigma_{T-1}^2\Big) - \frac{\omega}{\phi}. \tag{10}$$

- Consider $t = T-2$. From (8) we have

$$x_{T-2}^* = \beta\mathbb{E}_{T-2}[x_{T-1}^* - \frac{1}{2}(x_{T-1}^*)^2] - \frac{\omega}{\phi} \tag{11}$$

  Since $\lambda_{T-1}|T-2 \equiv \lambda_{T-2}$, then $\bar{\lambda}_{T-1}$ and σ_{T-1}^2 in (10) are constant at time $T-2$ and equal to $\bar{\lambda}_{T-1} = \bar{\lambda}_{T-2}$ and $\sigma_{T-1}^2 = \sigma_{T-2}^2$. The abatement rate at time $t = T-2$ is therefore

$$\begin{aligned} x_{T-2}^* &= \beta\mathbb{E}_{T-2}[x_{T-1}^* - \frac{1}{2}(x_{T-1}^*)^2] - \frac{\omega}{\phi} \\ &= \beta\Big[\beta\Big((\frac{\bar{\lambda}_{T-2}\eta}{\phi} - \frac{\omega}{\phi}) - \frac{1}{2}(\frac{\bar{\lambda}_{T-2}\eta - \omega}{\phi})^2 - \frac{1}{2}\frac{\eta^2}{\phi^2}\sigma_{T-2}^2\Big) - \frac{\omega}{\phi}.. \\ &.. - \frac{1}{2}(\beta\Big((\frac{\bar{\lambda}_{T-2}\eta}{\phi} - \frac{\omega}{\phi}) - \frac{1}{2}(\frac{\bar{\lambda}_{T-2}\eta - \omega}{\phi})^2 - \frac{1}{2}\frac{\eta^2}{\phi^2}\sigma_{T-2}^2\Big) - \frac{\omega}{\phi})^2\Big] - \frac{\omega}{\phi} \end{aligned} \tag{12}$$

  The same argument can be iterated for previous maturities $t = 0 \dots T-2$.

Upward-sloping abatement plan. We want to show that

$$\mathbb{E}_t[x_\tau^*] - \mathbb{E}_t[x_{\tau+1}^*] < 0. \tag{13}$$

for each maturity $\tau = t \dots T-1$. From (8) one has

$$x^*_\tau = \beta \mathbb{E}_\tau[x^*_{\tau+1} - \frac{1}{2}(x^*_{\tau+1})^2] - \frac{\omega}{\phi} \tag{14}$$

applying $\mathbb{E}_t$ on both sides and substituting it into (13), the inequality becomes

$$\mathbb{E}_t[x^*_\tau](\beta - 1) - \frac{1}{2}\mathbb{E}_t[(x^*_{\tau+1})^2] - \frac{\omega}{\phi} < 0 \tag{15}$$

from the jensen's inequality, a sufficient condition for (15) to hold is that

$$\mathbb{E}_t[x^*_\tau](\beta - 1) - \frac{1}{2}\mathbb{E}_t[x^*_{\tau+1}]^2 - \frac{\omega}{\phi} < 0 \tag{16}$$

the latter is satisfied in the admissible range $\{x^*_\tau, x^*_{\tau+1}\} \in [-1, 1]$ provided $\beta > \frac{1}{2}$ and $\omega, \phi > 0$.

Expected abatement rate $\mathbb{E}_t[x^*_\tau]$ as a function of belief's uncertainty σ^2_t. We want to show that

$$\frac{\partial \mathbb{E}_t[x^*_\tau]}{\partial \sigma^2_t} < 0 \quad \text{for } \tau \in t, \dots T-1. \tag{17}$$

Let us start with $\tau = T-1$. Recalling the expression in (10) one gets

$$\frac{\partial \mathbb{E}_t[x^*_{T-1}]}{\partial \sigma^2_t} = -\beta \frac{1}{2}\frac{\eta^2}{\phi^2} < 0 \tag{18}$$

Let us now assume that (17) is true for a certain $\tau = k+1$. Applying again (8) we get:

$$\begin{aligned}\frac{\partial \mathbb{E}_t[x^*_k]}{\partial \sigma^2_t} &= \beta\big(\frac{\partial \mathbb{E}_t[x^*_{k+1}]}{\partial \sigma^2_t} - \frac{1}{2}\frac{\partial \mathbb{E}_t[(x^*_{k+1})^2]}{\partial \sigma^2_t}\big) \\ &= \beta\big(1 - \frac{1}{2}\mathbb{E}_t[x^*_{k+1}]\big)\frac{\partial \mathbb{E}_t[x^*_{k+1}]}{\partial \sigma^2_t} - \beta\frac{1}{2}\frac{\partial Var(x^*_{k+1})}{\partial \sigma^2_t}\end{aligned} \tag{19}$$

since $\frac{\partial Var(x^*_{k+1})}{\partial \sigma^2_t} = 0$,[2] a sufficient condition for (17) to be satisfied is that $\mathbb{E}_t[x^*_{k+1}] < 2$.

B.2 Two-Firm Model

Solving the leader-follower model. At time $t = T$, the leader's value reads

$$V^l_T = \max_{x^l_T} \omega^l k^l_T - \frac{1}{2}\phi^l (x^l_T)^2 k^l_{T-1} + \gamma^l_T x^f_T \eta^l x^l_T k^l_{T-1} - \lambda^l \eta^l k^l_T \tag{20}$$

[2]The latter follows from our assumption that the anticipated beliefs distribution $\lambda_{k+1}|t \equiv \lambda_t$ and therefore the abatement rate x^*_{k+1} for $k+1 \leq T-1$ is a constant at time t (see the explicit expression for the abatement rate x^*_{T-2} in appendix B.1). However, the result holds more generally (i.e., for alternative beliefs distributions) as long as $\frac{\partial Var(x^*_{k+1})}{\partial \sigma^2_t} \geq 0$.

while the follower's value reads

$$V_T^f = \max_{x_T^f} \omega^f k_T^f - \frac{1}{2}\phi^f (x_T^f)^2 k_{T-1}^f + \gamma_T^f x_T^l \eta^f x_T^f k_{T-1}^f - \lambda^f \eta^f k_T^f. \tag{21}$$

Applying the FOC to (21), taking the leader's optimal rate $x_T^l = x_T^{*,l}$ as given, one gets

$$x_T^{*,f} = w_T^f x_T^{*,l} + \frac{\eta^f}{\phi^f}\lambda^f - \frac{\omega^f}{\phi^f} \tag{22}$$

with $w_T^f = \gamma_T^f \eta^f / \phi^f$. Substituting (22) into the leader's terminal profits in (20) and solving for the optimal $x_T^{*,l}$, one gets:

$$x_T^{*,l} = \frac{1}{1 - 2w_T^l w_T^f} \frac{\eta^l \lambda^l - \omega^l}{\phi^l} + \frac{w_T^l}{1 - 2w_T^l w_T^f} \frac{\eta^f \lambda^f - \omega^f}{\phi^f} \tag{23}$$

which gives the first result. Consider $t < T$. The Bellman equations for the leader-follower model read:

$$V_t^l = \max_{x_t^l} \omega^l k_t^l - \frac{1}{2}\phi^l (x_t^l)^2 k_{t-1}^l + \gamma_t^l x_t^f \eta^l x_t^l k_{t-1}^l + \beta^l \mathbb{E}_t[V_{t+1}^l], \tag{24}$$

and

$$V_t^f = \max_{x_t^f} \omega^f k_t^f - \frac{1}{2}\phi^f (x_t^f)^2 k_{\tau-1}^f + \gamma_t^f x_t^l \eta^f x_t^f k_{\tau-1}^f + \beta^f \mathbb{E}_t[V_{t+1}^f]. \tag{25}$$

Taking $x_t^{*,l}$ as given, the optimal $x_t^{*,f}$ is first derived following the same steps as in the baseline case with no externalities. It is simple to show that the optimal abatement rate of the follower satisfies:

$$x_t^{*,f} = w_t^f x_t^{*,l} + f_{t,t+1} \tag{26}$$

with $w_t^f = \frac{\gamma_t^f \eta^f}{\phi^f}$ and $f_{t,t+1}$ given by:

$$f_{t,t+1} = \beta^f \mathbb{E}_t\big[x_{t+1}^{*,f} - w_{t+1}^f x_{t+1}^{*,l} - \frac{1}{2}(x_{t+1}^{*,f})^2\big] - \frac{\omega^f}{\phi^f} \tag{27}$$

Substituting the follower's optimal abatement rate (26) into (24), and solving for the optimal leader's abatement rate, one gets

$$\begin{aligned}(1 - 2w_t^f w_t^l) x_t^{*,l} &= w_t^l f_{t,t+1} - \frac{1}{\phi^l}\beta^l \frac{\partial \mathbb{E}_t[V_{t+1}^l]}{\partial k_t^l} - \frac{\omega^l}{\phi^l} \\ &= w_t^l f_{t,t+1} - \frac{1}{\phi^l}\beta^l \frac{\mathbb{E}_t[\partial V_{t+1}^l]}{\partial k_t^l} - \frac{\omega^l}{\phi^l}.\end{aligned} \tag{28}$$

Following the same procedure as in the single-firm model, we then get:

$$\begin{aligned}\frac{\partial V_{t+1}^l}{\partial k_t^l} &= \omega^l(1-x_{t+1}^{*,l}) - \frac{1}{2}\phi^l(x_{t+1}^{*,l})^2 + \gamma_{t+1}^l\eta^l x_{t+1}^{*,l}x_{t+1}^{*,f}\dots \\ &\dots + (1-x_{t+1}^{*,l})\Big(-\omega^l - \phi^l x_{t+1}^{*,l} + \gamma_{t+1}^l\eta^l(x_{t+1}^{*,f} + w_{t+1}^f x_{t+1}^{*,l})\Big)\end{aligned} \tag{29}$$

where we used (26) to rewrite $\gamma_{t+1}^l\eta^l(2w_{t+1}^f x_{t+1}^{*,l} + f_{t+1,t+2}) = \gamma_{t+1}^l\eta^f(x_{t+1}^{*,f} + w_{t+1}^f x_{t+1}^{*,l})$. After rearrangement, this gives:

$$\frac{\partial V_{t+1}^l}{\partial k_t^l} = \frac{1}{2}\phi^l(1-2w_{t+1}^l w_{t+1}^f)(x_{t+1}^{*,l})^2 - \phi(1-w_{t+1}^l w_{t+1}^f)x_{t+1}^{*,l} + \gamma_{t+1}^l x_{t+1}^f. \tag{30}$$

Putting (30) back into (28) and solving for $x_t^{*,l}$, we finally get:

$$\begin{aligned}x_t^{*,l} &= \frac{w_t^l}{(1-2w_t^l w_t^f)}f_{t,t+1} + \beta^l\mathbb{E}_t\Big[\frac{(1-w_{t+1}^l w_{t+1}^f)x_{t+1}^{*,l} - w_{t+1}^l x_{t+1}^{*,f}}{(1-2w_t^l w_t^f)} + \dots \\ &\dots - \frac{(1-2w_{t+1}^l w_{t+1}^f)}{(1-2w_t^l w_t^f)2}(x_{t+1}^l)^2\Big] - \frac{\omega^l}{\phi^l(1-2w_t^l w_t^f)}\end{aligned} \tag{31}$$

substituting the expression for $f_{t,t+1}$ in (31) one gets

$$\begin{aligned}x_t^{*,l} &= \frac{w_t^l}{(1-2w_t^l w_t^f)}[\beta^f\mathbb{E}_t[x_{t+1}^{*,f} - w_{t+1}^f x_{t+1}^{*,l} - \frac{1}{2}(x_{t+1}^{*,f})^2]] + \beta^l\mathbb{E}_t\Big[\frac{(1-w_{t+1}^l w_{t+1}^f)x_{t+1}^{*,l} - w_{t+1}^l x_{t+1}^{*,f}}{(1-2w_t^l w_t^f)} + \dots \\ &\dots - \frac{(1-2w_{t+1}^l w_{t+1}^f)}{(1-2w_t^l w_t^f)2}(x_{t+1}^l)^2\Big] - \frac{\omega^l}{\phi^l(1-2w_t^l w_t^f)} - w_t^l\frac{\omega^f}{\phi^f(1-2w_t^l w_t^f)} \\ &= \beta^l\mathbb{E}_t\Big[\frac{(1-w_{t+1}^l w_{t+1}^f - w_t^l w_{t+1}^f\frac{\beta^f}{\beta^l})x_{t+1}^{*,l} + x_{t+1}^{*,f}(w_t^l\frac{\beta^f}{\beta^l} - w_{t+1}^l)}{(1-2w_t^l w_t^f)} + \dots \\ &\dots - \frac{(1-2w_{t+1}^l w_{t+1}^f)}{(1-2w_t^l w_t^f)2}(x_{t+1}^l)^2\Big] - \frac{\omega^l}{\phi^l(1-2w_t^l w_t^f)} - w_t^l\frac{\omega^f}{\phi^f(1-2w_t^l w_t^f)}\end{aligned} \tag{32}$$

which proves the result.

Proof of proposition 1. One notes that when $w_T^f = 0$ and $w_T^l = 0$, the leader and the follower's optimal abatement rates degenerate into the baseline case without externalities. Applying the expectation operator $\mathbb{E}_t[\cdot]$ in (23) and deriving with respect to $\bar{\lambda}_t^l$ we get:

$$\frac{\partial\mathbb{E}_t[x_T^{*,l}]}{\partial\bar{\lambda}_t^l} = \frac{\eta^l}{\phi^l}\frac{1}{1-2w_T^l w_T^f} > \frac{\eta^l}{\phi^l} = \frac{\partial\mathbb{E}_t[x_T^{*,l}]}{\partial\bar{\lambda}_t^l}\Big|_{w_T^f, w_T^l=0} \tag{33}$$

which proves the result for the leader provided the reputation weights satisfy $0 < w_T^f w_T^l < 1/2$.

Substituting the expression (23) into (22) and deriving $x_T^{*,f}$ with respect to λ_t^f, we then get:

$$\frac{\partial \mathbb{E}_t[x_T^{*,f}]}{\partial \bar{\lambda}_t^f} = \frac{\eta^f}{\phi^f}\big(\frac{w_T^f w_T^l}{1-2w_T^l w_T^f}+1\big) > \frac{\eta^f}{\phi^f} = \frac{\partial \mathbb{E}_t[x_T^{*,f}]}{\partial \bar{\lambda}_t^f}\Big|_{w_T^f, w_T^l=0} \tag{34}$$

which proves the result for the follower provided $0 < w_T^f w_T^l < 1/2$. Finally, one notes that

$$\frac{1}{1-2w_T^l w_T^f} > \frac{1-w_T^l w_T^f}{1-2w_T^l w_T^f} = \frac{w_T^f w_T^l}{1-2w_T^l w_T^f}+1 \tag{35}$$

which proves the second result given (33) and (34).

Proof of corollary 1. Deriving $\mathbb{E}_t[x_T^{*,l}]$ in (23) with respect to λ_t^f we get:

$$\frac{\partial \mathbb{E}_t[x_T^{*,l}]}{\partial \bar{\lambda}_t^f} = \frac{\eta^f}{\phi^f}\frac{w_T^l}{1-2w_T^l w_T^f} \tag{36}$$

which recalling (34), implies that

$$\frac{\partial \mathbb{E}[x_T^{*,l}]}{\partial \bar{\lambda}_t^f} > \frac{\partial \mathbb{E}[x_T^{*,f}]}{\partial \bar{\lambda}_t^f} \quad \longleftrightarrow \quad w_T^l > 1 - w_T^l w_T^f \quad \longleftrightarrow \quad w_T^l > (1+w_T^f)^{-1} \tag{37}$$

which proves the result for the terminal maturity T.

Consider now that the result holds for the near maturity $t+1 << T$.

Then from (31) and (26)

$$\begin{aligned}
\frac{\partial x_t^{f,*}}{\partial \bar{\lambda}_t^f} &= w_t^f \frac{\partial x_t^{l,*}}{\partial \bar{\lambda}_t^f} + \frac{\partial f_{t,t+1}}{\partial \bar{\lambda}_t^f} \\
&= \big(\frac{1-w_t^f w_t^l}{1-2w_t^l w_t^f}\big)\frac{\partial f_{t,t+1}}{\partial \bar{\lambda}_t^f} + w_t^f \beta^l \frac{\partial}{\partial \bar{\lambda}_t^f}\mathbb{E}_t\big[\frac{(1-w_{t+1}^l w_{t+1}^f)x_{t+1}^{l,*} - w_{t+1}^l x_{t+1}^{f,*}}{(1-2w_t^l w_t^f)} - \frac{(1-2w_{t+1}^l w_{t+1}^f)}{2(1-2w_t^l w_t^f)}(x_{t+1}^{l,*})^2\big] \\
&\approx \big(\frac{1-w_t^f w_t^l}{1-2w_t^l w_t^f}\big)\beta^f(1-\mathbb{E}_t[x_{t+1}^{f,*}] - \frac{w_{t+1}^f w_{t+1}^l}{1-w_{t+1}^f w_{\tau+1}^l})\frac{\partial \mathbb{E}_t[x_{t+1}^{f,*}]}{\partial \bar{\lambda}_t^f} - w_\tau^f \beta^l (\mathbb{E}_t[x_{t+1}^{l,*}]\frac{1-2w_{\tau+1}^l w_{\tau+1}^f}{1-2w_\tau^l w_\tau^f})\frac{\mathbb{E}_t[x_{t+1}^{l,*}]}{\partial \bar{\lambda}_t^f}
\end{aligned} \tag{38}$$

where for the last equation we used twice the approximation $\frac{\partial \mathbb{E}_t[x_{t+1}^{l\ *}]}{\partial \bar{\lambda}_t^f} \approx \frac{w_{t+1}^l}{1-w_{t+1}^l w_{t+1}^f}\frac{\partial \mathbb{E}_t[x_{t+1}^{f\ *}]}{\partial \bar{\lambda}_t^f}$ and the fact $Var(x^{f,*}_{t+1}), Var(x^{l,*}_{t+1}) = 0$ for $t+1 << T$[3]. From (31), we also get

$$\begin{aligned}\frac{\partial x^{l,*}_t}{\partial \bar{\lambda}^f_t} &= \frac{w^l_t}{1-2w^l_t w^f_t}\frac{\partial f_{t,t+1}}{\partial \bar{\lambda}^f_t} + \beta^l \frac{\partial}{\partial \bar{\lambda}^f_t}\mathbb{E}_t\Big[\frac{(1-w^l_{t+1}w^f_{t+1})x^{l,*}_{t+1} - w^l_{t+1}x^{f,*}_{t+1}}{(1-2w^l_t w^f_t)} - \frac{(1-2w^l_{t+1}w^f_{t+1})}{(1-2w^l_t w^f_t)2}(x^{l,*}_{t+1})^2\Big] \\ &\approx \frac{w^l_\tau}{1-2w^l_\tau w^f_\tau}\beta^f(1-\mathbb{E}_t[x^{f,*}_{t+1}] - \frac{w^f_{t+1}w^l_{t+1}}{1-w^f_{t+1}w^l_{\tau+1}})\frac{\partial \mathbb{E}_t[x^{f,*}_{t+1}]}{\partial \bar{\lambda}^f_t} - \beta^l(\mathbb{E}_t[x^{l,*}_{t+1}]\frac{1-2w^l_{\tau+1}w^f_{\tau+1}}{1-2w^l_\tau w^f_\tau})\frac{\mathbb{E}_t[x^{l,*}_{t+1}]}{\partial \bar{\lambda}^f_t}\end{aligned} \tag{39}$$

From (38) and (39), it derives that a sufficient condition for the current abatement rates to verify

$$\frac{\partial x^{l,*}_t}{\partial \bar{\lambda}^f_t} > \frac{\partial x^{f,*}_t}{\partial \bar{\lambda}^f_t} \tag{40}$$

is that $\mathbb{E}[x^{l,*}_{t+1}] < 0, \mathbb{E}[x^{f,*}_{t+1}] < 0$ and $w^l_t > (1+w^f_t)^{-1}$ continues to hold.

Terminal abatement rates in the simultaneous model. At any time period $t \leq T$, holding fixed the model parameters $\{\beta^f, \omega^f, \eta^f, \phi^f, \beta^l, \omega^l, \eta^l, \phi^l\}$, and payoff externality $|\gamma^f_T \gamma^l_T| \leq \phi^f \phi^l$, it is simple to show that the expected terminal abatement rates $\mathbb{E}_t[x^{*,l}_T]$ and $\mathbb{E}_t[x^{*,f}_T]$ solving a simultaneous equilibrium model satisfy

$$\mathbb{E}_t[x^{*,l}_T] = \frac{1}{1-w^l_T w^f_T}\frac{\eta^l \bar{\lambda}^l_t - \omega^l}{\phi^l} + \frac{w^l_T}{1-w^l_T w^f_T}\frac{\eta^f \bar{\lambda}^f_t|l - \omega^f}{\phi^f} \tag{41}$$

and

$$\mathbb{E}_t[x^{*,f}_T] = \frac{1}{1-w^l_T w^f_T}\frac{\eta^f \bar{\lambda}^f_t - \omega^f}{\phi^f} + \frac{w^f_T}{1-w^l_T w^f_T}\frac{\eta^l \bar{\lambda}^l_t|f - \omega^l}{\phi^l} \tag{42}$$

with $\bar{\lambda}^f_t|l$ ($\bar{\lambda}^l_t|f$) denoting firm l's (firm f's) belief about firm f's (firm l's) belief about the levy respectively at time t. From (41) and (42), one notes that

$$\begin{aligned}\frac{\partial \mathbb{E}_t[x^{*,l}_T]}{\partial \bar{\lambda}^l_t} &= \frac{1}{1-w^l_T w^f_T}\frac{\eta^l}{\phi^l} > \frac{\eta^l}{\phi^l} \\ \frac{\partial \mathbb{E}_t[x^{*,f}_T]}{\partial \bar{\lambda}^f_t} &= \frac{1}{1-w^l_T w^f_T}\frac{\eta^f}{\phi^f} > \frac{\eta^f}{\phi^f}\end{aligned} \tag{43}$$

meaning that firms' reactions to variations in their own beliefs are still greater than their baseline reactions in absence of externalities. On the other hand, under the same specification of beliefs as in the leader-follower model, i.e. $\lambda^f_t|l = \lambda^f_t$ and $\lambda^l_t|f = \lambda^f_t$, one has that

$$\frac{\partial \mathbb{E}_t[x^{*,l}_T]}{\partial \bar{\lambda}^f_t} > \frac{\partial \mathbb{E}_t[x^{*,f}_T]}{\partial \bar{\lambda}^f_t} \longleftrightarrow w^l_T > 1 + w^f_T \tag{44}$$

[3]Again, this derives from our assumption on the conditional distribution of future beliefs.

One notes that condition (44) is far more stringent than condition (37), required for Corollary 1 to hold. In particular, (44) does not hold in the case of a symmetric reputation externality, i.e. $w^f_T = w^l_T$, whereas (37) would still be admissible provided $w^f_T = w^l_T > 0.12$.

Proof of proposition 2. We outline first the case of the leader. Assume for simplicity $\beta^l \approx \beta^f \approx 1$ and consider an intermediate maturity $t < \tau < T-1$. The expression in (31) can be simplified to notation as

$$\mathbb{E}_t[x^{l,*}_\tau] = b^l \mathbb{E}_t[x^{l,*}_{\tau+1}] - a^l \mathbb{E}_t[x^{l,*}_{\tau+1}]^2 - c^l \tag{45}$$

where the coefficient of the linear term is $b^l = \frac{(1-w^l_{\tau+1}w^f_{\tau+1} - w^l_\tau w^f_{\tau+1})}{1-2w^l_\tau w^f_\tau}$, the coefficient of the quadratic term is $a^l = \frac{1}{2}\frac{1-2w^l_{\tau+1}w^f_{\tau+1}}{1-2w^l_\tau w^f_\tau}$ and the coefficient of the constant term is $c^l = \frac{\omega^l}{\phi^l(1-2w^l_\tau w^f_\tau)} + w^l_\tau \frac{\omega^f}{\phi^f(1-2w^l_\tau w^f_\tau)} - \frac{x^f_{\tau+1}(w^l_\tau - w^l_{\tau+1})}{(1-2w^l_\tau w^f_\tau)}$. We therefore have that

$$\mathbb{E}_t[x^{l,*}_\tau] > \mathbb{E}_t[x^{l,*}_{\tau+1}] \iff (b^l - 1)\mathbb{E}_t[x^{l,*}_{\tau+1}] - a^l \mathbb{E}_t[x^{l,*}_{\tau+1}]^2 - c^l > 0 \tag{46}$$

which holds whenever $\mathbb{E}_t[x^{l,*}_{\tau+1}]$ falls in the range

$$\mathbb{E}_t[x^{l,*}_{\tau+1}] \in \Big(0, b^l - 1 + \frac{\sqrt{(b^l-1)^2 - 4a^l c^l}}{2a^l}\Big] \tag{47}$$

A sufficient condition for the upperbound to be strictly positive, which in turns implies an inverted order of abatement $\mathbb{E}_t[x^{l,*}_\tau] > \mathbb{E}_t[x^{l,*}_{\tau+1}] > 0$, is that

$$(b^l - 1) > 2\sqrt{a^l c^l} \tag{48}$$

Assume that $w^f_{\tau+1} = \rho w^f_\tau$ and $w^l_{\tau+1} = \rho w^l_\tau$. The first thing to note is that, if $\rho \geq 1$ (i.e., reputation weights have an increasing time-path), then $b^l < 1$, $a^l > 0$ and $c^l > 0$, which implies that condition (46) is never satisfied. Consider $\rho < 1$. Condition (48) in turn requires that the weights satisfy

$$w^l_\tau w^f_\tau > \frac{1}{(2-(1+\rho^2))}\sqrt{2(1-2\rho^2 w^l_\tau w^f_\tau)(\frac{\omega^l}{\phi^l} + w^l_\tau(\frac{\omega^f}{\phi^f} - \mathbb{E}_t[x^{f,*}_{\tau+1}](1-\rho)))} \tag{49}$$

In the extreme case in which $\rho = 0$, the inequality above simplifies to

$$w^l_\tau w^f_\tau > \sqrt{(\frac{\omega^l}{\phi^l} + w^l_\tau(\frac{\omega^f}{\phi^f} - \mathbb{E}_t[x^{f,*}_{\tau+1}]))} \tag{50}$$

which is easily satisfied for a realistic range of model parameters $\{\omega^f, \omega^l, \phi^f, \phi^l, \beta^l, \beta^f, \eta^f, \eta^l\}$. The result then follows from continuity of the abatement rate as a function of ρ. In the case of the follower, recalling equation (26), one has that

$$\mathbb{E}_t[x_\tau^{f,*}] > \mathbb{E}_t[x_{\tau+1}^{f,*}] \quad \longleftrightarrow \quad w_\tau^f \mathbb{E}_t[x_\tau^{l,*}] + \mathbb{E}_t[f_{\tau,\tau+1}] - \mathbb{E}_t[x_{\tau+1}^{f,*}] > 0 \tag{51}$$

with $f_{\tau,\tau+1}$ as in (27). Recalling (13) and (15), it is simple to show that $f_{\tau,\tau+1} - x_{\tau+1}^{f,*} < 0$. Hence, a necessary condition for (51) to hold is that the leader's abatement rate $\mathbb{E}_t[x_\tau^{l,*}]$ is large *and* the reputation externality of the follower (i.e., the weight w_τ^f that the follower puts on the leader's actions) is also large.

C Construction of the Dataset

We employ detailed data on firms' voluntary disclosures from the Carbon Disclosure Project (CDP). CDP sends out environment-related questionnaires to firms each year, and we obtain firms' responses from 2011 to 2017. In total, over 3,000 publicly listed firms from different sectors and countries respond to the questionnaires. We focus on the CDP subsample of publicly listed North American firms that are also in the panel available from the CRSP/COMPUSTAT database between 2010 and 2016.[4] We find a total of roughly 700 CDP firms which match with the selected CRSP/COMPUSTAT sample,[5] but not all matched firms report all variables necessary for our analysis, and some provide inconsistent disclosures. As detailed below, we clean raw disclosures of climate risks, carbon emissions, and emissions reduction targets in order to get firm-level metrics of beliefs, actions, and plans that survive internal consistency checks, and can be validated against external data.

Construction of actions. Actions are measured as percentage changes in carbon emissions as reported by the firms in the dataset, Raw disclosures of CO_2 equivalent (CO_2e) emissions are from CDP data worksheets that pertain to emissions data. For each firm i and reporting year t, we compute emissions as

$$Emissions_{i,t} = Scope1_{i,t} + Scope2_{i,t} \tag{52}$$

Where $Scope1$ denotes direct emissions (e.g. for 2017 we look at the sheet "CC8. Emissions Data") and $Scope2$ denotes indirect emissions (e.g. for 2017 we look at the sheet "CC83a. Emissions Data"). In each reporting year, firms can provide multiple estimates of direct or in-

[4]We keep only firms in the CRSP/COMPUSTAT North America (Fundamental Annual) dataset with non-missing Tickers within the 2010–2016 accounting period. We lag the information from CRSP/COMPUSTAT by one year to account for a time window between the filling and the final release of the CDP questionnaires.

[5]Matches are computed at the Ticker level.

direct emissions, i.e., there are different vintages of the data. To avoid overlapping disclosures in the time-series, we select only disclosures of carbon emissions related to the latest accounting year: this can either be one year prior to, or the same year as, the reporting year, depending on the date of submission of the firm's data.

Construction of beliefs. Raw disclosures of regulation risk are from CDP sheets related to climate change risks driven by regulation (e.g. for 2017 we look at the sheet "CC5.1a" on risks driven by changes in regulation). Firms' descriptions of theregulation risk they face vary across firms and reporting years in the dataset. The word cloud in Figure 2 highlights some of the most frequent words that appear in the unstructured text field in the pooled dataset in which firms describe the specific regulation risks. As the figure shows, firms refer to several different types of climate regulation. Firms most frequently stated type of anticipated climate regulation is, as one might expect, fuel energy taxes and carbon taxes, to which follows a cap and trade system. Firms also refer to emissions reporting programmes as a third category of potential climate regulation. These text disclosures partly motivate our modelling choice, described in the paper, of regulation in the form of a carbon levy.

Unlike carbon emissions, risk disclosures always refer to the latest accounting year available. However, firms usually describe multiple types of regulation events as they differentiate, for example, at the plant or business unit levels. For each firm i and reporting year t, we therefore compute the aggregate belief metric as

$$\Lambda_{i,t} = \sum_{k=0}^{k_{it}} \beta_{i,t}^{H_k - t} M_k q_k \tag{53}$$

where $k = 0, \dots k_{it}$ varies over the number of events disclosed by firm i in reporting year t, while M_k and q_k are the magnitude and likelihood respectively of each event k.

Construction of plans. Raw disclosures of emissions reduction targets are from CDP sheets related to targets and initiatives (e.g. for 2017 we look at the sheet "CC3.1a" on absolute emissions reduction targets). As for climate risks, firms can provide multiple targets if they include emissions targets set in previous reporting years that might (or might not) be still active in the current reporting year. When this occurs, for each firm i and reporting year t, we compute the aggregate metric of abatement plans as:

$$plan_{i,t} = \sum_{k=0}^{k_{it}} \sum_{\tau=t+1}^{T_k} \beta^{\tau - t} \frac{\hat{x}_k}{T_k - t_k} \tag{54}$$

where $k = 0, \dots k_{it}$ ranges over the total number of targets reported by the firm that are still active in the reporting year t (i.e. $t < T_k$), while $\frac{\hat{x}_k}{T_k - t_k}$ is the average yearly rate of emissions reduction relative to target k, with $t_k \leq t$ the baseline year of the target. To get rid of inconsistent disclosures, we trim the distribution of the reduction rate $\hat{x}_k$ so that it lies between $0 \leq \hat{x}_k \leq 1$.

The final dataset (consisting of 446 unique firms that report carbon emissions and regulation risk for *at least* two consecutive years) is reported in the third column of Table 1.

Construction of the Out-of-Sample dataset. To conduct our out-of-sample validation exercise, we extend the CDP dataset in Table 1 to include U.S. public firms' responses from 2018 and 2019. Over these two years, CDP implemented a set of changes to make the questionnaires more aligned with the recommendations of the Task Force on Climate-Related Financial Disclosures (TCFD), established in 2016. Below, we describe the major changes to the dataset arising as a result of these changes, as well as adjustments that we implemented to our construction of the data as a result of these changes to make the out-of-sample data consistent with our treatment of the in-sample data.

First, regulation risk in the later period is part of a broader classification of climate-related risks, collectively labelled "climate transition risks". These risks include: marked shifts in consumer tastes, reputation risks from negative stakeholder feedback, technology risk due to forced substitution of products and services, and policy risk from new or existing regulations. To preserve continuity with the previous setting, we select firms' disclosures related only to this policy risk component. Second, time horizons of climate-related risks are not tied to numeric ranges as in the earlier data. That is, firms in the later period of the data choose from options: current, short-term, medium-term, and long-term horizons. To preserve continuity with the previous setting, we therefore translate these responses into the time ranges provided by CDP before 2018. More specifically, current horizon is translated into 0 to 1 years from the time of reporting, short-term horizon to 1 to 3 years from reporting, medium-term horizon to 3 to 6 years from reporting, and long-term horizon to beyond 6 years from reporting. Finally, while responses have remained unaltered as far as emissions reduction targets and total carbon emissions are concerned, a number of firms reporting CDP questionnaires in 2018 and 2019 have taken the option to hide their emissions data. As a consequence, of the 375 firms reporting emissions and risks in 2017 (see Table 1), only 330 (331) respectively report emissions in 2018 (2019 respectively), of which 177 (187) of these also report targets. We focus on the data for this reduced number of firms in our out-of-sample exercise.

D Calibration

We calibrate the model on the average firm with plans (firm p) and the average firm without plans (firm np) reporting in the selected dataset between $t = 2011$ and $t = 2016$.

Discount rate β. We measure the one-period discount rate as the inverse of the weighted average cost of capital of firm p and firm np respectively

$$\begin{aligned}\hat{\beta}^p &= \frac{1}{2016-2011}\sum_{t=2011}^{2016}\Big(\frac{1}{N_t^p}\sum_{i=0}^{N_t^p}\frac{1}{1+WACC_{i,t}}\Big)\\ \hat{\beta}^{np} &= \frac{1}{2016-2011}\sum_{t=2011}^{2016}\Big(\frac{1}{N_t-N_t^p}\sum_{i=0}^{N_t-N_t^p}\frac{1}{1+WACC_{i,t}}\Big)\end{aligned} \tag{55}$$

where $\mathrm{WACC}_{i,t}$ is firm i's cost of capital in reporting year t from Bloomberg Equity.

Pollution intensity η. We measure the pollution intensity as firm p and firm np's ratio of total reported emissions to total assets as

$$\begin{aligned}\hat{\eta}^p &= \frac{1}{2016-2011}\sum_{t=2011}^{2016}\Big(\frac{1}{N_t^p}\sum_{i=0}^{N_t^p}\frac{Emissions_{i,t}}{TotalAssets_{i,t}}\Big)\\ \hat{\eta}^{np} &= \frac{1}{2016-2011}\sum_{t=2011}^{2016}\Big(\frac{1}{N_t-N_t^p}\sum_{i=0}^{N_t-N_t^p}\frac{Emissions_{i,t}}{TotalAssets_{i,t}}\Big)\end{aligned} \tag{56}$$

with $Emissions_{i,t}$ firm i's reported emissions (Scope1 + Scope 2) as collected from CDP in reporting year t, and $TotalAssets_{i,t}$ firm i's total assets in reporting year t from CRSP/Compustat; **Productivity constant ω.** We measure the productivity constant ω as the net income from sales divided by total assets[6]

$$\begin{aligned}\hat{\omega}^p &= \frac{1}{2016-2011}\sum_{t=2011}^{2016}\Big(\frac{1}{N_t^p}\sum_{i=0}^{N_t^p}\frac{NetIncome_{i,t}}{TotalAssets_{i,t}}\Big)\\ \hat{\omega}^{np} &= \frac{1}{2016-2011}\sum_{t=2011}^{2016}\Big(\frac{1}{N_t-N_t^p}\sum_{i=0}^{N_t-N_t^p}\frac{NetIncome_{i,t}}{TotalAssets_{i,t}}\Big)\end{aligned} \tag{57}$$

with $NetIncome_{i,t}$ firm i's net income before extraordinary items in reporting year t from CRSP/Compustat.

[6]Recall that in the model specification we normalize operating costs to zero.

Date of the regulation event T. We either use a fixed terminal date $\hat{T} = 2020$, or a rolling terminal date $\hat{T}_t = t + \hat{h}$ with $\hat{h}$ the average target horizon reported by firm p in the dataset

$$\hat{h} = \frac{1}{2016 - 2011} \sum_{t=2011}^{2016} \Big(\frac{1}{N_t^p} \sum_{i=0}^{N_t^p} \frac{1}{k_{it}} \sum_{k=0}^{k_{it}} (T_k - t_k) \Big) \tag{58}$$

where k_{it} are the number of targets reported by firm i in group p in reporting year t.

Beliefs about regulation λ_t. For each reporting year $t = 2011, \ldots 2016$, we assume that the beliefs about the levy are a linear transformation of reported beliefs

$$\begin{aligned} \lambda_t^p &= \bar{\lambda}^p (1 + \alpha \bar{(} \Lambda_t^p - \bar{\Lambda}^p)) \\ \lambda_t^{np} &= \bar{\lambda}^{np} (1 + \alpha (\Lambda_t^{np} - \bar{\Lambda}^{np})) \end{aligned} \tag{59}$$

where $\alpha = 5$ is an amplification parameter to account for the fact that observed beliefs range between 0 and 5 whereas the levy is continuous. The distribution of reported beliefs is approximated by $\Lambda_t^p \approx \mathcal{N}(\Lambda_t^p, \Sigma_t^p)$ and $\Lambda_t^{np} \approx \mathcal{N}(\Lambda_t^{np}, \Sigma_t^{np})$, with mean

$$\bar{\Lambda}_t^p = \frac{1}{N_t^p} \sum_{i=1}^{N_t^p} \Lambda_{i,t}, \quad \Lambda_t^{np} = \frac{1}{N_t - N_t^p} \sum_{i=1}^{N_t - N_t^p} \Lambda_{i,t}, \tag{60}$$

and standard deviation

$$\Sigma_t^p = \sqrt{\frac{1}{N_t^p} \sum_{i=1}^{N_t^p} (\Lambda_{i,t}^2 - (\bar{\Lambda}_t^p)^2)}, \quad \Sigma_t^{np} = \sqrt{\frac{1}{N_t - N_t^p} \sum_{i=1}^{N_t - N_t^p} (\Lambda_{i,t}^2 - (\bar{\Lambda}_t^{np})^2)}. \tag{61}$$

The standard deviation inputed in the model is approximated by its linear interpolation to reduce measurement error (see Figure 12).

Moments Matching We estimate the model parameters so that to minimize squared distances with a set of observed moments. For the two-firm models, denoting $b = \{\lambda^p, \lambda^{np}, \gamma^p, \gamma^{np}\}$ the vector of parameters to estimate, and $\hat{\Theta} = \{\hat{\beta}^p, \hat{\beta}^{np}, \hat{\omega}^p, \hat{\omega}^{np}, \hat{\eta}^p, \hat{\eta}^{np}, \hat{\phi}^p, \hat{\phi}^{np}, \hat{h}\}$ the vector of parameters in input, the estimation exercise can be summarized as[7]

$$\hat{b} = arg \min_b g(b; \hat{\Theta})\, W g(b; \hat{\Theta}) \tag{62}$$

where W is the identity matrix, $g(b;\hat{\Theta}) = \frac{1}{N} \sum_{n=1}^{N} (y_n(b;\hat{\Theta}) - \hat{y}_n)$, with $y_n(b;\hat{\Theta})$ and $\hat{y}_n$ the 18-

[7]In the case of the two-firm model, the estimated $\hat{b}$ refers to the mean of the estimates across 100 draws of the initial guess b_0 from a uniform distribution with domain on a realistic range of the model parameters. In the case of the single-firm model, no initial guesses are needed.

dimensional model-implied and observed vector of moments respectively (e.g., 6 years of plans, actions of the leader, and actions of the follower respectively), and N the number of unique firms in the dataset. The standard errors reported in Table 2 are computed using the bootstrapping method. That is, we repeatedly take $s = 1, \dots S$ samples (of the same size as the original sample), with replacement, from our original sample. For each sample s, we then estimate $\hat{b}_s$ so as to solve (62). The standard errors of $\hat{b}$ are then computed as $se(\hat{b}) = \sqrt{\frac{1}{S-1}\sum_{s=1}^{S}(\hat{b}_s - \bar{b})^2}$, with $\bar{b}$ the mean of the estimates across the sample S. Table 2 reports the estimate $\hat{b}$ as well as the t-statistic $\hat{b}/se(\hat{b})$. At each iteration, the initial guess $b_0 = \{\lambda_0^p, \lambda_o^{np}, \gamma_0^p, \gamma_0^{np}\}$ which we feed into the minimization algorithm is drawn from a uniform distribution within an admissible range $\{\lambda_0^p, \lambda_0^{np}\} > 0$ and $\{\gamma^p > \gamma^{np} > 0\}$.

D.1 Accounting for time-varying beliefs

In Figure 12, we show that the standard deviation of the beliefs distribution shrinks moderately over time (i.e., the disagreement about the impact of the regulation de- creases, especially within the group of followers), suggesting that firms are learning something about regulation risk. However, given the small dataset we work with, it is difficult to disentangle the anticipated dynamics of future signals (i.e. the anticipated belief distribution) from the actual arrival of new information. We therefore choose to assume that the firm anticipates no learning about the levy up until the terminal date, which is the same as assuming that observed changes in beliefs of the leader and the follower are the result of unanticipated shocks.

To assess the extent to which accounting for future anticipated belief revisions might affect our inferences, we solve a simple single-firm version of the model with this feature.

More specifically, in the current model, in each reporting year $t = 2011, \dots 2016$, we input a belief over the levy λ_t which is extracted from the distribution of reported beliefs $\hat{\Lambda}_t$.[8] The predicted abatement choice x_t^* is then determined assuming the conditional distribution of future beliefs $\{\lambda_s | t = \lambda_t\}_{t<s<T}$. The alternative is to assume that in each reporting year $t = 2011, \dots 2016$, the firm expects to receive a stream of normally distributed signals about the levy $\{z_s\}_{t<s<T}$ with conditional distribution $z_s | t \sim \mathcal{N}(\bar{\lambda}_t, \sigma_t)$. This amounts to assuming that the firm's belief about the future signals is centered around the current expected levy. In turn, the conditional distribution of future beliefs will have a constant mean and decreasing variance $\lambda_s | t \sim \mathcal{N}(\bar{\lambda}_t, \frac{1}{s+1-t}\sigma_t)$.[9]

[8]Those differ across the leader and the follower, but we suppress the notation here as we consider the case of the single-firm model.

[9]Note that at time $s = t+1$, the posterior belief once observing the signal z_{t+1} is

$$\begin{aligned} \sigma_{t+1} &= \frac{1}{(\sigma_t)^{-1} + (\sigma_{t+1}^z)^{-1}} \\ \bar{\lambda}_{t+1} &= \frac{\sigma_t}{\sigma_{t+1}}\bar{\lambda}_t + \frac{\sigma_{t+1}^z}{\sigma_{t+1}} z_{t+1} \end{aligned} \tag{63}$$

Figure 20
Predicted Abatement Rates

The plots shows, in blue, the predicted rate x_t^* as a function of uncertainty σ_t when the anticipated signal $z_s|t \sim \mathcal{N}(\bar{\lambda}_t, \sigma_t)$ (right plot). In red, the predicted rate under the assumption of no anticipated signals, whereas in black the original rate where $\sigma_t = 0$. Other model parameters are $\eta = 0.0037, \beta = 0.93, \phi = 8, \omega = 0.038$, $\bar{\lambda}_t = 80$, $t = 0$, and $T = 8y$ (left plot) and $T = 3y$ (right plot) respectively.

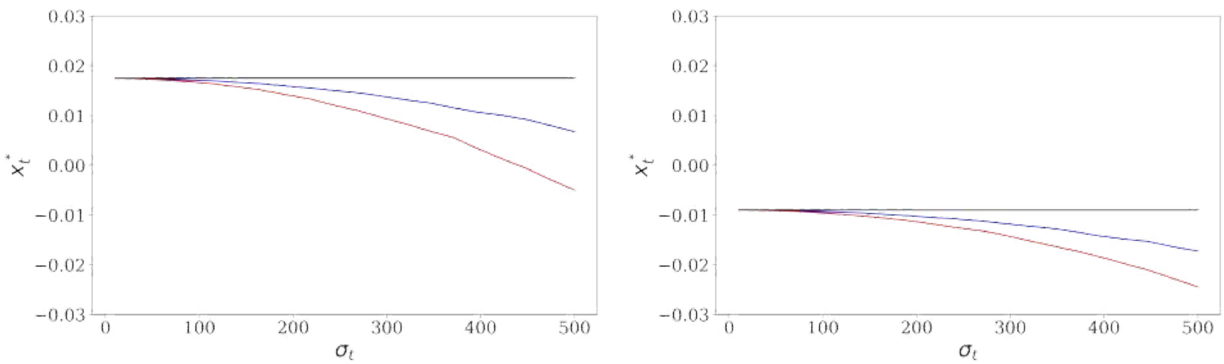

Using the model parameters calibrated on the average firm with plans in the dataset (Table 2, panel A), and setting the expected levy equal to the estimated level $\bar{\lambda}_t = 80$ \$/tCO$_2$e, we simulate in Figure 20 the predicted current abatement rate x_t^* for different values of the uncertainty σ_t and two values of the terminal date $T - t = 3yrs$ (left plot) and $T - t = 8yrs$ (right plot). The blue line is the predicted rate x_t^* when the anticipated distribution of future signals is $z_s|t \sim \mathcal{N}(\bar{\lambda}_t, \sigma_t)$. The red line is the predicted rate from our model, i.e. assuming no anticipated future signals, whereas the black line is the original rate which did not take into account of any uncertainty about the terminal levy (i.e. $\sigma_t = 0$ hence $\lambda \equiv \bar{\lambda}_t$).

With respect to the type of learning considered, our model seems to consistently underestimate the abatement action across all values of uncertainty, which could in turn result in an overestimation of the expected levy (i.e. the estimated levy could be lower than 80 \$/tCO$_2$e). However, the plots show that the errors become smaller as the uncertainty of the belief σ_t decreases. Importantly, since the variance of the reported beliefs $\hat{\Lambda}_t$ decreases over the time period $t = 2011 \ldots 2016$, we can expect that the magnitude of the bias becomes smaller around and after the Paris agreement announcement.

It is important to stress that this is only one possible specification of the conditional distribution of future signals $z_s|t$, and that other choices could be made. To the extent that the Paris regulatory announcement is not anticipated by the firm, we believe that our assumption is reasonable as it does not alter the main model predictions, but it does reduce the number of free parameters and make the two-firm model computationally easier to calibrate.

From which follows that $\mathbb{E}[\bar{\lambda}_{t+1}|t] = \bar{\lambda}_t$ since $\mathbb{E}[z_{t+1}|t] = \bar{\lambda}_t$ and $\sigma_{t+1} = \frac{1}{2}\sigma_t$ since $\sigma^z_{t+1}|t = \sigma_t$. Iterating (63) for $s = t + 2 \ldots T$ one gets that $\lambda_s|t \sim \mathcal{N}(\bar{\lambda}_t, \frac{1}{s+1-t}\sigma_t)$.

About ECGI

ECGI is an international non-profit membership organisation that provides a platform for debate and dialogue among academics, policymakers, and business leaders, with a focus on major corporate governance, ESG, and stewardship issues.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.