

Bank of England

Understanding climate-related disclosures of UK financial institutions

Staff Working Paper No. 1,017

March 2023

Jonathan Acosta-Smith, Benjamin Guin, Mauricio Salgado-Moreno and Quynh-Anh Vo

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.



Bank of England

Staff Working Paper No. 1,017

Understanding climate-related disclosures of UK financial institutions

Jonathan Acosta-Smith,⁽¹⁾ Benjamin Guin,⁽²⁾ Mauricio Salgado-Moreno⁽³⁾ and Quynh-Anh Vo⁽⁴⁾

Abstract

Climate-related disclosures reduce information asymmetries between firms and investors and help transition to a net zero economy. However, disclosure practices might differ across firms. We explore the determinants of firm disclosures by creating a unique, firm-level panel data set on climate-related disclosures of UK financial institutions. To that end, we apply Natural Language Processing techniques with Machine Learning classifiers on unique textual data which we hand-collected from their published reports. We document differences in disclosure levels across financial institutions with different sizes and over time. We show that climate-related policy communications in the form of regulatory guidance on future mandatory disclosures is associated with a catch-up by firms previously disclosing less.

Key words: Climate-related disclosures, market discipline, Task Force on Climate-Related Financial Disclosures (TCFD) and Natural Language Processing (NLP).

JEL classification: G2, C4, C8.

(1) Bank of England. Email: jonathan.smith@bankofengland.co.uk

(2) Bank of England. Email: benjamin.guin@bankofengland.co.uk

(3) Bank of England. Email: mauricio.salgadomoreno@bankofengland.co.uk

(4) Bank of England. Email: quynh-anh.vo@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful for comments from an anonymous referee for the Bank of England Staff Working Paper series, Marcus Buckmann, Angela Gallo, Linda French as well as seminar participants at the Bank of England, the Prudential Regulation Research Consortium Seminar Series (BoE PRR), the Columbia University/Bank Policy Institute 2022 Research Conference, the UK-China Workshop on Energy and Climate Change Economics, the IFABS 2022 conference in Naples and the 4th conference on 'Non-traditional Data, Machine Learning, and Natural Language Processing in Macroeconomics' in Stockholm. We are also grateful to Jiaqi Shi and Zahid Amadjarif whose preliminary work as part of the Bank of England Research Support Programme helped start us off on this project.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH

Email: enquiries@bankofengland.co.uk

©2023 Bank of England

ISSN 1749-9135 (on-line)

1 Introduction

The scale of investments required to support the transition of the global economy to a net-zero future is enormous. Relying on public finance alone is not sufficient to achieve that objective and private finance will have a critical role to play. It is therefore necessary to, as Mark Carney – the former UN special envoy for climate action and finance – noted, lay “the foundations for a system in which every financial decision takes climate change into account”¹. To do so, markets require information to operate effectively and comprehensive climate-related disclosures are thus becoming increasingly important.

Acknowledging a growing demand to provide markets with decision-useful, climate-related information, many authorities around the world are considering whether to make climate-related reporting mandatory. In the course of establishing a path towards that goal, understanding the extent of existing disclosures as well as factors that would affect firms’ incentives to disclose information is essential. This paper aims to shed light on these issues by examining voluntary climate-related disclosures by banks and insurers, hereafter referred to as firms for brevity, in the UK. The paper also examines the determinants of these disclosures. Specifically, it analyses the degree to which climate-related regulatory publications by one of the main UK financial regulators, i.e. the PRA, affects firms’ disclosure decisions.

This paper focuses on climate-related disclosures of financial sectors, namely banks and insurers, which is relevant for several reasons. First, they play a key role ensuring an orderly transition to a net-zero economy by channeling funds to investment projects as financial intermediaries. Second, the financial sector is a heavily regulated industry which allows us to examine the role of climate-related regulatory publications for firms’ disclosure decisions. Examining disclosures practices in the United Kingdom is appealing for several reasons. Indeed, looking at one country, we can hold other country-wide factors like government interventions constant. We consider it a neat way of controlling for a wide set of factors which may differ across countries. Moreover, in the time period considered, there was no mandatory disclosure regime affecting the firms in our sample. It allows us to examine the level and determinants of voluntary firm disclosures.

We rely on firm’s reporting that is in line with the recommendations of the Task Force on Climate-Related Financial Disclosures (TCFD). These recommendations, published in 2017, are organised into four themes: governance; strategy; risk management; as well as metrics and targets. Beneath these themes sit 11 individual recommendations that provide more granular detail on the information to be disclosed. TCFD recommendations are therefore a comprehensive disclosure standard that goes beyond disclosure of green house gas (GHG) emissions included in previous voluntary disclosure initiatives such as the Carbon Disclosure Project (CDP).² They constitute a basis for developing mandatory disclosure standards in many countries. This widely-accepted standard allow us to compare disclosures across firms and over time.

To determine the degree to which UK banks and insurers provide TCFD-aligned information, we use Natural Language Processing (NLP) techniques and Machine Learning (ML) classifiers to scan a large set of publicly available corporate reports published during the period from 2016 to 2020 by a sample composed of 35 bank and 34 insurer groups in the UK.

¹See “Mark Carney – Securing a Net Zero Future – IMF Finance & Development”.

²In 2018, the CDP redesigned its climate change questionnaire to align with TCFD recommendations.

Our contribution is threefold. First, the paper contributes to a strand of the literature examining how firms respond to regulations. While there is a growing literature examining the effects of changes in capital requirements (Gropp et al. 2019, De Jonghe et al. 2020, Arnould et al. 2020) or capital buffers (Jiménez et al. 2017, Basten 2020), research on climate-related regulations is nascent. Existing research examines the effects of mandatory disclosure regimes on emissions (Jouvenot & Krueger 2019) and stock returns (Bolton & Kacperczyk 2021). We contribute to this literature by analysing the role of regulatory actions for climate-related disclosure practices. We study a salient regulatory publication which sets out clear supervisory expectations on firms’ management of climate-related risks and their disclosures. This so-called “Dear CEO letter” also provides hints to a future mandatory disclosures regime. We document evidence of a significant effect of this action on firms’ decisions to disclose climate-related information. Specifically, we show that those firms which previously disclosed less catch up in terms of their disclosures.

Second, we demonstrate a novel application of Natural Language Processing (NLP) techniques. While such techniques are often used in the impact analysis of monetary policy announcements, this paper is amongst the first to illustrate their application in retrieving information from firms’ public reports. Bingler et al. (2022) use a deep neural network model developed in Webersinke et al. (2021) to identify climate-related financial information from TCFD reports. The methodology used in these two recent studies try to overcome the challenges identified in previous studies (see Varini et al. (2020)) suggesting that using a bag-of-words or similar approach in text classification tasks do not work as well in “niche language” contexts, such as in climate-related texts. We show that a much simpler multi-step approach, relying first on climate-keywords and stop-phrases and then two rounds of supervised machine learning models, is able to cope with the complex multi-labeling problem of identifying information in line with TCFD recommendations.³

Third, by doing so, we create a novel panel dataset on climate-related disclosures.⁴ Thereby, this paper is among the first formal analyses of the new disclosure standard introduced by the TCFD. This standard includes more comprehensive information that goes beyond carbon emission. Making use of such a disclosure standard that is uniform across firms allows us to compare climate-related disclosures both across firms and over time. We go beyond some existing analyses on TCFD disclosures by exploiting the comparability of firm disclosures in our dataset to examine the economic incentives of disclosure. Precisely, we investigate how climate-related financial disclosures relate to different financial characteristics of firms such as size, profitability, leverage and ownership structure. We find that disclosure level is positively correlated with firms’ size. We don’t find evidence of significant relationships between other firm characteristics and disclosure level. We observe a clear increasing trend in disclosures over time across all TCFD themes and across both

³There are further differences. While we focus on climate-related disclosures of the UK financial sectors, Bingler et al. investigate a wide range of sectors (such as Financial, Materials, Industrials, Utilities, Consumer staples, Real estate, Energy, IT, Transportation, Health Care, etc) across several regions (such as Europe, Asia, North America, etc). Another difference is that while we train 11 independent models for each of the 11 individual TCFD recommendations, Bingler et al. classify only by the TCFD theme (i.e., governance, strategy, risk management, or metrics and targets), which simplifies the complex classification task at hand.

⁴More precisely, our dataset includes binary variables that indicate, for each firm and each TCFD individual recommendation in each year, if that firm discloses any information in line with that recommendation in that year. We also construct, for each firm in each year, a variable that counts the number of TCFD recommendations that a firm discloses information in line with.

banking and insurance sectors. The increase is the most significant in 2019. However, there is still a significant number of firms that do not disclose any information.⁵

The rest of the paper is organised as follows. Section 2 describes our NLP approach to construct the TCFD disclosure dataset and the main disclosure measures we build. Then in Section 3 we examine how climate-related disclosures of UK financial institutions evolve over time and what characteristics can affect their decisions to disclose information. Section 4 is devoted to the analysis of the role of regulatory announcements for disclosure decisions. Section 5 contains additional robustness analyses and, finally, Section 6 concludes.

2 Climate-related disclosures

2.1 Background on TCFD disclosure standard

In June 2017, the Task Force on Climate-related Financial Disclosures (TCFD) - established by the Financial Stability Board (FSB) - announced a set of recommendations for voluntary, consistent climate-related financial risk disclosures for use by companies, banks and investors in providing information to stakeholders. The TCFD disclosure standard is structured around four thematic areas that represent core elements of how firms operate: governance; strategy; risk management; as well as metrics and targets. Beneath these themes sit 11 individual recommendations that provide more granular detail on the information to be disclosed.

Within the governance theme, the TCFD recommends to disclose the firms' governance around climate-related risks and opportunities, which includes the board's oversight of those risks and opportunities as well as the management's role in assessing them. Three recommendations of the strategy theme link to the actual and potential impacts of climate-related risks and opportunities on firms' businesses. They consist of information on firms' exposure to those risks and opportunities, their impact on firms' businesses and the resilience of firms' business strategy with respect to them. In terms of risk management, the TCFD considers useful to disclose how firms identify, assess and manage climate-related risks. The three recommended disclosures of this theme therefore focus on describing processes that firms use to identify those risks, manage them and how those processes are integrated into the overall risk management. Finally, for the metrics and targets theme, recommended information to be disclosed encompasses metrics used by firms to assess climate-related risks and opportunities, firms' green house gas (GHG) emissions across three scopes, and firms' targets to manage those risks.

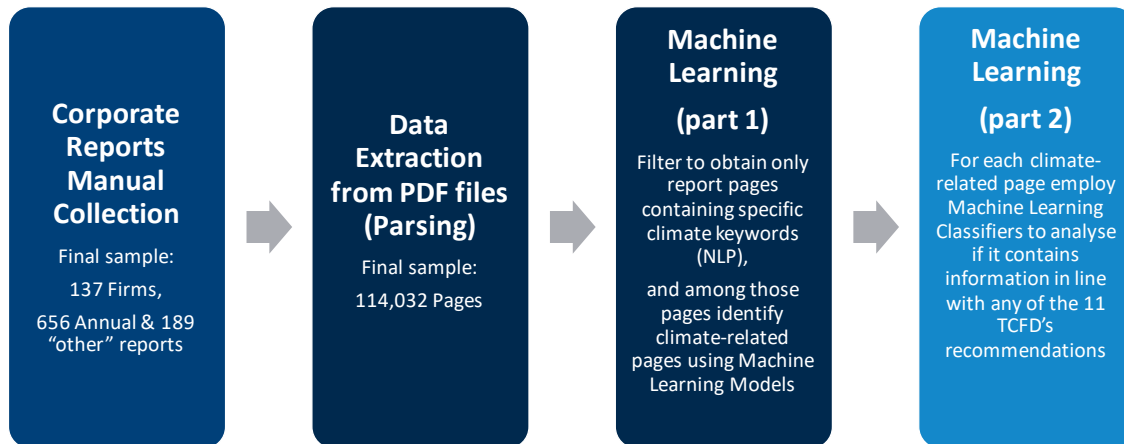
TCFD recommendations are the first comprehensive climate-related disclosure standard designed to solicit decision-useful, forward-looking information that can be included in mainstream financial filings. With more than three thousand global organisations from 95 countries having declared support for those recommendations as of June 2022 together with their consistency and comparability, they have become a good laboratory to study firms' incentives to disclose climate-related information.

⁵Despite methodical differences, [Bingler et al.](#) also find a positive and statistically significant increase in information disclosed in line with the TCFD recommendations between 2020 and 2017 when the final set of recommendations was published. Additionally, they also find substantial differences in climate-related disclosures across regions and different sectors in their sample, with financial firms disclosing more than almost any other sector and only behind sectors that are traditionally associated with a high carbon footprint, i.e, energy, utilities and transportation.

2.2 A novel data set on firms' disclosures

Our approach to create a novel data set with firm-level disclosures consists of four consecutive steps starting with the collection of corporate reports. Figure 1 summarises those steps.⁶

Figure 1: Machine Learning Pipeline



In the first step, we manually collect annual reports of the biggest banks, building societies and insurance companies between 2016 and 2020 in the UK. In addition, when applicable, we include in the sample the international parent (or group) company of UK firms, leaving us with 137 firms in total. These firms include many subsidiaries and firms belonging to the same group. Thus, if we consolidate them at the parent/group level, we are left with 35 banking and 34 insurance companies. Moreover, we also collect all types of publications, within our sample period, where firms might have disclosed climate-related information. For example, we gather, among others, Environment, Social and Governance (ESG) reports, Corporate Social Responsibility (CSR) report, Sustainability and TCFD reports. Our final collection includes over 820 corporate reports over five years.

In the second step, we parse over our collected reports using a single page as the unit of analysis to obtain our corpus encompassing over 110,000 pages. It is worth emphasising that our methodology is able to handle both text-formatted as well as image-formatted PDF files. To reduce the large dimensionality of our corpus, we first identify potential climate-related pages using NLP techniques such as climate keywords and stop-phrases. Using these simple techniques we are able to discard 81% of the initial observations leaving us with approximately 21,000 pages potentially containing climate-related information.

In step three, we continue reducing the size of our truncated sample by identifying truly climate-related pages using a first round of machine learning classifiers. Concretely, we rely on a supervised learning approach and a random forest model with keyword-only feature representation. We identify over 8,000 climate-related pages, which represent approximately 7% of our initial corpus.

Finally, to handle the complex multi-labeling problem at hand,⁷ we train 11 independent classifiers to

⁶Appendix A.1 provides a detailed description of our Machine Learning methodology.

⁷This is a multi-label problem since within each page there might be information related to any subset of the 11 TCFD recommendations.

identify the information disclosed on each corporate report’s climate-related page that is closely in line with each of the individual recommendations published by the TCFD. More specifically, we rely on a nested cross-validation approach to train our ML-models, and to find the best hyperparameters that optimize our performance measure AUC, i.e., the “Area under the Receiver Operating Characteristic Curve.”

We confirm the high performance of our fine-tuned ML-models in two complementary ways. First, we rely on the results from our ML-models on our test sample. Table A.2 in the appendix presents these results, which show that our models’ performance is indeed quite high giving an average AUC score of 0.88.⁸ Second, we compare the predictions of our ML classifiers against a manual classification exercise that the Bank of England’s Climate Hub conducted for a handful of the largest UK banks. This comparison yields similar results across both manual and ML methodologies in 85% of cases, while there is a similar number of false positive than false negative results under the remaining 15% of observations compared.

2.3 Disclosure measures

The fine-tuned ML-models described in the previous section allow us to measure the extent of firms’ voluntary disclosures in three different ways. First, we can count the number of pages that a firm uses to disclose information in line with TCFD recommendations. We can also construct, for each triple firm-recommendation-year, a binary indicator that takes value of 1 if that firm discloses information in line with that recommendation in that year and 0 otherwise. Using these indicators, we build two other measures of firms’ disclosure. The first measure, which we dub “Sum of Recommendations”, counts the number of TCFD recommendations that a firm, in each year, discloses in line with. This measure can take the value from 0 to 11. The second measure evaluates the extent to which a firm satisfies each of the four TCFD themes. This is constructed as the average value of the binary indicators that correspond to each recommendation within that theme. For example, for the *strategy* theme, which has three recommendations, we would measure disclosure for a firm in 2018 as the average across the three binary indicators corresponding to the three strategy recommendations in that year.

Table 1 reports summary statistics of our different disclosure measures for all firms, and for each of our two financial sectors, banks and insurers. For ease of presentation, we display only the average value across all years of each measure.⁹ We observe that while some firms still do not disclose any information at all, it would seem on average that we are approaching around 50% of recommendations being satisfied across the four TCFD-themes (Panel A). On average, firms disclosed information in line with about 5 recommendations and they published around 35 pages with climate-related information (Panel B).

Comparing banks and insurers, we see that firms in both financial sectors display very similarly disclosure levels across our two overall disclosure measures (“Sum of Recommendations” and “Pages with Recommendations”) (Panel B). Similarly, we observe very similar disclosure levels across the four TCFD-themes (Panel A).

⁸A model whose predictions are 100% wrong has an AUC of 0 while a model whose predictions are 100% correct has an AUC of 1.

⁹See below for year by year split.

Table 1: **Summary Statistics of Climate-related Disclosures**

	All Firms		Banks		Insurers	
	Mean/SD	Min/Max	Mean/SD	Min/Max	Mean/SD	Min/Max
A: Disclosure by TCFD-Theme						
Governance Avg.	0.424 (0.477)	0 1	0.436 (0.469)	0 1	0.411 (0.487)	0 1
Strategy Avg.	0.472 (0.435)	0 1	0.472 (0.434)	0 1	0.472 (0.436)	0 1
Risk Management Avg.	0.435 (0.421)	0 1	0.457 (0.418)	0 1	0.411 (0.425)	0 1
Metrics and Targets Avg.	0.589 (0.446)	0 1	0.617 (0.436)	0 1	0.560 (0.456)	0 1
B: Overall disclosure						
Sum of Recommendations	5.336 (4.259)	0 11	5.509 (4.158)	0 11	5.153 (4.370)	0 11
Pages with Recommendations	34.869 (60.289)	0 397	33.087 (60.167)	0 397	36.761 (60.546)	0 312

Note: Standard deviations (SD) in parenthesis. All statistics are taken at the group-level. The variable “Sum of Recommendations” is the sum of the 11 binary indicators for each individual TCFD recommendation. The variable “Pages with Recommendations” gives the total number of pages that contain information in line with a TCFD recommendation consolidated at the group level. In our analysis below the variable “Pages with Recommendations” is used with an inverse hyperbolic sine (IHS) transformation.

Finally, note that the firms with the maximal number of Pages with Recommendations in one year are several standard deviations above the mean in each sector. As we show below in more detail, these relative large values come from large international firms in the last year of our sample, i.e., 2020.


3 Evolution and determinants of climate-related disclosures

This section analyses the evolution over time of climate-related disclosures of UK financial institutions. We also examine which firm’s financial characteristics are important factors that affect disclosure.

3.1 Descriptive insights

There are three main insights that can be drawn from our dataset. First, when examining the consolidation level at which firms disclose climate-related information, we find that firms mostly provide information in reports published by the highest group consolidated entity.


Figure 2: TCFD Recommendations Disclosed



Note: Average number of TCFD recommendations disclosed per year by groups and subsidiaries.

Second, from the number of climate-related pages identified via ML and the number of disclosed TCFD recommendations (see Figure 3), we observe that climate-related disclosures by UK banks and insurers have increased over time. There are significant jumps of disclosure in 2019 and 2020. A simple t-test confirms that the increase across both measures of disclosure in 2019 and 2020 is statistically significant.

Figure 3: Climate-related Pages and TCFD Recommendations

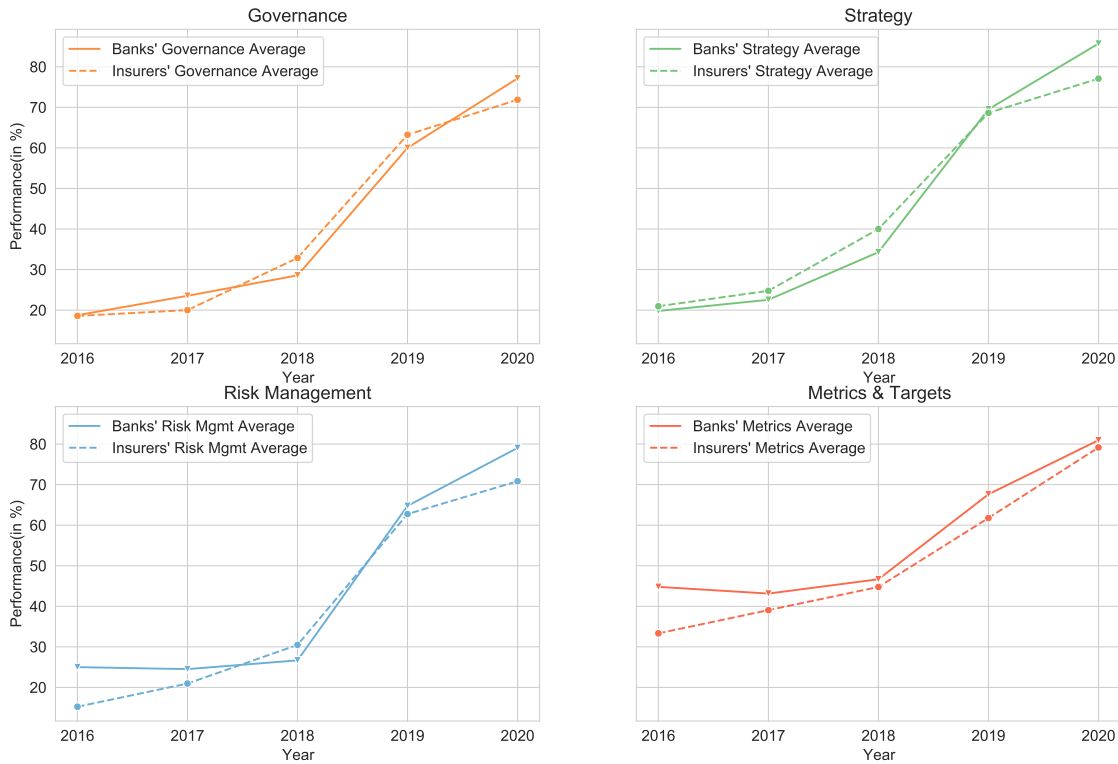


Note: Number of climate-related pages and number of TCFD recommendations disclosed per year.

Figure 4 which decomposes the TCFD disclosure by TCFD themes and sectors. We see that there is an increase in disclosure over time which is uniform across all themes and across both sectors. For example, the upper left panel shows that in 2017, 27% of banks in our sample disclosed information related

to the governance theme, while less than 20% of insurers disclosed information in line with the same theme. However, there is still a significant number of firms that do not disclose any information in line with TCFD recommendations.

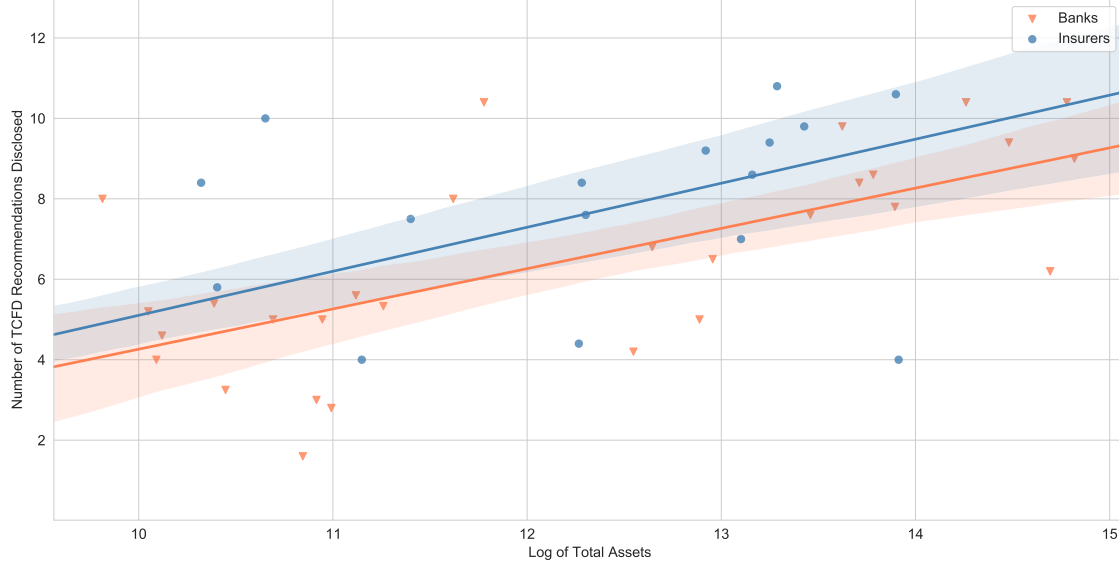
Figure 4: TCFD Disclosure Themes at the Group Level



Note: Group-Level Performance (in %) by sector of disclosed TCFD-theme recommendations per year.

Third, taking into consideration firm characteristics, firm size seems to particularly matter. Figure 5 illustrates the strong positive correlation between the level of disclosure and the log of firms' total assets. This might be because bigger firms are exposed to greater societal and therefore market pressure to disclose information. Hence, they would disclose more information for this reason. Also, practically, it is likely that larger firms have a greater ability to produce such information.

Figure 5: Firm’s Size and TCFD Disclosures by Sector



Note: (Log of) Total Assets (average over time) vs average number of TCFD recommendations disclosed per group by sector.

3.2 Regression analysis

This subsection examines more formally which firms’ characteristics affect their decision to disclose climate-related information. To do so, we collect additional data on firm financial characteristics from CapitalIQ. Since our dataset includes both banks and insurers, we consider those financial characteristics that are common across the two. They are the size of an institution (measured using total assets), the profitability of that institution (measured using the return on equity (ROE)) and its leverage (measured by the ratio of total debt to total assets). Table 2 reports the key summary statistics across these variables. There is substantial variation across all the variables. For example, there are both highly profitable firms as well as loss-making ones.

Table 2: Descriptive Statistics: Firm’s Financial Characteristics

	All Firms		Banks		Insurers	
	Mean/SD	Min/Max	Mean/SD	Min/Max	Mean/SD	Min/Max
Total Assets (billions)	488.25 (706.02)	0.85 3384.76	684.51 (864.26)	12.41 3384.76	247.56 (302.92)	0.85 1294.59
ROE	7.29 (11.72)	-36.07 56.27	5.07 (7.31)	-36.07 25.42	10.02 (15.11)	-32.41 56.27
Leverage	0.18 (0.22)	0.00 0.86	0.28 (0.25)	0.00 0.86	0.05 (0.06)	0.00 0.35
Observations	236		130		106	

Note: Standard deviations (SD) in parenthesis.

Hypotheses development There exist two main theories¹⁰ that explain disclosure decisions: sociopolitical theories of disclosure (e.g. Gray et al. (1995) and Hahn & Luelfs (2014)) and economic theories of (voluntary) disclosure (e.g. Verrecchia (1983) and Clarkson et al. (2008)). The former theories view disclosures as firms’ response to social, political and stakeholder pressures. The two main anchors of this group of theories were stakeholder and legitimacy theory, which differ mainly in terms of their focus on actors. According to stakeholder theory, disclosure can be explained as a response to demand for information from stakeholders who can be investors, contractors etc. Legitimacy theory suggests that environmental disclosures can be used to maintain the implicit social contract between a company and society.

Economics-based theories of disclosures indicate that firms voluntarily disclose information to interested actors based on an evaluation of costs and benefits. Disclosed information can help reducing the problem of asymmetric information in the financial market. Disclosures can also lead to different costs for firms. The direct costs of disclosures include the preparation, certification and dissemination of information. Those fixed costs from information production and dissemination can make certain disclosures particularly burdensome for smaller firms. Disclosures have indirect costs because information provided to capital-market participants can also be used by other parties such as competitors, regulators.

Those above-described theories therefore suggest that characteristics that are relevant for disclosure decisions include size, profitability, leverage, ownership and governance.

- Bigger firms are usually more visible and so pressured to legitimise their actions. Bigger firms are also better equipped to bear the fixed costs of disclosures as well as proprietary costs.
- More profitable firms can be expected to have more financial resources to implement required system to produce and disseminate information.
- Leverage can be important since it affects the relevance of asymmetric information problem for firms. For example, the corporate finance theory suggests that in the presence of moral hazard (e.g. incentives for entrepreneurs to exert effort) or of asymmetric information problem in terms of semi-verifiable cash flows, the optimal financing contract takes the forms of debt contract. In that spirit, higher leverage will reduce the relevance of asymmetric information problems for firms.
- Institutional investors usually have higher demand for information and may require disclosures of some information as condition for investment. Therefore, institutional ownership can be a characteristic relevant for firms’ disclosure decisions.

Analysis To examine robustly how climate-related financial disclosures relate to the characteristics of firms, we estimate a simple multivariate regression model in which we relate group-level disclosures to a set of explanatory variables.

$$\text{Disclosure}_{i,c,t} = \alpha + \beta_1 \text{year}_t + \beta_2 \text{PRA}_{i,c} + \beta_3 \text{Sector}_{i,c} + \beta_4 \text{RBBM}_{i,c} + \beta_5 \text{NLBM}_{i,c} + \gamma' \mathbf{X}_{i,c,t} + u_{i,c,t} \quad (1)$$

where $\text{Disclosure}_{i,c,t}$ is our dependent variable wherein we use the six different variables shown in Table 1,

¹⁰See Hahn et al. (2015) for a complete literature review.

namely: the four TCFD recommendation themes; the sum of all recommendations; and the total number of pages that contain TCFD recommendations. $Year_t$ is a linear time trend expressed in years. $PRA_{i,c}$ is a dummy variable equal to 1 if firm i is PRA-regulated, 0 if not. $Sector_{i,c}$ is a dummy equal to 1 if it is a bank, but 0 if it is an insurer. $RBBM_{i,c}$ and $NLBM_{i,c}$ are dummy variables equal to 1 if a firm’s business model is retail banking or non-life insurance respectively, 0 otherwise. $u_{i,c,t}$ is an idiosyncratic error term.

$\mathbf{X}_{i,c,t}$ is a vector of firm-level financial characteristics that can be relevant for firms’ disclosure decisions. Based on the above-described theories, we include in the baseline regression size - proxied by total assets, profitability measured by ROE and leverage measured by the ratio of total debt to total assets as shown in Table 2. We also collect data on institutional ownership of firms in our sample. Due to data availability, including ownership characteristic in the regression reduces the number of observations significantly. Therefore we don’t include it in the baseline regression but in a robustness analysis.

Table 3 shows the results from the baseline regression. We report results with leverage as an explanatory variable. Column (1) shows results for the sum of recommendations. Columns (2)-(5) show results when we consider each of the TCFD themes individually.¹¹ Lastly, column (6) shows results when considering the total number of pages that mention the recommendations.

Table 3: **Climate-related Disclosures and Firm Characteristics**

	Sum of Rec.	Governance Avg.	Strategy Avg.	Risk Mgmt. Avg.	Metrics and Targets Avg.	Pages with Rec. (IHS)
<i>Time Trend</i>						
Year	1.499***	0.150***	0.164***	0.147***	0.089***	0.654***
<i>Financial Variables</i>						
Total Assets (logs)	1.031***	0.109***	0.090***	0.109***	0.073***	0.505***
ROE	0.018	0.004**	0.000	0.002	0.001	0.007
Leverage	0.114	-0.038	-0.077	0.134	0.007	-0.310
<i>Indicator Variables</i>						
Sector	-1.131*	-0.162	-0.138*	-0.176**	0.046	-0.514
PRA-Regulated	-1.853***	-0.039	-0.238***	-0.168**	-0.186*	-0.887**
Bank Bus. Model	1.526*	0.098	0.151**	0.188***	0.105	0.499
Non-Life Bus. Model	1.674**	0.089	0.128	0.117	0.254**	0.688*
Constant	-3029.278***	-303.823***	-330.664***	-296.929***	-179.618***	-1320.992***
Observations	236	236	236	236	236	236
Adjusted R^2	0.572	0.399	0.566	0.549	0.244	0.569

Note: This table shows the results of the regression model in Eq. (1). The variable “Pages with Rec.” is used with an inverse hyperbolic sine (IHS) transformation. Stars denote p-values * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Looking across the specifications, we can see there is a significant increasing trend in disclosures over time. This compares well to our descriptive charts above in which we see a consistent increase in disclosures across all themes and sectors over time. In fact, on average, one year is associated with an average increase of 1.5 recommendations (column (1)) and an increase of 0.7 pages (column (6)). Similarly, the coefficient on size is positive and statistically significant across all specifications. In fact, a 1 percent increase in total assets

¹¹The dependent variables are defined as the average for firm X in year Y across the recommendations of the TCFD-theme (e.g., for the *strategy* theme, it would be the average across the three strategy recommendations for firm X in year Y).

is associated with an average increase of 1.0 recommendations (column (1)) and about 0.5 pages (column (6)). This fits with the prior that larger firms are more exposed to societal pressure to disclose information. Whereas, it is not clear that this societal pressure consistently hits a particular sector or business model across the different indicators. These insights are valid both when leverage is and is not controlled for. Furthermore, leverage itself does not seem to matter for the disclosure decision.

4 The role of regulatory announcements for climate disclosures

In this section of the paper, we examine the role of regulatory announcements on firms' decision to disclose in line with TCFD. Specifically, we examine whether firms that had previously not disclosed as much catch up in terms of their disclosures following regulatory announcements relative to firms that had previously demonstrated high levels of disclosures. To that end, we employ a Difference-in-Differences design which allows us to compare disclosure levels of these two group over time at different stages of regulatory announcements.

4.1 Background of regulatory announcements

In the UK, the Prudential Regulation Authority (PRA) is responsible for the prudential regulation and supervision of around 1,500 banks, building societies, credit unions, insurers and major investment firms. It has a general objective to promote the safety and soundness of the firms it regulates. Its rules require PRA-regulated financial firms to maintain sufficient capital and have adequate risk controls in place.¹² By contrast, the Financial Conduct Authority (FCA) regulates the conduct of around 51,000 firms. Moreover, it is the prudential supervisor for 49,000 firms and it sets specific standards for around 18,000 firms.¹³ As our sample does not include FCA-regulated firms, we focus on the sequencing of PRA regulatory publications for the purpose of our analyses.

Since 2018, the PRA has issued a set of announcements and publications to facilitate the management of climate-related financial risk. Table 4 outlines the sequencing of the most material announcements.

- In October 2018, the PRA opened a public consultation by publishing a consultation paper (CP23/18) on managing the financial risks from climate change.¹⁴ In this CP, the PRA proposes that firms should develop and maintain an appropriate approach to disclosure of the financial risks from climate change.¹⁵
- In early 2019, after the end of the public consultation period, the PRA published a policy statement

¹²<https://www.bankofengland.co.uk/prudential-regulation>, accessed on 3 November 2021.

¹³<https://www.fca.org.uk/about/the-fca>, accessed on 3 November 2021.

¹⁴<https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/consultation-paper/2018/cp2318.pdf?la=en&hash=8663D2D47A725C395F71FD5688EC48E08>, accessed on 29 June 2022.

¹⁵ Consultation papers are the formal document by which we set out draft proposals and invite comments on them from the public. CPs normally deal with one discrete issue of significance.

(PS11/19) which provided feedback on consultation responses, alongside the publication of a supervisory statement.¹⁶

- In this supervisory statement (SS3/19)¹⁷, the PRA set first expectations for firms to promote sound approaches to managing financial risks from climate change by asking firms to take a strategic approach.¹⁸ In terms of disclosure, the SS stated that the PRA expects firms to engage with wider initiatives on climate-related financial disclosures and to take into account the benefits of disclosures that are comparable across firms. It mentions the TCFD framework as an example for firms to consider when developing their approach to climate-related financial disclosures.
- In July 2020, it issued a “Dear CEO letter” which addressed risk management practices directly with PRA-regulated firms.¹⁹ This recent letter to the CEOs of all PRA-regulated firms builds on the expectations set out in SS3/19, provides observations on good practice, and sets out next steps for implementation. In terms of timelines, the letter made clear that “firms should have fully embedded their approaches to managing climate-related financial risks by the end of 2021”. The letter more explicitly encouraged firms to engage with the TCFD framework and it flagged that the PRA was exploring the appropriateness of making climate-related disclosure reporting mandatory.

Table 4: **Climate-related Regulatory Publications by the PRA**

Date	Authority	Type		Focus	Firms affected
2018 October	PRA	Consultation (CP23/18)	Paper	Managing the financial risks from climate change	PRA-regulated firms
2019 April	PRA	Policy (PS11/19)	Statement	Managing the financial risks from climate change	PRA-regulated firms
2019 April	PRA	Supervisory (SS3/19)	Statement	Managing the financial risks from climate change	PRA-regulated firms
2020 July	PRA	Dear CEO letter		Managing the financial risks from climate change	PRA-regulated firms

4.2 Empirical strategy

In our empirical analyses, we aim to examine whether any of these publications affected firms’ climate-related disclosures. Ideally, we would like to estimate the effect of a policy communication by calculating the difference in the average climate-related disclosure by a firm that has been affected by this announcement (i.e. our *treatment group*) and the average hypothetical, counterfactual disclosure of the same firm if it had

¹⁶<https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/policy-statement/2019/ps1119.pdf?la=en&hash=CD95D958ECDA4C7CF94337DAFD8AD962DE>, accessed on 29 June 2022.

¹⁷<https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2019/ss319>, accessed on 29 June 2022.

¹⁸Supervisory statements set flexible frameworks for firms, incorporating new and existing expectations. They focus on our expectations and are aimed at facilitating firm and supervisory judgement in determining whether they meet those expectations. They do not set absolute requirements – these are contained in rules.

¹⁹<https://www.bankofengland.co.uk/prudential-regulation/letter/2020/managing-the-financial-risks-from-climate-change>, accessed on 29 June 2022.

not been affected (i.e. our *control group*).²⁰

Unfortunately, we cannot directly observe such counterfactual outcomes. In theory, *randomized controlled trials (RCT)* can address this problem.²¹ However, it is easy to see that in case of regulatory announcements we examine, the treatment assignment is not random. Instead, they affect a well-defined group of firms, those being under the scope of either of the regulators. Hence, we employ a Difference-in-Differences design which has been widely used in the applied literature (e.g. Card & Krueger 1994, Brown et al. 2016).²² It allows us to analyse how firms affected by regulatory announcements change their climate-related disclosures relative to firms not directly affected by those announcements.

In the spirit of Bolton & Kacperczyk (2021), we construct the treatment group as the subset of those groups whose level of disclosure was below the average level of disclosure in 2017. Additionally, for a firm being considered part of treatment group, we require that a firm is indeed regulated by the PRA. Vice versa, the control group is formed by those groups whose level of disclosure was at least the average level of disclosure in 2017. The intuition of this approach is that firms already disclosing would not have been affected by the regulatory announcements as they had decided to disclose their risk before. Vice versa, firms in the treatment group had not yet disclosed and regulatory announcements are more relevant for them.

We then compare the level of disclosures across these two groups over time. Our empirical strategy follows Rodnyansky & Darmouni (2017) and Arnould et al. (2020) that use a difference-in-differences setup with multiple treatment dates to examine the effects of prudential policy announcements on bank behaviour.²³ In our main specification, we estimate the following regression model using OLS.

$$\text{Disclosure}_{i,c,t} = \delta' \mathbf{Post}_t + \gamma' \text{Treated}_{i,c} \mathbf{Post}_t + c_{i,c} + \omega' \mathbf{Z}_{c,t} + \gamma' \mathbf{X}_{i,c,t} + u_{i,c,t} \quad (2)$$

where the subscript $i = 1, \dots, n$ indexes the firm, $c = 1, \dots, C$ indexes the industry in which a firm operates, and $t = 2016, \dots, 2020$ indexes the years of our sample period. An interpretation of the model components is given below:

- *Disclosure* stands for the dependent variables that we use. We employ six different ones. We employ the average across four TCFD recommendation themes (*Governance, Strategy, Risk Management, Metrics & targets*) as well as the sum of all recommendations and the total number of pages which include climate disclosures.
- $\mathbf{Post}_t = \{Post_{2018}, Post_{2019}, Post_{2020}\}$ is a set of binary time dummy variables, which become equal

²⁰Or, in more technical terms, the effect of the policy on the outcome variable, our case the disclosure, can be then estimated by comparing the identified counterfactual outcomes to the observed outcomes under the policy intervention (Fougère & Jacquemet 2020).

²¹In such empirical designs, properties would have the same chance of being allocated to the treatment group, i.e. being affected by the policy intervention, or to the control group, i.e. not being affected by it. RCTs ensure that both groups are equivalent for comparison, and that there are not any pre-existing differences. In such a setting, we would not expect any confounding variables, e.g. neither the type, size, nor the location of a firm would influence the outcome and confound the results of the analysis.

²²For example, Gropp et al. (2019) and Arnould et al. (2020) study the effects of prudential policy. Rodnyansky & Darmouni (2017) and Fatouh et al. (2021) examine the effects of monetary policy. Garbarino & Guin (2021) study banks' reaction to climate risks.

²³See, e.g., Fatouh et al. (2021) for an application in the context of quantitative easing.

to one after each regulatory announcement.

- $\text{Treated}_{i,c}\mathbf{Post}_t$ is an interaction term between the time dummies and an indicator of a firm disclosures in 2017 which takes the value of one if disclosure levels were below the cross-sectional average before the final set of TCFD recommendations was published in 2017.
- c_i is the intercept for the i th firm, that controls for unobserved variation at a firm level;
- $\mathbf{Z}_{c,t}$ is the vector of industry x time fixed effects, controlling for unobserved variation constant within a certain industry in a given year.
- $\mathbf{X}_{i,c,t}$ is the vector of time-varying firm-level control variables. Following the literature on bank lending (e.g., [Kashyap & Stein \(2000\)](#)), it includes group size via total assets and return on assets (ROA) as a benchmark for profitability. We also control for firm leverage, measured by equity normalized by total assets.
- $u_{i,c,t}$ is the idiosyncratic error term.

We cluster standard errors at the firm level. Of primary interest is the estimated coefficients of our vector of interaction variables, $\text{Treated}_{i,c}\mathbf{Post}_t$, which measures the change of the average climate-related disclosure by firms affected relatively more by the regulatory announcements, compared to relatively unaffected firms.

4.3 Results

Before showing the more formal regression results, we examine how climate-related disclosures changed over time across our treatment group and control group. [Figure 6](#) shows how our main measure, the sum of all recommendations, has evolved over time. Prior to the first publication in 2018, firms in the treatment group published only on average about two recommendations in line with TCFD. By contrast, firms in the control group published on average close to eight recommendations. Following the first announcements by the PRA in October 2018, firms in the control group slightly increased their average number of disclosures by more than two recommendations. By contrast, firms in the treatment group increased their average number of disclosures even more, by almost four recommendations on average. We observe a similar pattern in the following year 2019 which saw the PRA publishing its Supervisory Statement (SS3/19). In this time period, the average number of disclosures of our control group remained at ten, while firms in our treatment group increased their disclosures on average by about one recommendation to a total of seven.

The sum of all recommendations shown in [Figure 6](#) comes with a potential drawback in that there is a well-defined upper bound. This could be a concern for our Difference-in-Difference analyses if there are many firms in the control group that are close to this bound. Hence, we show the evolution of an alternative measure, the total number of pages with recommendations in line with TCFD. It mitigates this potential concern as there is no upper bound in terms of the number of pages that can be published. [Figure 7](#) suggests that this measure shows a similar pattern with both, firms in the treatment group and control group, disclosing with a parallel time trend prior to 2018. Only after 2018, firms in the treatment group catch up in terms of their

published number of pages with the gap narrowing by about 1 page in the period from 2018 until 2020 to a difference of only 2 pages in 2020.

Figure 6: Sum of Recommendations

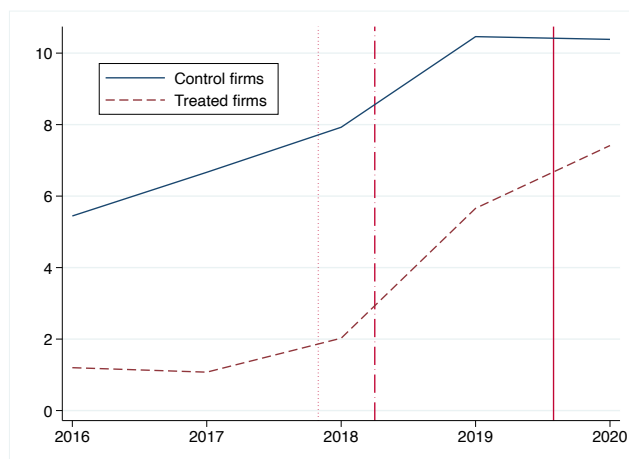
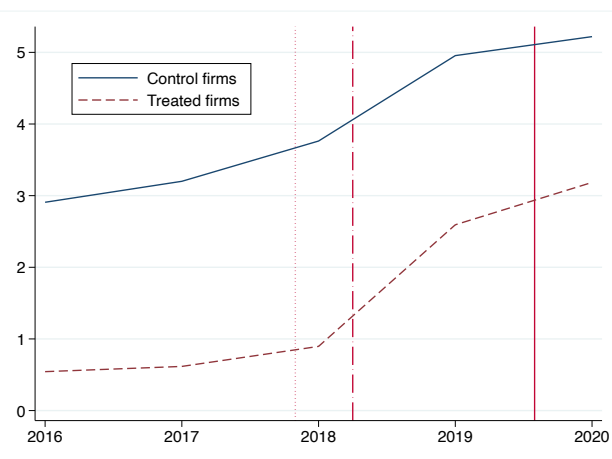


Figure 7: Pages with Recommendations (IHS)



While there is clear descriptive evidence that firms in our treatment group catch up in terms of their disclosures following the regulatory announcement, we would like to understand if these effects are statistically significant. To that end, we estimate the regression model presented in Equation 2. Table 5 summarises these regression results controlling for firm leverage. Column 1 shows the sum of recommendations. In columns 2-5, we break the sum down into individual categories. Column 6 shows the total number of pages with climate-related disclosures.

The first row shows the estimated coefficient of Post (2018), the time dummy that becomes one in the year 2018 (and later) which is the first year of regulatory announcements. The average number of recommendation increased by about 1.3 in that year relative to the pre-intervention period (column (1)). Also the number of published pages increased by about 0.4 pages (column (6)). By contrast, the estimated coefficient on the interaction term Treated x Post (2018) suggests that there is no strong evidence for a differential effect among the firms in our treatment group.

The second row analyses the coefficient of Post (2019), the time dummy that becomes one in the year 2019 (and later) which is the year of the publication of the Supervisory Statement (SS3/19). Again, there is evidence of a strong time trend with average number of recommendations increasing by another 2.5 (column (1)) and the average number of pages increasing by 1.1 (column (6)) in that year. The estimated coefficient on the interaction term Treated x Post (2019) suggests some catching of firms in our treatment group but only in terms of the strategy recommendations. On average, these firms increased the number of disclosed strategy recommendations by about 0.2 more (column (3)) and the total number of pages by about 0.6 more (column (6)) relative to unaffected firms.

The third row analyses the coefficient of Post (2020), the time dummy that becomes one in the year 2020 which is the year of the publication of PRA’s “Dear CEO letter”. This communication is relevant to firms as it contains a clear deadline for the implementation of the supervisory expectations published earlier.

Interestingly, the coefficient of the time dummy is not positive anymore, which suggests that there is no further increase in disclosure levels across all firms. Looking at the interaction term Treated x Post (2020), there is evidence of a significant effect on treated firms’ decisions to disclose climate-related information across all four TCFD recommendation themes. On average, these firms increased the number of disclosed recommendations by another 2.6 relative to unaffected ones (column (1)) with the largest catch-ups in the governance recommendations (column (2)). In total, affected firms increase about another 0.6 pages more than unaffected firms (column (6)).

Table 5: **Difference-in-Differences Results**

	Sum of Rec.	Governance Avg.	Strategy Avg.	Risk Mgmt. Avg.	Metrics and Targets Avg.	Pages with Rec. (IHS)
<i>Time variables</i>						
Post 2018	1.296***	0.103	0.187***	0.039	0.133***	0.390***
Post 2019	2.521***	0.283***	0.255***	0.281***	0.080*	1.132***
Post 2020	-0.782**	-0.136*	-0.091**	-0.022	-0.046	0.046
<i>Treatment variables</i>						
Treated x Post 2018	-0.497	-0.009	-0.080	-0.040	-0.097	-0.155
Treated x Post 2019	1.118	0.034	0.217**	0.147	0.177	0.646*
Treated x Post 2020	2.581***	0.428***	0.223***	0.264***	0.288*	0.623*
<i>Controls</i>						
Total Assets (logs)	0.791	0.013	0.240**	0.094	0.001	0.253
ROE	0.001	0.004	0.000	-0.002	0.000	0.004
Leverage	2.395	-0.232	-0.090	0.612	-0.202	0.459
Constant	-5.708	0.157	-2.492*	-0.910	0.576	-0.876
Observations	236	236	236	236	236	236
Number of Groups	54	54	54	54	54	54
Number of Treated Groups	42	40	43	40	28	37
Adj. R ² (within)	0.674	0.387	0.609	0.570	0.303	0.680

Note: This table shows the regression model in eq. (2). The variable “Pages with Rec.” is used with an inverse hyperbolic sine (IHS) transformation. Stars denote p-values * p < 0.1, ** p < 0.05, *** p < 0.01.

5 Robustness

We run several robustness analyses for both analyses on disclosure determinants and impact of policy communications. In relation to the former, our results are robust when using ROA or pretax profit as alternative measure for firms’ profitability. We further check the robustness of our results by including all possible combinations of our indicator variables (i.e., PRA-Regulated, Bank Bus. Model, and Non-life Bus. Model) and find no significant difference to our baseline estimates.²⁴

As mentioned before, our baseline estimates exclude ownership data given the poor data coverage. Nevertheless, Table 6 reports the results when including ownership variable in the regression. We observe that the significant increasing trend in disclosures over time still holds. Size still significantly matters across most of disclosures. However, the coefficient of ownership variable is not significant. Therefore, we will keep our

²⁴These results are available upon request.

baseline specification for the remaining robustness tests.

Table 6: Climate-related Disclosures and Firm Characteristics (with Ownership)

	Sum of Rec.	Governance Avg.	Strategy Avg.	Risk Mgmt. Avg.	Metrics and Targets Avg.	Pages with Rec. (IHS)
<i>Time Trend</i>						
Year	1.468***	0.161***	0.158***	0.142***	0.083***	0.668***
<i>Financial Variables</i>						
Total Assets (logs)	0.838***	0.120***	0.082***	0.098***	0.020	0.399***
ROE	0.020	0.007**	-0.000	0.002	0.000	0.009
Leverage	3.091***	0.417**	0.143	0.247	0.362*	0.716
Ownership	0.004	0.000	0.001	0.001	-0.001	-0.001
<i>Indicator Variables</i>						
Sector	-1.882***	-0.372***	-0.193**	-0.212**	0.026	-0.647*
PRA-Regulated	-2.537***	-0.013	-0.312***	-0.162	-0.363***	-1.284***
Bank Bus. Model	3.671***	0.314***	0.311***	0.321***	0.383**	1.398***
Non-Life Bus. Model	3.212***	0.140	0.266**	0.140	0.571***	1.557***
Constant	-2965.510***	-325.511***	-318.217***	-286.480***	-166.800***	-1348.845***
Observations	163	163	163	163	163	163
Adjusted R^2	0.556	0.418	0.538	0.458	0.242	0.561

Note: This table shows the regression model in eq. (1). The variable “Pages with Rec.” is used with an inverse hyperbolic sine (IHS) transformation. Stars denote p-values * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Regarding the study on the impact of regulatory actions, we also consider a cleaner way to define the treated and control groups. Precisely, in the baseline specification described in Section 4, all firms whose level of disclosure was at least the average level of disclosure in 2017 are included in the control group. Here, we impose an additional requirement according to which firms are only considered as part of control group if they are also PRA-regulated firms. Table 7 show our results.

Table 7: **Difference-in-Differences Results (only PRA-firms)**

	Sum of Rec.	Governance Avg.	Strategy Avg.	Risk Mgmt. Avg.	Metrics and Targets Avg.	Pages with Rec. (IHS)
<i>Time variables</i>						
Post 2018	1.050**	0.060	0.163*	0.057	0.120**	0.453**
Post 2019	1.616*	0.100	0.210*	0.162*	0.043	0.881***
Post 2020	-0.630	-0.162	-0.091	0.090	-0.066	0.393*
<i>Treatment variables</i>						
Treated x Post 2018	0.021	0.099	-0.034	-0.037	-0.067	-0.269
Treated x Post 2019	2.017**	0.161	0.296**	0.272**	0.235*	1.187***
Treated x Post 2020	2.410***	0.459***	0.223***	0.154*	0.261*	0.275
<i>Controls</i>						
Total Assets (logs)	1.318*	0.220	0.232*	0.148	0.116	0.812**
ROE	0.017	0.005	0.003	0.000	-0.001	0.002
Leverage	3.270	-0.701	0.027	0.753*	0.094	0.553
Constant	-12.130	-2.112	-2.418*	-1.578	-0.768	-7.273*
Observations	153	153	153	153	153	153
Number of Groups	37	37	37	37	37	37
Number of Treated Groups	23	22	24	22	13	16
Adj. R ² (within)	0.738	0.451	0.681	0.636	0.330	0.782

Note: This table shows the regression model in eq. (2). The variable “Pages with Rec.” is used with an inverse hyperbolic sine (IHS) transformation. Stars denote p-values * p < 0.1, ** p < 0.05, *** p < 0.01.

We further consider defining treated and control groups as in Bolton & Kacperczyk (2021), i.e., without additionally requiring firms to be PRA-regulated to be considered part of the treated firms, as we do in our baseline results on table 5. Table 8 shows these results.

Table 8: **Difference-in-Differences Results (Bolton & Kacperczyk 2021)**

	Sum of Rec.	Governance Avg.	Strategy Avg.	Risk Mgmt. Avg.	Metrics and Targets Avg.	Pages with Rec. (IHS)
<i>Time variables</i>						
Post 2018	1.034**	-0.028	0.098*	0.047	0.111***	0.340***
Post 2019	1.856***	0.216**	0.199***	0.190***	0.062	0.949***
Post 2020	-0.381	-0.044	-0.082*	0.009	-0.033	0.142
<i>Treatment variables</i>						
Treated x Post 2018	-0.058	0.191	0.066	-0.050	-0.022	-0.062
Treated x Post 2019	2.333***	0.151	0.301***	0.335***	0.257**	1.141***
Treated x Post 2020	1.786***	0.218*	0.196***	0.184**	0.231*	0.400
<i>Controls</i>						
Total Assets (logs)	0.472	0.042	0.172	0.081	0.035	0.261
ROE	0.014	0.006	0.002	-0.002	0.001	0.007
Leverage	3.033	-0.288	-0.025	0.619	-0.276	0.164
Constant	-2.204	-0.184	-1.731	-0.768	0.186	-0.943
Observations	236	236	236	236	236	236
Number of Groups	54	54	54	54	54	54
Number of Treated Groups	49	49	50	46	32	42
Adj. R ² (within)	0.714	0.422	0.674	0.624	0.363	0.722

Note: This table shows the regression model in eq. (2). The variable “Pages with Rec.” is used with an inverse hyperbolic sine (IHS) transformation. Stars denote p-values * p < 0.1, ** p < 0.05, *** p < 0.01.

The key insights from our baseline regression still hold across all these alternative specifications.

Finally, given our dependent variables consist of count data, we also run several Poisson regressions to check robustness. Those results stay qualitatively similar, and are reported in Appendix A.2.

6 Conclusion

In this paper, we investigate climate-related voluntary disclosures in the UK banking and insurance sectors. We first use NLP techniques and Machine Learning classifiers to retrieve, from firms' public reports, climate-related information disclosed in line with different TCFD recommendations. We then examine which firms' characteristics are related to their disclosure decisions. We also study how firms' climate-related disclosures relate to the publication of climate-related regulatory documents.

We find that climate-related disclosures by UK banks and insurers have increased over time with significant jumps of disclosures in 2019 and 2020. This increasing trend in disclosures is uniform across all TCFD themes and across both sectors. However, there is still a significant number of firms that do not disclose any information. In terms of disclosure determinants, disclosure level is found to be positively associated with firms' size. Finally, we find evidence of a significant effect of regulatory announcements on firms' decisions to disclose climate-related information.

Our findings provide some interesting insights for policy makers when considering to make climate-related disclosure mandatory. Our results suggest that prior to regulatory interventions only a fraction of firms disclosed climate-related information in line with TCFD.²⁵ and these were on average larger firms. This gap in voluntary disclosures creates a case for regulatory interventions encouraging smaller firms to disclose too. Indeed, our results suggests that regulators setting clear timelines for mandatory disclosures can help accelerate the trend in disclosures, which leads to convergence across firms.

²⁵Whether it is a lack of willingness or ability to produce information is hard to say given our analyses.

References

- Arnould, G., Guin, B., Ongena, S. & Siciliani, P. (2020), '(When) do banks react to anticipated capital reliefs?', *Bank of England Staff Working Paper* **889**.
- Basten, C. (2020), 'Higher bank capital requirements and mortgage pricing: Evidence from the counter-cyclical capital buffer', *Review of Finance* **24**(2), 453–495.
- Bingler, J. A., Kraus, M., Leippold, M. & Webersinke, N. (2022), 'Cheap talk and cherry-picking: What climatebert has to say on corporate climate risk disclosures', *Finance Research Letters* **47**, 102776.
URL: <https://www.sciencedirect.com/science/article/pii/S1544612322000897>
- Bolton, P. & Kacperczyk, M. (2021), 'Carbon Disclosure and the Cost of Capital', *Available at SSRN* **3755613**.
URL: <https://dx.doi.org/10.2139/ssrn.3755613>
- Brown, M., Guin, B. & Kirschenmann, K. (2016), 'Microfinance Banks and Financial Inclusion', *Review of Finance* **20**(3), 907–946.
- Card, D. & Krueger, A. B. (1994), 'Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania', *American Economic Review* **84**(4), 772–793.
- Chawla, N. V., Bowyer, K. W., Hall, L. O. & Kegelmeyer, W. P. (2002), 'Smote: synthetic minority over-sampling technique', *Journal of artificial intelligence research* **16**, 321–357.
- Clarkson, P. M., Li, Y., Richardson, G. D. & Vasvari, F. P. (2008), 'Revisiting the relation between environmental performance and environmental disclosure: an empirical analysis', *Accounting, Organizations and Society* **33**(4-5), 303–327.
- De Jonghe, O., Dewachter, H. & Ongena, S. (2020), 'Bank capital (requirements) and credit supply: Evidence from pillar 2 decisions', *Journal of Corporate Finance* **60**, 101518.
- Fatouh, M., Giansante, S. & Ongena, S. (2021), 'Economic support during the COVID crisis. Quantitative easing and lending support schemes in the UK', *Economics Letters* **209**, 110138.
URL: <https://www.sciencedirect.com/science/article/pii/S0165176521004158>
- Fougère, D. & Jacquemet, N. (2020), 'Policy Evaluation Using Causal Inference Methods', *IZA Institute of Labor Economics Discussion Paper Series* .
- Garbarino, N. & Guin, B. (2021), 'High water, no marks? Biased lending after extreme weather', *Journal of Financial Stability* **54**, 100874.
- Gray, R., Kouhy, R. & Lavers, S. (1995), 'Corporate social and environment reporting: a review of the literature and a longitudinal study of uk disclosure', *Accounting, Auditing & Accountability Journal* **8**(2), 47–77.

- Gropp, R., Mosk, T., Ongena, S. & Wix, C. (2019), ‘Banks response to higher capital requirements: Evidence from a quasi-natural experiment’, *The Review of Financial Studies* **32**(1), 266–299.
- Hahn, R. & Luelfs, R. (2014), ‘Legitimizing negative aspects in gri-oriented sustainability reporting: a qualitative analysis of corporate disclosure strategies’, *Journal of Business Ethics* **123**(3), 401–420.
- Hahn, R., Reimsbach, D. & Schiemann, F. (2015), ‘Organizations, climate change, and transparency: Re-viewing the literature on carbon disclosure’, *Organization & Environment* **28**(1), 80–102.
- Jiménez, G., Ongena, S., Peydró, J.-L. & Saurina, J. (2017), ‘Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the Spanish dynamic provisioning experiments’, *Journal of Political Economy* **125**(6), 2126–2177.
- Jouvenot, V. & Krueger, P. (2019), ‘Mandatory corporate carbon disclosure: Evidence from a natural experiment’, *Available at SSRN 3434490* .
- Kashyap, A. K. & Stein, J. C. (2000), ‘What Do a Million Observations on Banks Say about the Transmission of Monetary Policy?’, *American Economic Review* **90**(3), 407–428.
URL: <https://www.aeaweb.org/articles?id=10.1257/aer.90.3.407>
- Luccioni, A. & Palacios, H. (2019), Using natural language processing to analyze financial climate disclosures, *in* ‘ICML 2019 Workshop on Climate Change: How Can AI Help?’.
- Rodnyansky, A. & Darmouni, O. M. (2017), ‘The Effects of Quantitative Easing on Bank Lending Behavior’, *The Review of Financial Studies* **30**(11), 3858–3887.
URL: <https://doi.org/10.1093/rfs/hhx063>
- Varini, F. S., Boyd-Graber, J., Ciaramita, M. & Leippold, M. (2020), ‘ClimaText: A dataset for climate change topic detection’, *arXiv preprint arXiv:2012.00483* .
- Verrecchia, R. E. (1983), ‘Discretionary disclosure’, *Journal of Accounting & Economics* **5**, 179–194.
- Webersinke, N., Kraus, M., Bingler, J. A. & Leippold, M. (2021), ‘ClimateBert: A Pretrained Language Model for Climate-Related Text’.
URL: <https://arxiv.org/abs/2110.12010>

A Appendix

A.1 Machine Learning Methodology

Our approach consists of four consecutive steps starting with the collection of corporate reports and ending with our final TCFD-disclosure dataset. In the first step, we manually collect corporate reports of the biggest banks, building societies and insurance companies between 2016 and 2020 in the UK. Next, in step two, we parse over 800 reports to obtain our corpus with over 100 thousand pages. Hence, in step three, to reduce the large dimensionality of our corpus we first identify climate-related pages using NLP techniques and machine learning classifiers. Finally, we identify the information disclosed on each corporate report’s climate-related page that is closely in line with each of the individual recommendations published by the TCFD.

In the remaining of this section we describe these steps in more detail.

Step 1: Data Selection (Corporate Reports Collection)

The first step, before the manual collection of corporate reports, is defining the universe of financial firms to be analysed. Our firm-level sample consists on the largest 52 building societies and banks, as well as the biggest 65 insurance companies in the UK, which add together to a total of 117 UK-based firms. Additionally, when applicable, we include the international parent (or group) company of firms in our UK sample, leaving us with a final firm-level list encompassing 137 companies.

We then proceed to manually collect annual reports (in PDF format) between 2016 and 2020 from the companies’ own web sites, and in cases where these reports were not available we search and collect reports from “Company Houses.”^I

Additional to firms’ annual reports, we also collect any corporate reports that might include climate-related information. For example, among others we include in our sample so-called ESG, CSR, Sustainability and even TCFD reports. For simplicity, henceforth, we refer to these additional reports as “other” reports. All in all, across the five years in our sample we collect 824 corporate filings, of which 641 are annual reports and 183 are other reports.

Step 2: NLP & Parsing (Data extraction)

In order to extract the information contained in the textual reports, we parse over reports whose PDF files are either text-formatted or image-formatted. For text-formatted PDF files, we used Tika-Python.^{II} It extracts all the text content and stores it as long string. While this approach has some minor difficulties when parsing files with complex visual design, the performance is good enough for our purposes here. For image-formatted PDF files, however, parsing is more complicated. Thus, we employ Optical Character Recognition (OCR) tools. Briefly, this approach first converts each document page into a black and white image that is then scanned using the OCR tool called Pytesseract to obtain the information in the image-based reports.


^ICompany Houses is an executive agency, sponsored by the UK Department for Business, Energy & Industrial Strategy, where limited companies’ information is registered and made available to the public in the UK, see <https://www.gov.uk/government/organisations/companies-house/about>.

^{II}Tika-Python is a port of the Apache Tika library that allows us to use Apache Tika as a Python library.

Finally, all information from both text-formatted and image-formatted PDF files is stored in a JSON file, which can then be accessed as a Pandas data frame. Concretely, to simplify information retrieving we use this pipeline, which first, locates table of contents of annual reports, and then parse the table of contents to extract section headings using multiple rules to achieve better performance. Then, in this pipeline, we match the extracted headings with actual page numbers for every annual report’s page. Finally, the table of contents are mapped to an expanded data frame where each observation (row) represents individual pages from corporate reports.

Thus, our unit of analysis are individual corporate report pages, and our corpus contains 112,354 pages. We chose to use pages as the unit of analysis instead of paragraphs, since otherwise we would have around 1 million paragraphs, depending on the method used to distinguish paragraphs. This is an order of magnitude larger and would increase the number of samples to be manually labelled in our supervised learning approach described below. Non-automated inspection of the documents in our corpus would require close to 3,750 reading hours under certain conservative assumptions.^{III} Moreover, such a labour-intensive approach is not only impractical, but it is also susceptible to the judgment of the person responsible for reading the information disclosed in each page, increasing thus, the possibility of inconsistencies in such a labour-intensive manual classification task.

Figure A.1: Climate-related pages in corporate reports



Note: Total number of pages, and climate-related pages in full sample.

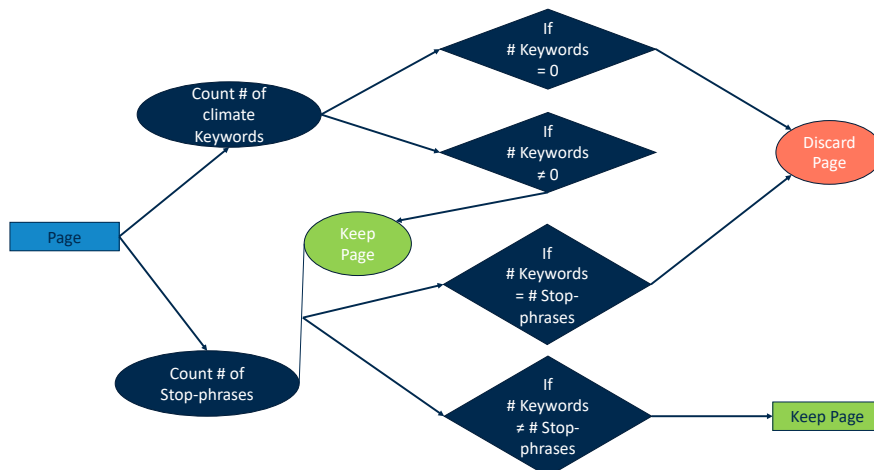
Nevertheless, as one would expect, the vast majority of pages in our sample contain no climate-related information at all. Hence, to reduce the workload for the following tasks using simple rules to identify and discard pages that are highly unlikely to contain climate-related information.

The simple rules employed here rely on climate keywords and stop-phrases. Figure A.2 depicts our pipeline to reduce the original sample with these rules. We use the list of eight climate keywords introduced

^{III}We assume an average reading speed of 250words/minute, and that each page contains on average 500 words.

by Luccioni & Palacios (2019) together with five further keywords to count the number of times that our climate keywords appear in each page. Pages containing none of our 13 climate keywords are discarded.^{IV} Additionally, we further reduce our sample using a list of stop-phrases.^V Concretely, we discard all pages where the number of keywords is the same as the number of stop-phrases. Some examples of the stop-phrases used are “economic environment”, “sustainable returns” and “climate of low interest rates”.^{VI}

Figure A.2: Natural Language Processing Pipeline



Note: The figure depicts our NLP pipeline, i.e., how we reduce the dimensionality of our original corpus using climate keywords and stop-phrases. To do so, each page’s number of keywords and stop-phrases are counted, and those pages without any keywords or where the number of keywords and stop-phrases coincide are discarded.

These two simple rules reduce our sample to 19% of the initial observations, leaving us with 21,478 pages containing climate keywords.

Step 3: Machine Learning I (Identifying Climate-related Pages)

With our truncated sample encompassing only the 21 thousand pages that contain climate keywords, we proceed by identifying pages that disclose truly climate-related information. To do this, we rely on a supervised learning approach, where 300 pages were randomly selected and manually labelled, i.e., each page was read and classified as being truly climate-related or not.

Concretely, to account for class imbalances, we first split the set of pages in into low and high keyword groups according to the distribution of keywords. Then, our training/validation set is composed by randomly selecting 100 pages from the group with lower keyword counts and 200 pages from the group with higher

^{IV}Our full list of climate keywords are ‘biodiversity’, ‘carbon’, ‘climate’, ‘ecology’, ‘environment’, ‘emission’, ‘pollution’, ‘sustainable’, as well as ‘environmental’, ‘ESG’, ‘TCFD’, ‘sustainability’ and ‘green’.

^VStop-phrases, or stop-words, are words that are filtered out (ignored) to prevent them from being considered by search engines. In our specific case, stop-phrases are instances where one of our climate keywords are used in a context that is not related to climate-related information.


^{VI}Our full list of stop-phrases is: ‘economic environment’, ‘operating environment’, ‘geopolitical environment’, ‘control environment’, ‘interest rate environment’, ‘changing environment’, ‘rate environment’, ‘technology environment’, ‘risk environment’, ‘external environment’, ‘internal environment’, ‘commercial environment’, ‘market environment’, ‘business environment’, ‘sustainable value for shareholders’, ‘sustainable returns’, ‘sustainable dividend’, ‘sustainable capital strength’, ‘climate of low interest rates’.

keyword. Moreover, our test set is built from 100 additional pages. In total, our training set contains 124 positive samples and 176 negatives, while in our test set there are 27 and 73 pages, respectively.

In order to mitigate the information leakage problem of an over-sampled data set, we further split our 300 pages into a training (240 pages) and validation (60 pages) set. Furthermore, in our training set 40% of samples are negative. We address this imbalance by using the Synthetic Minority Over-sampling Technique (SMOTE) by Chawla et al. (2002) to up-sample the positives samples. Hence, the balanced data set contains 141 positive and 141 negative samples.

We chose a Random Forest Classifier with keyword-only feature representation to identify truly climate-related pages, since the testing sample’s accuracy of this model is 95%.^{VII} Our preferred keyword-only Random Forest classifier identifies close to 1/3 of our pages containing climate-keywords to be truly climate-related. More specifically, out of our 21,478 pages in our truncated sample, our machine learning model classifies 7,402 pages as climate-related and 14,076 pages as climate-irrelevant. Note that this implies that only 6.6% of our initial sample contains climate-related information.

Figure A.3: **Evolution of the Share of Climate-related Pages**



Note: Share of climate-related pages per year.

Step 4: Machine Learning II (TCFD-disclosure dataset)

Next, we identify disclosed information that is line with the recommendations by the TCFD in our sample of climate-related pages. Given the complex nature of the classifying problem at hand we chose to develop 11 independent machine learning models to cope with our multi-label classification problem.^{VIII} To do so

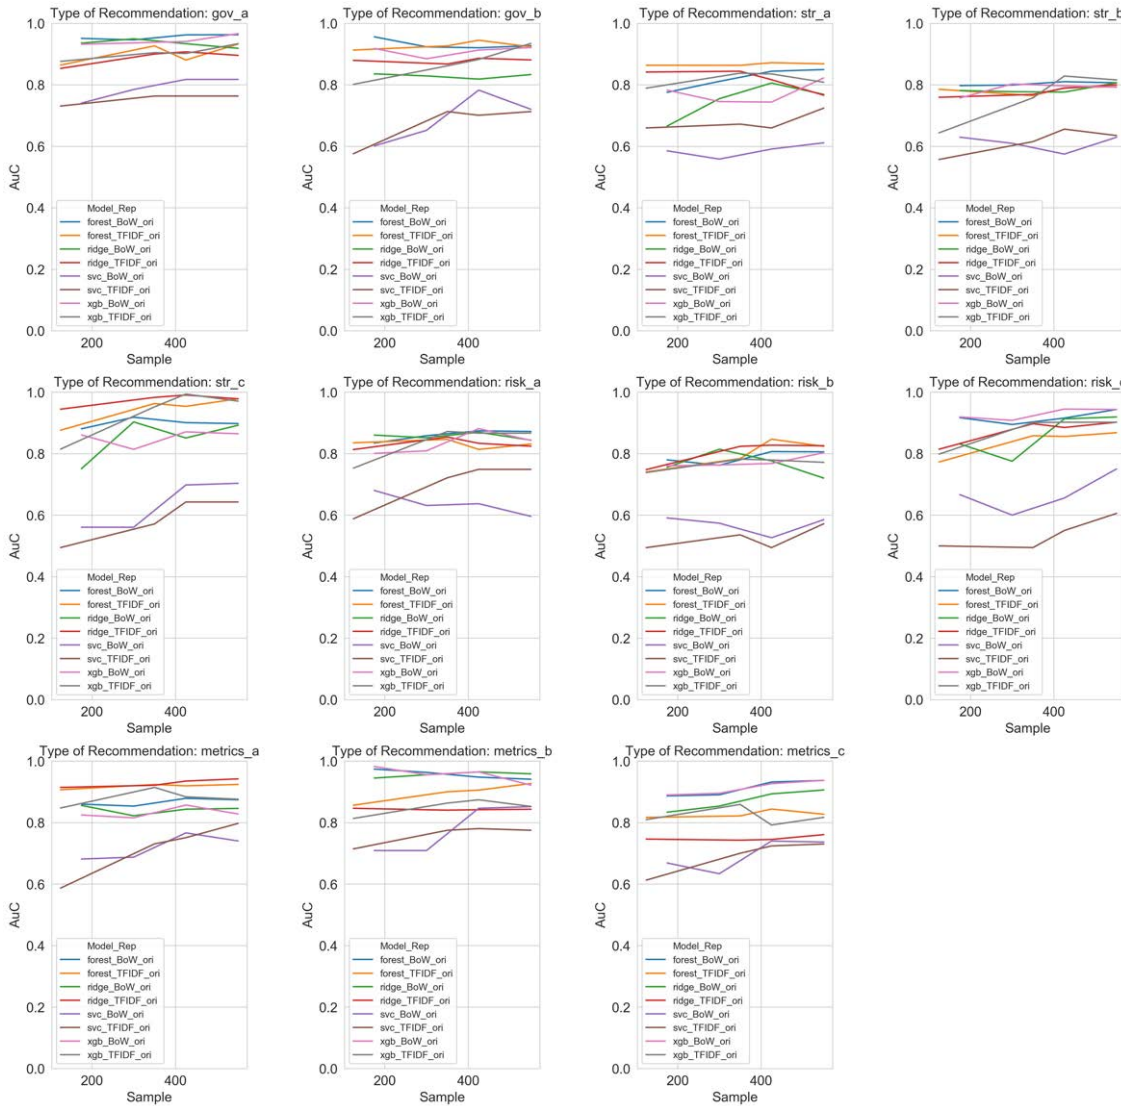
^{VII}We also trained and tested alternative models and feature representations. Concretely, besides the Random Forest model we also considered a Linear Support-Vector Machine classifier and an Extreme Gradient Boosting (XGBoost) classifier. Moreover, besides the keyword-only representation, we considered Bag-of-Words (BoW) and Term Frequency-Inverse Document Frequency (TF-IDF) representations, as well as reduced forms of both BoW and TF-IDF representations obtained using Latent Semantic Analysis (LSA).

^{VIII}This is a multi-label classification problem since in each observation (page) in our sample there exist the possibility that firms disclose information related to multiple TCFD recommendations simultaneously.

we rely again on a supervised-learning approach, where we manually label each page in our training sample according to the individual TCFD recommendations.

Our training sample consists of 9% of our climate-related pages. We chose to stop the manual labelling process at 650 pages given that the performance among our candidate machine learning classifiers did not improve after 400 training samples. This can be seen in figure A.4 where we present the results of our learning curve analysis for (a subset of) the models in consideration.

Figure A.4: Learning Curve Analysis



Note: Learning Curve Analysis for original feature representations.

More concretely, the models that we consider for each of our 11 recommendations are a (Balanced) Random Forest classifier, an Extreme Gradient Boosting (XGBoost) classifier as well as a Ridge classifier. Similar to our previous machine learning problem above, we consider four alternative feature representations for each classifier and model for each individual label/recommendation. These are BoW and TF-IDF feature representations, as well as LSA-reduced forms of both BoW and TF-IDF feature representations.

We train and validate our 11 independent classifiers using a nested cross-validation approach. To do so, we first split our 650 manually labelled observations into a training-validation sample (550 pages), and a testing set (100 pages). Next, in the outer cross-validation we chose five folds ($K=5$), where at each iteration four folds are used for training and 1 for validation. Furthermore, for the inner cross-validation we select 3 inner folds ($k=3$) on the training set, i.e., using four K -folds at the time, to fine-tune our model’s hyperparameters. Our tuning approach consist on a grid search to obtain the hyperparameters to optimize our performance measure, i.e., the “Area under the Receiver Operating Characteristic Curve” or AUC for short, using the inner 3- k folds to find the best estimator, and then validate this specification on the holdout outer K -fold. This is then repeated until all of the five outer K -folds have been used for validation.

All in all, for our multi-label problem at hand with 11 independent recommendations ($L=11$) and three alternative models ($M=3$) with four different features representations ($F=4$), our nested-cross-validation approach considers 1980 models in total ($= K \times k \times M \times F \times L = 5 \times 3 \times 3 \times 4 \times 11$), out of which we can select the best 11 performances out of 660 model-feature combinations ($= K \times M \times F \times L = 5 \times 3 \times 4 \times 11$).

We select, for each of the 11 labels, the best model-feature combination according to the average validation AUC score. Table A.1 shows for each of the 11 TCFD recommendations the best model as well as the validation accuracy, F1 and AUC scores. Only in one case (governance A) a reduced form representation is selected. Moreover, the Balanced Random Forest classifiers performs best in six cases, whereas the Ridge and the XGB classifiers outperform the alternatives in three and two instances, respectively.

Table A.1: Nested Cross-Validation Results

Recommendation	ML-Model	Features Representation	AUC	F1-Score	Accuracy
Governance A	Ridge	TF-IDF (reduced)	0.949818	0.706852	0.92
Governance B	Ridge	TF-IDF (original)	0.893122	0.551529	0.88
Metrics & Targets A	BRF	BoW (original)	0.875043	0.575158	0.80
Metrics & Targets B	XGB	TF-IDF (original)	0.957699	0.591826	0.95
Metrics & Targets C	XGB	BoW (original)	0.859370	0.525609	0.84
Risk Management A	BRF	TF-IDF (original)	0.859253	0.488876	0.77
Risk Management B	BRF	TF-IDF (original)	0.890731	0.455261	0.77
Risk Management C	Ridge	TF-IDF (original)	0.893415	0.458095	0.92
Strategy A	BRF	TF-IDF (original)	0.840566	0.649228	0.78
Strategy B	BRF	BoW (original)	0.828412	0.566092	0.75
Strategy C	BRF	BoW (original)	0.936061	0.441299	0.85

Note: BRF stands for “Balanced Random Forest” and XGB stands for “XGBoost”, which is an optimized extreme gradient boosting framework. TF-IDF stands for “term frequency–inverse document frequency” and BoW for “bag of words”. Finally, AUC is an abbreviation for “Area under the Receiver Operating Characteristic Curve.”

The test sample is then used to assess the performance of the 11 independent fine-tuned machine learning classifiers, similar to an econometric out-of-sample forecast exercise. Table A.2 presents the testing sample results, which imply that our models’ performance is quite high with an average AUC score of 0.88.

Table A.2: Test Sample Results

Recommendation	ML-Model	Features Representation	AUC	F1-Score	Accuracy
Governance A	Ridge	TF-IDF (reduced)	0.888704	0.666667	0.91
Governance B	Ridge	TF-IDF (original)	0.839431	0.500000	0.84
Metrics & Targets A	BRF	BoW (original)	0.914452	0.578947	0.84
Metrics & Targets B	XGB	TF-IDF (original)	0.995392	0.857143	0.98
Metrics & Targets C	XGB	BoW (original)	0.794454	0.344828	0.81
Risk Management A	BRF	TF-IDF (original)	0.906250	0.564103	0.83
Risk Management B	BRF	TF-IDF (original)	0.845982	0.600000	0.84
Risk Management C	Ridge	TF-IDF (original)	0.904509	0.600000	0.92
Strategy A	BRF	TF-IDF (original)	0.766798	0.538462	0.76
Strategy B	BRF	BoW (original)	0.862786	0.654545	0.81
Strategy C	BRF	BoW (original)	0.960938	0.421053	0.89

Note: BRF stands for “Balanced Random Forest” and XGB stands for “XGBoost”, which is an optimized extreme gradient boosting framework. TF-IDF stands for “term frequency–inverse document frequency” and BoW for “bag of words”. Finally, AUC is an abbreviation for “Area under the Receiver Operating Characteristic Curve.”

Finally, we use our 11 independent models to label the remaining 6752 climate-related pages, i.e., to identify in each of these pages whether firms disclosed information in line with any of the recommendations published by the TCFD or not. This generates our novel TCFD-disclosure dataset, consisting of 11 indicator variables representing each of the TCFD recommendations.

A.2 Poisson regressions

In this appendix, we report the results of several Poisson regressions. First, we check if the analysis on the firm characteristics of the determinants of disclosure presented in table 3 using a pooled OLS model remain unchanged when estimated using a Poisson model. Table A.3 shows that the main determinant across all dependent variables remains (log) total assets and the time trend.

Table A.3: **Climate-related Disclosures and Firm Characteristics (Poisson)**

	Sum of Rec.	Governance Avg.	Strategy Avg.	Risk Mgmt. Avg.	Metrics and Targets Avg.	Pages with Rec.
<i>Time Trend</i>						
Year	0.244***	0.319***	0.302***	0.299***	0.129***	0.536***
<i>Financial Variables</i>						
Total Assets (logs)	0.179***	0.243***	0.176***	0.244***	0.111***	0.361***
ROE	0.004	0.014***	0.002	0.007	0.001	0.009
Leverage	0.016	-0.097	-0.227	0.337	0.021	-0.447
<i>Indicator Variables</i>						
Sector	-0.181*	-0.312	-0.211	-0.371**	0.052	-0.652***
PRA-Regulated	-0.395***	-0.203	-0.625***	-0.464**	-0.297	-0.754*
Bank Bus. Model	0.372**	0.330	0.496**	0.563**	0.207	0.824**
Non-Life Bus. Model	0.358**	0.238	0.451*	0.338	0.399*	0.590
Constant	-493.072***	-646.362***	-611.620***	-605.618***	-260.049***	-1082.239***
Observations	236	236	236	236	236	236
Pseudo R^2	0.244	0.160	0.186	0.196	0.060	0.600

Note: This table shows the (Poisson) regression model in eq. (1). Stars denote p-values * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Finally, we estimate our baseline difference-in-differences model presented in equation 2 using a Poisson model. To estimate this model a few changes in our baseline specification are necessary. First, the dependent variables are modify such that instead of reflecting the TCFD theme average, they now count how many of each recommendation within a theme are reported by a given company on a given year.^{IX} Additionally, due to data limitations we employ firm and time fixed effects, instead of firm and sector-time fixed effects.

We find that our main results are robust under a Poisson regression as shown in table A.4. In fact, the results under this specification show that for most of our dependent variables (excluding metrics and targets) even the policy publications in 2018 and 2019 had a significant positive effect on our treated group’s disclosures. Nevertheless, we still prefer our baseline specification given that under the Poisson model we lose a significant amount of observations (and even groups), even if the results stay qualitatively similar.

Table A.4: **Difference-in-Differences Results (Poisson)**

	Sum of Rec.	Governance	Strategy	Risk Mgmt.	Metrics and Targets	Pages with Rec.
<i>Time variables</i>						
Post 2018	0.367***	0.382**	0.481***	0.461***	0.191**	0.792***
Post 2019	0.237***	0.337***	0.189***	0.371***	0.087**	0.691***
Post 2020	-0.066**	-0.038	-0.080*	-0.075	-0.001	0.301**
<i>Treatment variables</i>						
Treated x Post 2018	0.341**	16.111***	1.330***	1.476***	0.828	-0.038
Treated x Post 2019	0.601***	0.959*	0.921***	1.130***	1.499	0.397**
Treated x Post 2020	0.290***	0.426*	0.296**	0.465***	0.746**	0.693***
<i>Controls</i>						
Total Assets (logs)	0.043	-0.104	0.133*	-0.385***	0.089*	0.122
ROE	-0.007	0.005	-0.010	-0.020**	0.001	0.002
Observations	247	200	233	225	218	247
Number of Groups	52	41	48	46	46	52
Number of Treated Groups	42	40	43	40	28	45

Note: This table shows the regression model in eq. (2). Stars denote p-values * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

^{IX}We also don’t apply the IHS transformation to the ”Pages with Rec.” dependent variable.