A Greenwashing Index *

Elise Gourier
ESSEC Business School and CEPR

Hélène Mathurin
ESSEC Business School

Avril 2025

Abstract

We construct a news-implied index of greenwashing salience. Our index reveals that greenwashing has become more prominent in the past five years. This increase was driven by skepticism towards the financial sector, specifically ESG funds, ESG ratings and green bonds. Furthermore, we build a measure of general anti-ESG sentiment from opinion pieces, and show that increases in greenwashing salience tend to be followed by increases in anti-ESG sentiment. In line with this finding, greenwashing salience affects investors' behavior. Following unexpected increases of greenwashing salience, retail investors invest less in all funds, and institutional investors invest less in funds with high ratings and self-labelled as sustainable.

JEL Classification: G11, G12, Q51

Keywords: greenwashing, climate risk, text analysis, fund flows

*Email: elise.gourier@essec.edu, helene.mathurin@essec.edu. We would like to thank Svetlana Bryzgalova, Alan de Genaro, Javier Gil-Bazo, Samuel Hartzmark, Christian Julliard, Elisabeth Kempf, Johannes Klausmann, Stefano Lovo, Alessandro Melone, Ludovic Phalippou, Alberto Quaini, Zacharias Sautner, and Gustavo Schwenkler for feedback and suggestions on our project. We also thank the participants of seminars at ESCP, the University of Liverpool, SKEMA, Vienna Graduate School of Finance, HEC-HKUST Sustainable Finance Webinar and participants to the Sciences-Po HEC Summer Workshop in Sustainable Finance, NEOMA Conference on Sustainable Finance, the Spring Finance Workshop 2-24, the 17th Financial Risks International Forum, the 8th ESSEC Asset Pricing Breakfast, and the 5th International Workshop in Financial Econometrics for comments. This research has been supported by the JCJC ANR grant ANR-22-CE26-0010-01 ("GreenLies") and a grant from the "Institut Europlace de Finance and Labex Louis Bachelier". Finally, we would like to thank the research assistants who helped with this project.

1 Introduction

The mounting pressure from activists, investors and government agencies to tackle climate change has led some firms to provide misleading information about the environmental soundness of their practices. This phenomenon is often referred to as 'greenwashing'. Greenwashing plausibly hinders the measurement of climate risks and the actions to reduce these risks. It has thus recently been the focus of regulators and policy makers, and the subject of a heated policy debate.[1] In this paper, we build an index of greenwashing salience. We show that greenwashing has become increasingly salient since 2018, and that this increase was driven by suspicions in the financial sector. Furthermore, we find that more greenwashing salience is followed with a stronger anti-ESG sentiment, and by decreases in inflows in green funds. Our results underline the link between media reports on greenwashing and investors' trust in green assets.

Our greenwashing index measures the fraction of climate-related news articles alluding to firms' greenwashing. We apply an advanced Natural Language Processing algorithm to the history of paper-based Wall Street Journal articles between January 1986 and June 2022 (nearly one million articles). The algorithm is trained to identify the articles that allude to greenwashing. Identifying these articles is challenging for several reasons. First, the word "greenwashing" became widely used only recently: it was used in fewer than 10 articles before 2007 and in fewer than 35 articles before 2020. In contrast, it was used in nearly 40 articles between January 2020 and June 2022. Our algorithm is able to learn the words or combinations of words that characterize articles related to greenwashing, even if they do not contain the word itself. Second, the data set is large and unlabelled. Third, it is highly imbalanced: the fraction of all Wall Street Journal articles that are greenwashing-related is small. To make these issues less critical, we proceed in two steps. First, we identify articles related to climate risk, and second, we pinpoint, among these articles, those that mention greenwashing (or equivalent terms). For each step, we use a method that is well adapted for large unlabelled data, namely the keyword and document set discovery method of King, Lam, and Roberts (2017). This method starts from a small set of articles that are highly likely to be climate risk-related, or greenwashing related (e.g., because they contain the word "greenwashing"). It learns the patterns that characterize these articles and expands the initial set in an iterative manner. It yields as output a classification rule of all articles. We evaluate the out-of-sample performance of this classification rule on a set of nearly 1500

[1] New rules have recently been proposed by the Securities and Exchange Commission in the U.S., and by the European Parliament and Council, to reduce greenwashing. We provide details on the actions and regulations adopted in the U.S. and in Europe over time in Section 3.3.

articles that were labelled by hand.

The out-of-sample performance of the algorithm critically depends on how words are embedded into a numerical vector. The simplest approach is to use a bag-of-words representation, which counts the frequency of each word or group of words in an article. We find that in the first step, articles related to climate risk are well identified using this approach. However, it does not allow for the precise identification of greenwashing-related articles in the second step. Indeed, there is no list of words (besides the word "greenwashing" itself) that uniquely characterizes greenwashing-related articles, making this step more challenging. We use the Word2Vec embedding, which accounts, to some extent, for the context of words, and show that it provides a parsimonious solution to this issue. We compare the performance of our method to other rule-based and machine learning methods from the literature, and find that it provides the best trade-off between parsimony and performance.

Similarly to news-implied indices of climate risk, our index can serve as a proxy for actual greenwashing in aggregate, under the assumption that the more firms greenwash, the more they are caught and the media report about it. By definition, companies that are greenwashing try to hide it, which makes the measurement of greenwashing particularly challenging. Our index is an attempt to overcome this challenge.[2]

Articles classified as greenwashing related represent a small part, on average 0.07%, of all selected articles. Although one could see this number as an indication that greenwashing does not matter, it is important to note that this fraction has been steadily increasing since mid-2018, with several peaks near 1% and an average value of 0.5% between 2018 and 2022. In comparison, there is on average 2.4% of all Wall Street Journal articles that are climate risk-related. In the past few years, greenwashing salience has therefore grown to a magnitude of the same order as that of climate risk.

We find that there are three waves during which greenwashing accounted for more than 5% of the text on climate risk. The first wave was from 1990 to 1992, the second from 2006 to 2010 and the third from 2018. The first two waves coincided with accusations of greenwashing involving mostly companies in the oil and gas industry (e.g., Exxon Mobil) and the consumer goods industry (e.g., Proctor & Gamble). Occasional and punctual peaks were furthermore triggered by incidents such as the Volkswagen scandal in 2015-2016. The last wave, from 2018, coincided with greenwashing accusations towards investment firms (e.g., Pax World Management, Blackrock, Deutsche Bank), and with the repeted criticisms of the

[2]It is possible that greenwashing-related reports only uncover the tip of the iceberg. In April 2023, the Wall Street Journal reported that "nearly three-quarters of corporate leaders say most organizations in their industry would be caught greenwashing if they were investigated thoroughly". See https://www.wsj.com/articles/global-executives-say-greenwashing-remains-rife-10a0e273.

metrics used to measure these firms' green credentials. Since 2018, the greenwashing index has consistently been above 8%. We conclude that in the past five years, greenwashing has substantially affected the discussion on climate risk, in a way that is unprecedented. In order to understand what is different in these past five years, we analyze the firms and industries that are mentioned in articles used to build our greenwashing index, as well as the topics of these articles.

We list the names and industries of firms that are mentioned in greenwashing-related articles. We find that investment firms (Blackrock, JPMorgan Chase, Goldman Sachs etc.) are the most mentioned by far. Companies in the oil & gas industry are also prominent, including Exxon Mobil and Chevron. Mapping these firms to the 17 Fama French industries and analyzing the time series evolution of industry mentions confirms that the financial sector has by and large driven the recent increase in greenwashing salience. In contrast, the shares of other industries in the index have remained comparable over time.

In order to gain a more granular understanding of the topics that are most associated to greenwashing in the news, we perform topic analysis on the set of greenwashing-related articles. We use an extension of the standard Latent Dirichlet Allocation (LDA) method, namely the keyATM method of Eshima, Imai, and Sasaki (2023). It delivers, for each topic, the proportion of text that belongs to this topic over time. In line with the previous results, we find that the Financial sector is the main theme in greenwashing-related articles: it accounts for 17.3% of the text in these articles. It is composed of three topics: Asset management (specifically, ESG funds), ESG ratings and green bonds. While the first topic dominates the two others, they have all grown in importance over the past five years at a similar pace. Two other sectors stand out in the topic analysis: the Energy and Construction sectors. However, we find that in the recent years, these two sectors have received, in total, about a fourth of the weight of the financial sector.

We next study two possible effects of greenwashing on financial markets: Does greenwashing fuel anti-ESG sentiment? Does it affect investors' choices?

The answer to both questions is yes. To study the link between greenwashing salience and anti-ESG sentiment, we build an anti-ESG sentiment index from opinion pieces of the Wall Street Journal. They were excluded from the construction of the greenwashing index, but contain useful information on the general opinion towards specific topics. Our sentiment index counts the fraction of climate-related opinion pieces that have negative tonality. We find that there is a large and statistically significant increase in anti-ESG sentiment following an increase in greenwashing salience.

We further study the effect of greenwashing on investors' choices by examining the re-

sponse of market demand in funds with different sustainability and performance ratings, to unexpected shocks in the greenwashing index since August 2018. As these shocks are mainly driven by media concerns about the financial sector, and specifically ESG funds, they may trigger a decrease in demand for funds in general, or for some of them with specific features. We find, using panel regressions, that both retail and institutional investors react to unanticipated shocks in the greenwashing index by decreasing their flows into funds. The targeted funds however differ. Retail investors invest less in all funds, irrespective of their ratings and names, over the three months following the shock in greenwashing salience. This reaction reflects a loss of trust in the financial sector as a whole. It is significant but small in magnitude: a one standard deviation unexpected increase in the greenwashing index is followed by a decrease in retail fund inflows by 3%, in total, over the following three months. In contrast, institutional investors invest less in funds that have a high sustainability and performance rating, and which are self-labelled as sustainable. Their flows into other funds remain largely unaffected. The effect is much larger than for retail investors: a one standard deviation unexpected increase in the greenwashing index is followed by a decrease in targeted fund inflows by more than 20% each month, over the three months following the shock.

2 Literature Review

We build on a literature that uses text analysis to quantify climate-related risks. Engle, Giglio, Kelly, Lee, and Stroebel (2020) apply a rule-based method and compute the similarity of articles of the Wall Street Journal to climate-related texts, to build an index of climate risk. They show how climate risk exposure can be hedged using mimicking portfolios. Ardia, Bluteau, Boudt, and Inghelbrecht (2022) expand the set of news to a large number of newspapers including the New York Times, as well as newswires such as Reuters News. They calculate a lexicon-based climate concern index based on the articles that are categorized as related to climate-change. They show that green stocks outperform brown stocks on days when the climate concern index unexpectedly increases. Sautner, van Lent, Vilkov, and Zhang (2023a) modify the King, Lam, and Roberts (2017) algorithm to identify which fraction of sentences of earning calls are linked to climate risk for each company in their dataset, each year. In a follow-up paper (Sautner, van Lent, Vilkov, and Zhang (2023b)), they test whether this climate change exposure is priced in the cross-section of stocks, and find that the answer varies depending on the time period considered. Faccini, Matin, and Skiadopoulos (2023) use a rule-based method to build factors of physical and transition risks from Reuters news, and study their price.

Building a climate risk time series is for us an intermediary step in the construction of the greenwashing index. We benchmark this climate risk time series to the indices built in the above mentioned papers. We find that our time series has a correlation above 50% with all of these indices. This high correlation indicates that our index captures well the common component of climate risk, and is not much contaminated by the idiosyncracies due to the algorithm and the data source.

All these papers abstract from the existence of greenwashing when measuring climate risk. Cartellier, Tankov, and Zerbib (2024) show that in equilibrium, firms face a trade-off between improving their fundamental environmental value (through, e.g. emission reductions) and improving their environmental communications (using, e.g., greenwashing). Focusing on the latter leads investors to have a biased view of firms' climate risk exposure, and to choose suboptimal portfolio allocations. Inderst and Opp (2024) study, in equilibrium, the conditions under which a taxonomy of ESG investments would prevent greenwashing.

Several recent papers have tried to identify greenwashing firms. Bingler, Kraus, Leippold, and Webersinke (2022a) implement a BERT model that is trained on climate resources, to determine whether companies' TCFD disclosures are mostly cheap talk or not. In a parallel paper, the same authors (Bingler, Kraus, Leippold, and Webersinke (2022b)) analyze companies' annual reports. Finally, recent papers have tried to highlight greenwashing in the financial sector. Brandon, Glossner, Krueger, Matos, and Steffen (2022) find that signatories of the Principles for Responsible Investment outside the United States have better ESG ratings than non-signatories, but that this no longer holds in the United States. Kacperczyk and Peydró (2022) find evidence of banks' greenwashing. Parise and Rubin (2023) show that the timing of purchases and sales of sustainable assets by mutual funds is chosen to manipulate sustainability ratings. Heath, Macciocchi, Michaely, and Ringgenberg (2023) find that socially responsible investment funds do not follow on their promise to impact. Along the same lines, Atta-Darkua, Glossner, Krueger, and Matos (2023) raise doubts about the effectiveness of investor-led initiatives in reducing corporate emissions and helping an all-economy transition to "green the planet". Dumitrescu, Gil-Bazo, and Zhou (2023) propose a definition of greenwashing for mutual funds, and find that according to their definition, a third of self-labelled ESG funds are greenwashing.

In contrast to these papers, we do not attempt to measure actual greenwashing at the firm or fund level, but instead focus on greenwashing salience in the media.

3 Greenwashing

In this section, we highlight the absence of a legal definition of greenwashing, and provide an overview of the latest definitions proposed. We briefly review the main drivers of greenwashing given in the literature, as well as the recent actions and regulations taken in the E.U. and the U.S. to prevent greenwashing.

3.1 Definition

The term "greenwashing" was first used by Jay Westerveld, a New York environmentalist, in a 1986 essay about the hotel industry's practice of encouraging guests to reuse towels to "save the environment". Early reports of greenwashing include the SEC charging Exxon Corporation with violating the antifraud provisions of the federal securities laws for making false and misleading statements about the environmental impact of the 1989 Exxon Valdez oil spill in Alaska. This was one of the many incidents that exposed the gap between corporate environmental rhetoric and reality.

These incidents have led to a significant effort from both academics and practitioners to better understand greenwashing, with the goal of reducing it. This objective is however, by the very nature of greenwashing, difficult to achieve. Surprisingly, there is as of now yet no legal definition of greenwashing. In 2020, the E.U. Taxonomy Regulation defines greenwashing, in the context of the regulation, as "the practice of gaining an unfair competitive advantage by marketing a financial product as environmentally friendly, when in fact basic environmental standards have not been met". The European Securities and Markets Authority (ESMA) has recently made a step towards adopting a legal definition of greenwashing, but received negative reactions from industry. For example, the U.S.-based Investment Company Institute (ICI), which represents investment funds, responded that "Seeking to adopt a general definition of greenwashing or enshrine it in legislation would be counterproductive.", Reuters (2023).

Despite this pushback, in its recently released progress report on greenwashing (European Securities and Markets Authority (2023)), the ESMA defines it as "a practice where sustainability-related statements, declarations, actions, or communications do not clearly and fairly reflect the underlying sustainability profile of an entity, a financial product or financial service. This practice may be misleading to consumers, investors, or other market participants." The report highlights that marketing materials, labels and voluntary reporting are most exposed to greenwashing risk, and identifies several ways that can be used for greenwashing, such as cherry-picking, omission, ambiguity, empty claims (including exaggeration),

and misleading use of environmental, social and governance (ESG) terminology.

With the increasing popularity of the ESG framework, the definition of greenwashing has in the past years often been broadened to misleading claims on ESG issues.

3.2 Drivers of greenwashing

The drivers of greenwashing have been extensively studied in the management literature. Delmas and Burbano (2011) classify drivers into three categories: external, organizational, and individual. External drivers include pressures from both non-market actors (regulators and NGOs) and market actors (consumers, investors, and competitors). Organizational drivers include firms' incentive structure and ethical climate, effectiveness of intra-firm communication, and organizational inertia. Individual drivers include cognitive biases such as narrow decision framing, hyperbolic intertemporal discounting and optimistic bias. They find that the limited regulation of greenwashing, combined with the limited enforcement of the existing regulation, is the key driver of greenwashing. In a recent paper, Cartellier, Tankov, and Zerbib (2024) build a general equilibrium model and derive the optimal greenwashing strategy of firms. They show that greenwashing can be curbed by enhancing transparency, and by fostering technological innovation. We briefly present in Section 3.3 the regulations in place and the current actions that aim to prevent greenwashing.

The role of transparency and disclosure has been studied in the accounting literature, see, e.g., Deegan (2002). Flammer (2021) and Ilhan, Krueger, Sautner, and Starks (2023) further highlight investors' increased demand for transparency. As of now, however, the disclosure of climate-related claims suffers from both a lack of consensus on the format of disclosure, and requirements that vary throughout countries. Some reporting frameworks, such as for example those of the Sustainability Accounting Standards Board (SASB), propose material topics and metrics that measure financial impact, and are thus designed solely for investors. Other frameworks such as the ones of the Global Reporting Initiative (GRI), the Task Force on Climate-related Financial Disclosures (TCFD) and the Carbon Disclosure Project (CDP), aim to measure firms' impact on all stakeholders, under the principle of double-materiality. These four frameworks are the most widely used, but in many cases, it is up to firms to choose which framework(s) they want to follow.

Consistent with the need for better rules on climate disclosure, we find when analyzing the main topics underlying greenwashing-related articles, that the topic of disclosure is one of the main topics and has gained importance over time.

3.3 Actions and Regulations

To combat greenwashing, many countries have introduced or proposed new regulations that aim to ensure that environmental claims are based on scientific evidence and verified by independent third parties.

The U.K. was probably the first country to implement disclosure requirements: premium listed companies and regulated financial firms (banks, insurers and asset managers) have been required by the Financial Conduct Authority (FCA) to disclose under the TCFD framework since 2021. The U.K. government has also amended the Companies Act of 2006 to implement the TCFD recommendations for large and listed companies, since 2022.

In the E.U., In France, companies that have more than 500 employees and more than 100 million euros in turnover or balance sheet total have had to file a document known as Greenhouse Gas Emission Balance ("Bilan des Emissions de Gaz à Effet de Serre") every four year since 2010, as well as an annual Declaration of Extra-Financial Performance (DPEF) since 2017. These documents were partially in compliance with TCFD recommendations. An ordonnance was signed in December 2023 to strengthen these reporting requirements in compliance with the E.U. Taxonomy, the French Sustainability Reporting Standards (FSRS) and thereby the TCFD recommendations. In the US, the SEC amended in September 2023 the "Names Rule" initially adopted in 2001, which regulates the names of registered funds to ensure they are marketed to investors truthfully. The rule ensures that ESG funds have at least 80% of their assets be ESG assets, and improves transparency on the criteria chosen to select the assets.[3]

In the EU, a series of actions have been taken in the past few years to ensure better disclosure of firms. The Sustainable Finance Disclosure Regulation (SFDR) aimed to improve transparency in the market for sustainable investment products, by requiring investment firms to make detailed disclosures in relation to the products and services with environmental or social characteristics. Initial requirements became applicable in 2021. The Taxonomy Regulation that started coming into force in 2022 further required firms to disclose how aligned their products or services were with what is considered "green" under the taxonomy. Furthermore, a Regulation on European green bonds was proposed in 2021. It was adopted in October 2023, and started being applied in October 2024. On March 22, 2023, the European Commission published its proposed Directive on Green Claims.[4] The proposal is part of the EU's Green Deal, which aims to make Europe the first climate-neutral continent by 2050. The proposed Directive introduces new rules on how companies can substantiate their

[3] See, https://www.sec.gov/news/press-release/2023-188.

[4] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2023%3A0166%3AFIN

environmental claims. The proposal also requires that all environmental claims and labels be independently verified and certified before being publicized. In 2023, a new Corporate Sustainability Reporting Directive (CSRD) also entered into force. It introduced more detailed and standardized reporting requirements compared to the previous Non-Financial Reporting Directive (NFRD). It applies to all public interest entities operating in the E.U., such as listed companies, banks, and insurance companies. It covers around 49,000 companies and requires them to report on ESG topics using the European Sustainability Reporting Standards (ESRS), which are based on the TCFD recommendations. Moreover, in September 2023, the European Parliament and the European Council reached a provisional agreement on new rules to ban misleading advertisements and provide consumers with better product information.[5]

Finally, the European Commission requested in May 2022 the three European Supervisory Authorities (the European Banking Agency, the European Insurance and Occupational Pensions Authority and the European Securities and Markets Authority – ESAs), to provide input on greenwashing risks and occurrences in the EU financial sector and on the supervisory actions taken and challenges faced to address those risks. The final report was released on June 4, 2024, by the European Securities and Markets Authority (2023).

In the US, until recently, there was no specific federal law that regulated greenwashing. There were several laws that could be used to address it, such as the Federal Trade Commission Act (FTC Act), which prohibits unfair or deceptive acts or practices in commerce. The FTC issued guidelines for environmental marketing claims, known as the Green Guides, which provide examples of how to make truthful and substantiated claims about the environmental attributes of products or services. The FTC can take enforcement actions against companies that violate the FTC Act or the Green Guides, such as issuing cease and desist orders, imposing civil penalties, or requiring corrective advertising. In 2021, for example, the FTC fined Walmart and Kohl's, a combined $5.5m for mislabeling rayon as "sustainable" bamboo. Additionally, some states, e.g., California, have enacted their own laws to prevent greenwashing. Finally, in September 2023, the U.S. Securities and Exchange Commission (SEC) amended the "Names Rule" initially adopted in 2001, which regulates the names of registered funds to ensure they are marketed to investors truthfully. The rule ensures that ESG funds have at least 80% of their assets be ESG assets, and improves transparency on the criteria chosen to select the assets.[6]

[5]See https://www.europarl.europa.eu/news/en/press-room/20230918IPR05412/eu-to-ban-greenwashing-and-improve-consumer-information-on-product-durability for details.

[6]See, https://www.sec.gov/news/press-release/2023-188.

4 Measuring Greenwashing Salience

Greenwashing is by nature challenging to measure. We do not claim to be able to identify individual incidents of greenwashing. Our goal, instead, is to measure greenwashing salience. The hypothesis is that investors become aware of greenwashing through the news, and integrate this information in their investment behavior. In this section, we describe the data and the algorithm that we use to quantify greenwashing salience. We show that our algorithm provides a good trade-off between out-of-sample performance and parsimony.

4.1 Data

We use the history of daily paper-based Wall Street Journal articles from January 1986 to June 2022.[7] We filter out articles that belong to the opinion section of the Wall Street Journal as well as pieces that we deem irrelevant.[8] We also remove articles with less than 200 words or more than 2000 words. The number of remaining articles over time is depicted on Figure 1. In total, our database contains 860,427 articles.

4.2 What do we want to measure?

We proxy greenwashing salience as the volume of news that are related to greenwashing. Below are two typical examples of news that we classify as greenwashing-related:

> *2018-10-25 Business News: New York Sues Exxon Over Climate Change. "Exxon built a facade to deceive investors into believing that the company was managing the risks of climate-change regulation to its business when, in fact, it was intentionally and systematically underestimating or ignoring them, contrary to its public representations," New York Attorney General Barbara Underwood said in a statement.*

> *2021-09-27. In August, The Wall Street Journal reported that DWS struggled to define and implement an ESG strategy, at times painting a rosier-than-reality*

[7] Since 2000, the Wall Street Journal is also available online. We do not use the online articles as most of them are redundant with paper articles.

[8] The Opinion section of the Wall Street Journal contains the following subsections: editorials, letters to the editors, opinions and commentaries. We also filter out the following types of news: obituaries, reviews, advertorials, blogs, calendar of events, country profiles, headline listings, headline-only content, personal announcements, prospectus, rankings, contest and lottery, horoscopes, real estate listings, recipes, traffic and weather news.

picture to investors, according to Ms. Fixler and internal emails and presentations.

We slightly extend the definitions of greenwashing given in Section 3.1, to include articles that explicitely mention confusion, or deception, created around green products. This confusion may affect consumers and investors' trust in these products in the same way as traditional greenwashing. In particular, this extended definition classifies the following articles as greenwashing-related: 1) articles explicitly mentioning investors being confused about the risk-return profile of ESG assets because of asset managers' communications, 2) articles explicitly mentioning the confusion created by firms' ESG disclosures (or lack of), 3) articles explicitly mentioning the confusion around the definition of what is green or not (for example, the confusion on whether nuclear energy is green).

With this extended definition, the following text is considered greenwashing-related:

2018-09-18. The problem here isn't the ESG ratings, but that they are used as though they were some sort of objective truth. In reality they are no more than a series of judgments by the scoring companies about what matters – and investors who blindly follow their scores are buying into those opinions, mostly without even knowing what they are.

4.3 Algorithm

The challenge of dealing with our dataset is threefold: it is large, unlabelled, and most articles are neither climate risk-related nor greenwashing-related so it is highly unbalanced. Furthermore, there may be many articles mentioning the greenwashing without using the "greenwashing" word specifically.

We identify greenwashing-related articles in two steps. First, we identify articles that are related to climate risk, and filter out all the other articles. Second, we identify, among the climate risk-related articles, those that are related to greenwashing. The keyword and document discovery algorithm of King, Lam, and Roberts (2017) provides an ideal setup to overcome the challenges of dealing with our dataset. The method is a semi-supervised learning method. It starts from a small number of documents that are highly likely to be in a given category (for example, climate risk-related), and uses text analysis methods to learn the patterns that most characterize these documents. Based on this training, the algorithm then predicts which other documents are likely to be climate-related, enlarges the initial set with these articles and iterates. Sautner, van Lent, Vilkov, and Zhang (2023a) use this

algorithm at the level of the sentence to identify climate mentions in earning calls. We adjust it to our purpose and use it in the two steps of our method, to identify, first, climate risk-related articles, and second, greenwashing-related articles.

In order to fine-tune our algorithm's hyperparameters and to test the out-of-sample performance, we create validation and testing sets for each of the algorithm's two steps. We recruited around 40 research assistants, who each received the task to read 150 articles, and to evaluate whether these articles were climate risk- and greenwashing-related. A climate risk (greenwashing) -related document is defined as a document which contains at least one mention of climate risk (greenwashing). It needs not use the bigram "climate risk" ("greenwashing"), and can use other words deemed equivalent by the student.

As the review process leads itself to errors, each article was reviewed by two students. When the two students did not report the same classification, a third student was asked to read and classify the article, and to flag "questionable" articles, i.e., articles for which the classification was unclear. We reviewed independently ourselves a random selection of non-questionable articles, and all questionable ones.

A total of 1479 articles were assigned to students and read by at least two students. The selection process of these articles was designed to address the unbalanced number of climate risk-related articles, compared to non climate risk-related, as well as the increasing volume of articles over time. Details on this selection process are provided in Appendix A. Among these articles, 875 we identified as climate risk-related and 604 as non climate-related. We then sampled randomly half of the articles in each category (i.e., 437 climate risk-related articles and 302 non climate risk-related articles). They were assigned to the validation set, and the other half to the testing set. Similarly, the 875 climate risk-related articles were classified into 60 greenwashing-related articles and 815 non greenwashing-related articles. Half of these (i.e., 30 greenwashing-related articles and 407 non greenwashing-related articles) went into the validation set of the second step of our algorithm, and the other half into the testing set. A summary of the validation and testing sets is provided in Table 1.

We describe below each step of the algorithm in detail.

4.3.1 Step 1: Identification of climate risk-related articles

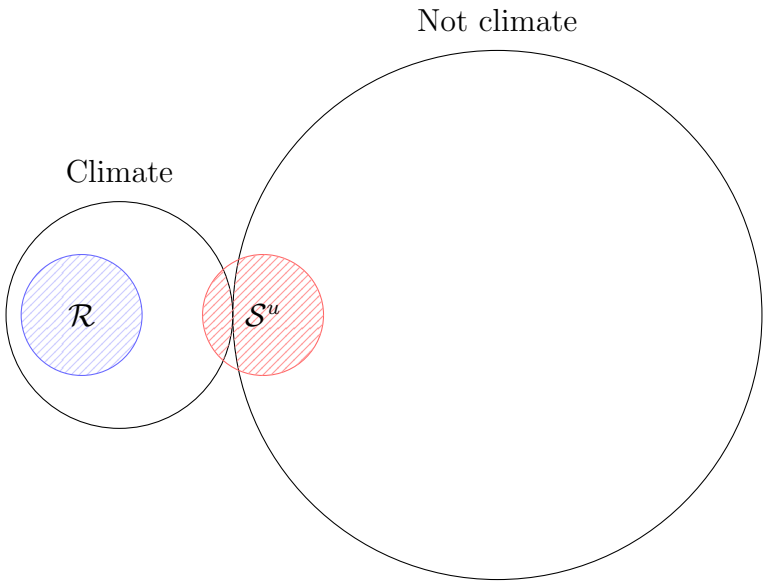
The algorithm of King, Lam, and Roberts (2017) is based on a small set of articles, namely the *reference set* $\mathcal{R}$, which are highly likely to belong to a given category. We use a dictionary-based method to build this reference set. We start by defining a short list of words C_0 that are unambiguously related to climate risk. We build this list from the dictionaries of Engle,

Giglio, Kelly, Lee, and Stroebel (2020) and Sautner, van Lent, Vilkov, and Zhang (2023a), removing the words whose meaning depends on the context. For example, the unigram "environment" could be used in the phrase "low interest-rate environment". Therefore we remove it from the list. We construct a preliminary reference set $\mathcal{R}^{\text{prelim}}$ with all the articles that contain at least one of these words.

To address the possibility that we have missed important unigrams and bigrams, we search for the n-grams that appear in $\mathcal{R}^{\text{prelim}}$, and try and detect those that discriminate this set most from the other articles. Specifically, we use the term frequency - inverse document frequency (TF-IDF) to identify the bigrams that appear frequently in $\mathcal{R}^{\text{prelim}}$ but less frequently in the rest of the articles. We list the n-grams with the highest TF-IDF, after validating their relevance. We finally rank the resulting bigrams based on how well they discriminate the reference set from the rest of the database. Specifically, we compute the likelihood metric suggested in Sautner, van Lent, Vilkov, and Zhang (2023a) for each bigram. The final list of bigrams C_S contains the bigrams with a top 5% likelihood. The final climate dictionary contains the initial dictionary C_0, enlarged with C_S. The word cloud in Figure 2 illustrates the resulting climate dictionary $C = C_0 \cup C_S$.

All articles in the database that contain more than three n-grams in C are added to the reference set $\mathcal{R}$. By definition, all of the articles in $\mathcal{R}$ are thus very likely to be climate risk-related. The set $\mathcal{R}$ contains 1631 articles.

Figure 1: Decomposition of the training set



The strength of the algorithm lies in that it then discovers new climate-related articles by

learning the keywords/patterns that distinguish the documents in the reference set $\mathcal{R}$ from the remaining articles of the training set (*search set* $\mathcal{S}$). We randomly select 1631 articles in $\mathcal{S}$, to create an under-sampled search set $\mathcal{S}^u$ of the same size as $\mathcal{R}$.[9] The decomposition of the training set is illustrated in Figure 1.

We train three classifiers to identify the differences between the articles in the reference set $\mathcal{R}$ and the articles in the under-sampled search set $\mathcal{S}^u$. This step relies on the assumption that most articles in $\mathcal{S}^u$ are not climate risk-related. Similar to Sautner, van Lent, Vilkov, and Zhang (2023a), we choose a random forest classifier, a support vector classifier and a Naive Bayes classifier.

These classifiers take as inputs a numerical representation for each word of the article to classify. There are several ways to encode text. The simplest way is a bag-of-words approach, which simply counts how many times each word appears in a document. This approach ignores the order of words, as well as their context, i.e., which words tend to be around one another. More involved approaches such as Word2Vec, are based on word vectors, which take into account some elements of context by modelling how likely words are to appear near one another. They encode text into numerical vectors such that words that are often used together will have similar representations. For example, the n-grams "greenwashing", "window dressing" and "misleading claims" will be encoded with vectors that are close to one another. Several recent papers have used such embeddings, e.g., Ash, Chen, and Ornaghi (2021); Cao, Kim, Wang, and Xiao (2020) and van Binsbergen, Bryzgalova, Mukhopadhyay, and Sharma (2023).

We average the embedding vectors provided by Word2Vec for each word of a document, to obtain the numerical representation of the entire document. We replace the Naive Bayes classifier by an eXtreme Gradient Boosting (XGBoost) classifier because the Naive Bayes algorithm can only take as input an integer representation of documents and Word2Vec provides a float representation instead. If several terms that are close to "greenwashing" appear in the document, this will be reflected in the average. We test both bag-of-words and Word2Vec methods and find that in this first step, the bag-of-words representation provides the best trade-off between simplicity and performance.

The classifiers' hyperparameters are fine-tuned by maximizing the F1-score over the validation set of manually labelled articles. The F1-score is the average of the precision and

[9] Using all the articles of the search set would result in an extremely unbalanced dataset, as the reference set is small compared to the search set. This issue is usually tackled using over-sampling of the minority class (the reference set) or under-sampling of the majority class (search set). Under-sampling methods are in our case better suited, as there are still climate-related articles in $\mathcal{S}$. We performed tests with respect to the number of articles contained in the under-sampled set $\mathcal{S}^u$ and chose $\mathcal{S}^u$ to be of the same size as $\mathcal{R}$.

recall, where the precision refers to the fraction of all the positive predicted (i.e. all articles classified as climate risk-related) that are true positives (i.e., rightly predicted), and the recall is the fraction of all positives (i.e., all climate risk-related articles) that are true positives (i.e., rightly predicted). Both performance measures should be as close to 1 as possible. A low precision means that many articles are classified as climate risk-related wrongly, that is, the volume of climate risk-related articles is over-estimated. A low recall means that few of the climate risk-related articles are classified as such, that is, the volume of climate risk-related articles is under-estimated. Optimizing the F1-score allows reaching a trade-off between these two situations.[10]

We then estimate the probability of each article in the full search set $\mathcal{S}$ to be climate risk-related. The intuition is that if an article of the search set is similar enough, in terms of vocabulary, to articles in the reference set, then it is assigned a high probability to be climate risk-related. We label an article as climate risk-related if two of the three classifiers assign this article to the climate-related class.

The procedure is iterated, taking the new total set of articles identified as climate-risk related as reference set $\mathcal{R}$, until convergence.

The output of this first step is a classification of all articles in the training set as climate risk-related or not. Over the 857,548 articles of the training set, 20,846 are classified as climate risk-related. Together with the manually labelled validation and test sets, we obtain in total 21,719 climate-related articles (2.5% of the database).

4.3.2 Step 2: Identification of greenwashing-related articles

In the second step, we automatically classify as non greenwashing-related all articles that are not classified as climate risk-related in the first step. The training set therefore contains the 20,846 articles that were classified as climate risk-related in the first step. It represents 96% of all the climate risk-related articles. The validation set (437 articles) and the testing set (438 articles) both account for 2% of the climate risk-related articles.

We apply a similar procedure as the one of the first step to identify which climate-related articles are greenwashing-related. We build the reference set by selecting all articles in the training set that contain the n-grams "greenwash" (or alternatively, "green-wash"), "green-

[10]The accuracy, which measures the fraction of predictions which are right, is often the criterium that is maximized when fine-tuning a classifier's hyperparameters. As our dataset is highly unbalanced, and contains much more non-climate risk-related articles than climate risk-related articles, a high accuracy can be achieved by classifying all articles as non-climate risk-related. When working with such an unbalanced dataset, maximizing the F1-score is therefore a better method.

washing" (or alternatively, "green-washing"), "window dressing", "deceptive advertising", "false advertising", "misleading advertising", "false claim" or "misleading claim". The reference set contains 119 articles, i.e., 0.6% of the total training set. We build the search set by randomly selecting 300 articles from the climate-related articles of the training set that are not in the reference set.

As in Step 1, we encode the text using both a bag-of-words approach and Word2Vec. Parameters are fine-tuned by maximizing the F0.5-score, which is a weighted average of the precision and recall with double weight on precision. This choice is made following the observation that greenwashing classifiers tend to either have a high precision or a high recall. To be conservative and avoid over-estimating the volume of greenwashing-related articles, we put more weight on precision than on recall.

We keep the same aggregation rule: we classify as greenwashing-related articles which are considered as such by at least two of the three classifiers.

The output of this second step is a classification of all articles as greenwashing-related or not. Over the 20,846 articles of the training set, 646 are classified as greenwashing-related. Together with the manually labelled validation and test sets, we obtain in total 706 greenwashing-related articles (3.3% of the climate-related articles).

4.4 Out-of-sample Performance

We evaluate the out-of-sample performance of each classification step using the manually labelled articles in the testing set. The first line of the two cells in Table 2 lists various measures of performance that each step of our algorithm achieves.

All measures of performance are strikingly high for the first step of the algorithm, which identifies climate risk-related articles. Above 90% of all articles are classified in the right category (accuracy). About 90% of all articles classified as climate risk-related are indeed in this category (precision). Over 95% of all climate risk-related articles are well categorized (recall). These two statistics yield an F1-score of 92.5%. The last statistic we display is the area under the Receiver Operating characteristic Curve, commonly called the Area Under the Curve (AUC), which measures the relation between the true positive rate (fraction of all climate risk-related articles that are well classified) and the false positive rate (fraction of non-climate risk-related articles that are classified as climate-related). A perfect classifier would reach a true positive rate of 1 and a false positive rate of 0, yielding an AUC of 1. Our classifier reaches an AUC of nearly 90%, confirming the high quality of the identification of climate risk-related articles.

Changing the numerical representation of words in articles from a simple bag-of-words approach to a Word2Vec approach (second line) does not substantially improve results. It allows reaching a slightly higher recall, at the expense of a slightly lower accuracy, precision, F1-score and AUC. Therefore we choose the simpler bag-of-words approach.

We compare the performance of our algorithm to the ones of Engle, Giglio, Kelly, Lee, and Stroebel (2020) and Sautner, van Lent, Vilkov, and Zhang (2023a). More details on our implementation of these methods are provided in Appendix B. We are able to classify more than 10% more articles in the right category than these two papers, which in our tests achieve comparable statistics. In particular, our classification of climate-related articles is more precise (over 10% gain). Our recall is also slightly higher, and our AUC more than 10% higher.

The statistics of the second step of the algorithm are displayed in the bottom half of the table. Our method displays an extremely high accuracy of 98%, which means that 98% of the articles are classified in the right category. As greenwashing-related articles only represent a small fraction of all articles[11], this number would however likely be achieved by classifying all articles as non-greenwashing-related. It is therefore reassuring to see that the precision is nearly 90%, meaning that very few articles classified as greenwashing-related are wrongly classified. The recall is however only of 56.7%, meaning that out of all greenwashing-related articles, our classifier is missing 40% of them. This statistic shows that identifying greenwashing-related articles is much more challenging that identifying climate risk-related articles.

These statistics are obtained using Word2Vec as embedding method. Using a simple bag-of-words approach indeed a significantly lower recall of 33.3%. Word2Vec therefore allows better capturing greenwashing-related concepts, so that less articles are missed. This finding is interesting to contrast to the first step of the algorithm, in which a bag-of-words approach yielded comparable results to Word2Vec. This difference between the two steps sheds some light on why identifying greenwashing-related articles is more difficult than identifying climate risk-related artices. Indeed, whereas a fair number of bigrams are specific to climate risk (e.g., global warming, greenhouse gas, Kyoto protocole etc.), few words and bigrams are specific to greenwashing. Accounting for the context of words therefore brings value to the second step while it does not to the first step.

[11] The statistics of the second step are displayed as functions of all articles, whether they are climate-related or not.

5 Greenwashing Index

Using the classification algorithm described in Section 4, we identify the articles of the Wall Street Journal that are climate- and greenwashing-related. In this section, we analyze the evolution of climate risk and greenwashing, and describe the resulting greenwashing index. We investigate which firms, industries and topics are most related to greenwashing.

5.1 Climate Risk

Figure 3 displays the percentage of climate risk-related articles and the percentage of greenwashing-related articles, out of all articles. Our measure of climate risk (blue line) has been increasing since the end of the nineties and exhibits peaks around climate-related events, such as the UN Climate Change conferences and the Paris Agreement. Importantly, the peaks seem to be contemporaneous with crises, with a first peak in mid-2008 (around the UN Climate Change conference, COP12 in Bali), and a second peak around the time of the UN Climate Change Conference in Glasgow (COP26), end of 2021. Climate risk-related articles then reached nearly 10% of all articles published in the Wall Street Journal. The years from 2010 to 2020 show lower values climate risk, even though these values are still higher than the pre-2000 values.

The overall shape of our measure of climate risk is in line with the indices obtained by Engle, Giglio, Kelly, Lee, and Stroebel (2020); Ardia, Bluteau, Boudt, and Inghelbrecht (2022) and Faccini, Matin, and Skiadopoulos (2023). Table 3 displays the correlation matrix of these monthly indices, and reveals that our measure has a correlation of above 50% with all other indices. These high correlations were not obvious, given the various methods (databases, algorithms and time periods) used to build the other indices. Indeed, none of the other indices has a correlation above 50% with all indices. This indicates that our measure captures well the common component of climate risk.

Figure 3 shows that greenwashing-related articles represent a small fraction, on average 0.07%, of all articles. Although this fraction is small, it has been steadily increasing since mid-2018. In particular, it exhibited several peaks close to 1% in 2021 and 2022, and an average value of 0.5% between 2018 and 2022. In comparison, there is on average 2.4% of Wall Street Journal articles that are climate risk-related. In the past year, greenwashing has therefore grown to a magnitude of the same order as that of climate risk. Furthermore, greenwashing and climate risks are highly correlated, with a 66% correlation.

5.2 Index Description

Figure 4 displays our greenwashing index, built as the percentage of climate risk-related articles (instead of the percentage of all articles as above) that are greenwashing-related over time. This index is designed to measure how much greenwashing contaminates the discussion on climate risk. The correlation between the climate index and this greenwashing index is of 42%. This correlation drops to 38% if we clean the climate index by removing from it all articles that are greenwashing-related. All correlation values therefore indicate that there is a strong comovement between climate risk and greenwashing.

Greenwashing-related articles represent between 0 and 5% of climate risk-related articles, with occasional peaks, until 2018. These peaks coincide with various companies being accused of greenwashing, and sometimes charged by the SEC. They are clustered in two waves. The first wave is from 1990 to 1992, as several companies came under scrutiny. Exxon was accused of trying to look greener than it actually was following the 1989 Exxon-Valdez oil spill in Alaska. Procter & Gamble was criticized, first because of claims on disposable diapers, and second for lobbying against environmental measures while showing off their green credentials. Burger King was accused of falsely claiming that their packages were degradable. Toyota and General Motors were criticized for releasing ads which implied that their cars were protecting the environment. In 1991, greenwashing-related articles reach 10% of all climate risk-related articles. Following these incidents, the FTC started developing the Green Guides, issued in 1992. These guidelines provided recommendations and best practices for marketers to avoid making misleading or unsubstantiated environmental claims.

The following years gave less importance to greenwashing until BP was charged with misleading claims following a pipeline leak in 2002. But it is only from 2006 that greenwashing increased again, in a second wave. In particular, it exhibits peaks when Exxon Mobil was accused of funding climate denier groups in the end of 2006, and when computer makers Dell and HP came under fire because of exaggerated claims that their computers were energy-savers, in 2007.

It is only around 2015 that the greenwashing index increased again, briefly triggered by the Volkswagen emissions scandal, which shook the credibility of the car industry. From 2018, a last greenwashing wave started, of larger amplitude than the two others. It coincided with greenwashing accusations towards investment firms (e.g., Pax World Management, Greenlight Capital, Blackrock, Deutsche Bank) and repeted criticisms of the metrics used to measure their green credentials. Since 2018, the greenwashing index has regularly been above 8%. In the past five years, greenwashing has therefore contaminated substantially the discussion on climate risk, in a way that is unprecedented.

5.3 Greenwashing-related firms and industries

Greenwashing salience has recently been increasing. But which firms and industries are most linked to this increase? To answer this question, we extract the names of companies that are mentioned in the greenwashing-related articles, and map them to the 17 Fama-French industries. The main challenge of this exercise is that the same company can be known using several aliases, e.g., Proctor and Gamble can also be referred to as P&G, or Procter & Gamble Co; similarly, Apple and Apple Inc. refer to the same company.

We address this challenge using a method that proceeds in two steps. The first step performs Named Entity Recognition: it identifies all the company names present in the greenwashing-related articles. The second step perfoms Entity Linking: it links the different ways to refer to the same company.

For now, we restrict our analysis to companies currently traded on the NYSE, the NASDAQ and the AMEX. The list of these companies, as well as all the different names that can be used to refer to them, is obtained using Wikidata. For each company, Wikidata contains a list of aliases, e.g., of names commonly used. These aliases can be collected using a Wikidata SQL query.[12]

This procedure allows us to map each mention of a company to the main name of that company, its ticker and the stock exchange on which it is traded. The SIC code of the company is then recovered from CRSP. The allocation of company mentions to industry is finally performed by mapping SIC codes to Fama-French industries.[13]

These two steps allow us to count how many times each company is mentioned in the set of greenwashing-related article. Figure 5 shows in a word cloud the companies that are mentioned. The size of each company's name increases with the number of mentions. Blackrock clearly appears as the company that has been mentioned most, together with many other companies of the financial sector, e.g., JPMorgan Chase, Goldman Sachs, Bank of America, Deutsche Bank, Citigroup etc. Companies in the oil & gas industry are also prominent, including ExxonMobil and Chevron.

We map companies to the 17 Fama-French industries using their SIC codes. Figure 6 shows the evolution of the most prominent industries over time. Here again, it appears that the financial sector (banks, insurance companies and other financials) are in large part driving the recent increase in greenwashing. In contrast, the shares of other industries have

[12]Wikidata SQL queries can be performed on https://query.wikidata.org/. The detail of the query is provided in Appendix C.

[13]The mapping from SIC codes to Fama-French industries is obtained from the webpage of Wenzhi Ding, https://wenzhi-ding.com/pages/project.html

remained comparable over time.

5.4 Greenwashing Topics

As the definition in Section 3.1 indicates, companies can greenwash in many different ways. In order to better understand the nature of greenwashing over time, we perform topic analysis. We use an extension of the standard Latent Dirichlet Allocation (LDA) method introduced by Blei, Ng, and Jordan (2003), which allows, through human validation, to overcome some well-known issues of the LDA.

The LDA method is a probabilistic model, i.e., it posits a set of latent topics, multinomial distributions over words, and assumes that each document can be described as a mixture of these topics. It is an unsupervised learning method, which only uses unlabelled data. While it has been successfully used as a tool to explore the topics of a corpus of documents, it has also been shown to often inadvertently create multiple topics with similar content and combine different themes into a single topic (Chang, Boyd-Graber, Gerrish, Wang, and Blei (2009); Newman, Bonilla, and Buntine (2011); Morstatter and Liu (2016)). The keyword Assisted Topic Model (keyATM) introduced by Eshima, Imai, and Sasaki (2023) is a semi-supervised topic model[14], which incorporates prior information the researcher has on the existing topics. This prior information has the form of lists of keywords for the topics of interest. The method can be iteratively applied, by inputting at each iteration the topics that were well identified at the previous iteration. Eshima, Imai, and Sasaki (2023) show that this prior information improves the performance of the topic model and better serves the purpose of measurement of topic importance.

We first estimate the standard LDA. The twenty most important topics are displayed in Appendix D. They reveal many topics that are difficult to interpret and highly redundant. In particular, the keyword "fund" is in nineteen out of the twenty topics. Therefore we apply the keyATM procedure, starting with one topic defined by the keyword "fund". After several iterations, we obtain a list of twelve topics as displayed in Table 4, and for each of them, the ten most relevant keywords. The names of the topics were chosen based on the list of keywords given as output of the keyATM methodology.

The most important topic is "Disclosure". This comes to no surprise as it has been shown to be a major driver of greenwashing, as detailed in Section 3.2. The topic is characterized by keywords such as "report" and "data". The second most important topic is the one that we labelled "Asset management" (with keywords related to investments, e.g., "fund", "asset",

[14]The keyATM model is an open-source package available at https://keyatm.github.io/keyATM/.

"ETF" and "portfolio"). There is no surprise here either, given the prominence of the word "fund" in the initial results of the LDA. A reading of the articles most representative of this topic reveals that ESG funds' stock-picking methods are described as a source of skepticism for investors.

Half of the topics are specific to three industries: the financial, the energy and the construction industries. They represent, respectively, 17.3%, 12.4% and 2.2% of the text in greenwashing-related articles. One of the strengths of the keyATM method is that it also provides as output a dynamic representation of the topics' importance over time. Figure 7 describes the evolution of these three sector-based clusters of topics and reveals that the Financial topics have grown in two steps. Before 2000, their associated topic proportion was essentially zero. They grew to 5 to 10% between 2000 and 2010, peaked at 40% in 2017, and remained above 20% until the end of the time series. In contrast, the Energy topics represented about 10% of the text until 2009, but then slightly decreased in importance. The construction topic was important starting in 1998, for about a decade.

There are three Financial topics: "Asset management" (with a focus on funds), "ESG ratings" (with keywords including "ESG", "score" and "Morningstar"), and "Green bonds" (with keywords such as "credit, debt" and "green" and "project"). Greenwashing concerns about asset management mostly relates to the difficulty to know which assets are really green. Some funds are self-labelled as green (or equivalent terms), but as we will see in Section 6.2, some of these funds have received poor Morningstar globe ratings, indicating that their holdings are not aligned with Morningstar's sustainability criteria. The confusion created by green bonds is mostly due to the lack of standards in the definition of green projects and the absence of measures of their impact. Finally, ESG ratings are often described as misleading because of the low correlation between one another. Figure 8, Panel A, shows that there three topics have followed the same patterns and become more prominent together in the last ten years. This topic analysis therefore confirms our previous results, i.e., that the financial sector is driving the increase in greenwashing in the recent years. ESG funds, ESG ratings and green bonds are all responsible for this increase.

Two topics are part of the Energy sector: "Alternative energies" (with keywords such as "clean", "technology" and "solar"), and "Fossil fuels" (with keywords such as "car", "gas" and "coal"). Figure 8, Panel B, shows the initial skepticism associated to alternative energies until the mid-nineties: their topic proportion was much larger than the one of fossil fuels. Since then, it appears that greenwashing is associated both to alternative energies and to fossil fuels, in comparable magnitudes. The scale of the figure however indicates that these topics have had, in the recent years, a small weight in the greenwashing index, of the order

of a quarter the weight of financial topics.

Three topics are cross-sectors: "Emissions" (with keywords including "carbon", "target" and "greenhouse"), "Labels", (with keywords such as "marketing" and "brand"), and "Recycling" (with keywords including "plastic", "packaging" and "waste"). Reading through the main articles on emissions highlights the skepticism linked to emission measurements, and the lack of trust of these measurements. Figure 9, Panel A, shows that this concern has been growing over the years, to reach nearly 10% of topic importance in the past three years. In contrast, labels were an important subject of concern in the nineties, due to the multiplication of green labels, to their opacity and lack of regulation. In the last ten years however, this concern faded, with a topic importance that decreased from 10-15% (depending on the year) to less than 5%. Similarly, the "Recycling" topic was predominant in the nineties, but has in the past ten years almost disappeared from the greenwashing index, with a topic importance of 1 to 2%. Only one cross-sector topic therefore remains of importance: the "Emissions" topic.

Finally, three topics are related to firms' public relations. First, the "Disclosure" topic, as mentioned above, is the most important greenwashing topic, with an overall topic importance of 12.4%. Figure 9, Panel B shows that this topic has been of importance throughout the time period we analyze, with a tendency to slightly grow over time. Such tendency indicates the shortcomings of the many disclosure frameworks proposed to address this concern, as detailed in Section 3.2. Related to this, the second public relations topic that has also been growing over time is a topic on "Shareholder activism" (with keywords such as "corporate", "board", "proposal" and "vote"). This result is consistent with the increasing number of shareholder proposals pressuring companies to their sustainable transition, and more transparent in their disclosures, Fisch (2022). Finally, we find a "Law suits" topic (with keywords including "claim", "case" and "federal"), whose importance has been decreasing over time. This result may reflect the shortcomings of laws in addressing companies' actions and disclosures. The increasing climate risk has pushed companies to find new ways to look green. Whereas law suits may have been sufficient, in the nineties, to counterbalance incentives to greenwash, this may not be the case anymore.

To sum up, our analysis underlines two main conclusions. First, the sector that has drawn most concern from market participants, in recent years, is the financial sector. This sector includes ESG funds, green bonds, and ESG ratings. Second, the main object of concern is disclosure.

6 Impact of Greenwashing on Financial Markets

6.1 Does Greenwashing Fuel Anti-ESG Sentiment?

The media have uniformly reported an ESG backlash over the past few years. For example, in its December 4, 2023 article entitled "The real impact of the ESG backlash", the Financial Times gives examples of companies, such as Blackrock, which changed their communications over the years as a response to the anti-ESG sentiment. The article mentions greenwashing as a source of concern, indicating that it might be partly responsible for the ESG backlash.[15] In line with this anectotal evidence, Section 5.4 revealed that the financial sector was a key driver of the greenwashing index in the past few years. Specifically, the Asset Management topic was the second most prominent topic, with key words linked to ESG funds and ESG ratings. In this section, we study whether increases in the greenwashing index are indeed followed by increases in anti-ESG sentiment.

We build an anti-ESG sentiment index from the history of opinions of the Wall Street Journal. Since 1986, the Wall Street Journal has published nearly 80,000 opinion pieces. These pieces were filtered out in the construction of the index as by nature, they reflect sentiment rather than facts, and are targeting the general population rather than investors. Using the iterative algorithm described in Section 4.3, we classify these opinions into climate-related or not. The number of climate-related opinions increases over time, with 3,708 opinions since 1986, among which 3,381 were published after 2000.

We associate a tonality to each opinion piece. To do so, we identify the sentences that contain climate-related words, and measure the tonality of these sentences. Climate-related words were obtained from our database of climate risk-related WSJ articles, as the list of words with a TF-IDF above a chosen threshold. The threshold was chosen by hand, to only keep words that were unambiguously related to climate. Tonality words were obtained from the Loughran-MacDonald dictionary, Loughran and McDonald (2011). Each sentence containing at least one climate-related word was classified negative if it contained more negative than positive words. Opinion pieces were finally classified as negative if they contained more negatively toned than positively-toned sentences. Finally, the anti-ESG sentiment is built as the fraction of climate-related opinions that have negative tonality. The index, aggregated at the quarterly frequency, is displayed in Figure 10, starting in 2005.[16] The index exhibits a clear and steady upward trend. In 2005, around 5% of climate-related opinion pieces were negatively toned where in the past few years, this number increased to nearly 15%, on aver-

[15] See https://www.ft.com/content/a76c7feb-7fa5-43d6-8e20-b4e4967991e7.

[16] Before 2005, most months did not include negatively-toned opinion pieces.

age. In line with the ESG backlash described in the media, there is a decrease in the trend between 2018 and 2021, followed by a sudden rebound and a large peak in the index. This peak is however not the first one, and is similar to the one that started around 2005 and ended at the end of the financial crisis.

We study the impact of greenwashing on the anti-ESG sentiment using monthly indices. At the monthly level, the anti-ESG sentiment index exhibits little persistence, with an autocorrelation of 0.17. It is positively correlated with the greenwashing index, with a correlation coefficient of 0.10. To fully remove the persistence of these time series, we transform them into unexpected shocks by taking the innovations of AR(1) models applied to each index. The two resulting time series of shocks are almost perfectly uncorrelated, with a correlation of 0.03. We denote greenwashing shocks by GS_t, and anti-ESG sentiment shocks by $AntiESG_t$. We regress innovations in the anti-ESG sentiment at time t on past values of greenwashing innovations:

$$AntiESG_t = a_0 + \sum_{j=1}^{3} b_j GS_{t-j} + \epsilon_t. \tag{1}$$

The results of the regression are displayed in Table 5. The coefficients b_j are positive, and for most of them statistically significant.

We conclude that the anti-ESG sentiment is, as suspected, fueled by an increase in greenwashing salience.

6.2 Does Greenwashing Affect Investors' Behaviors?

In this section we study whether the general anti-ESG sentiment fueled by greenwashing translates into changes in how investors trade funds. The question of the impact of greenwashing salience on fund flows is relevant as flows have been shown to have a strong effect on returns, Van der Beck (2023).

We consider the universe of Morningstar open-end funds and ETFs domiciled in the United States from August 2018 until June 2022.[17] We focus on the evolution of their monthly flows. Flows in dollars are provided across share classes. The different share classes refer to the same portfolio, offered under different terms to different types of investors. Some of the share classes are flagged as targeting retail investors, the others are offered to

[17]Morningstar updated their globe rating methodology in 2018. Data on funds' globe ratings before this methodology change are not available. Our analysis of fund flows starts in August 2018 as it strongly relies on the availability of these ratings.

institutional investors. On average, each fund has 2.7 share classes.

Following Chevalier and Ellison (1997) and Sirri and Tufano (1998), net flows into share classes are defined as the net growth rate in total net assets (TNA) that is not due to dividends and capital gains on the assets under management. These net flows are then normalized by share class size as in Hartzmark and Sussman (2019):

$$\text{Share class flows}_{i,t} = \frac{A_t^i - A_{t-1}^i(1 + r_t^i)}{A_{t-1}^i} \tag{2}$$

where A_t^i are the share class's total net assets and r_t^i is the monthly return between t - 1 and t as reported on CRSP.

We filter out funds with size less than a million dollars. Both relative flows and share class sizes are skewed to the right. To avoid outliers, we winsorise them at the 99% level. We also winsorise relative flows at the 1% level. We further filter out funds which did not receive Morningstar sustainability and star ratings. Table 6 summarizes the sample. The final dataset contains 7,613 funds and 20,480 share classes.

Each share class in our sample has two ratings, updated by Morningstar on a monthly basis: a globe and a star rating. Globe ratings assess the sustainability of a fund. An aggregated score is calculated for each fund by taking a weighted average of the ESG ratings of their holdings, obtained from Sustainalytics. Funds are then ranked based on these scores. They are given five globes (four, three, two and one respectively) if they are in the top 10% of funds within their category (between 10% and 32.5%, between 32.5% and 67.5%, between 67.5% and 90%, and above 90%). Star ratings are allocated at the share class level, based on the past risk-adjusted returns of the portfolio. Star ratings are thus a backward-looking measure of performance.

To study how greenwashing affects fund flows, we focus on share classes that are either sustainable, or performant, or both. Share classes that have the highest sustainability rating (5G), but not the highest star rating (5S), are denoted $5G \backslash 5S$. Similarly, share classes that have the highest star rating but not the highest sustainability rating are denoted by $5S \backslash 5G$. Finally, share classes that are both performant and sustainable are in $5S \cap 5G$. If greenwashing becomes more salient, investors may invest less in funds that are sustainable but not performant (the former share classes), and more into funds that are performant, but not necessarily sustainable (the latter share classes).

Furthermore, some funds have made their sustainability more salient through the choice of their name, which contains an ESG-related word. We identify these funds to study whether the reaction of their flows to increases in greenwashing salience is different than

for other funds. To do so, we make a list of words that are commonly used in fund names and are linked to sustainability. This list includes words such as "environment", "green", "sustainable", "ESG", "clean", and "renewable". We find 392 funds with an ESG-related word. There is no one-to-one relation between having a 5-globe rating and an ESG name: nearly 40% of the funds advertised as sustainable have a time-average globe rating below 4, and about 10% of the funds that are not advertised as green have a time-average rating above 4.[18]

We run panel regressions of share class flows, as defined in equation (2), on lagged greenwashing shocks, crossed with dummy variables. We indicate whether share classes belong to one of the groups described above using the dummy variables $D_{5G\setminus 5S}$, $D_{5S\setminus 5G}$, $D_{5G\cap 5S}$ and Advertised green. The general regression is the following:

$$\begin{aligned}\text{Fund flows}_{i,t} = &\sum_{h=1}^{3} a_h GS_{t-h} \times D_{5G\setminus 5S,i,t} + \sum_{h=1}^{3} b_h GS_{t-h} \times D_{5S\setminus 5G,i,t} \\ &+ \sum_{h=1}^{3} c_h GS_{t-h} \times D_{5S\cap 5G,i,t} + \sum_{h=1}^{3} d_h GS_{t-h} \times D_{5S\cap 5G,i,t} \times \text{Advertised green}_i \\ &+ \text{Fund fixed effects}_i + \text{Month fixed effects}_t + \text{Controls}_{i,t-1}. \qquad (3)\end{aligned}$$

We use the following controls: flows over the past year, past returns over 1 and 2 years, lagged returns, lagged size and age of the share class.

Solomon et al. (2014) show that the reaction of fund flows to news depends on investors' sophistication. They find that investors that are less sophisticated react more to news than other investors. Therefore we run the regressions for all share classes, irrespective of which type of investor they target, and then separately for share classes targeting retail and institutional investors.

We first run regressions with time fixed effects, excluding the trend variables from the controls. Time fixed effects are designed to capture general trends that hold across funds. Such trend can be due to macroeconomic events, or systemic events such as the Covid pandemic. Results are displayed in Table 7, in the first three columns. The first column shows the results for all investors, the second column only for retail investors and the third one only for institutional investors. They reveal several interesting facts.

First, irrespective of the greenwashing index, the star rating has a strong impact on

[18]The Names Rule, which helps ensure that a fund's name accurately reflects the fund's holdings, was amended in September 2023 by the SEC. It used to apply only to fund names that suggested a focus on a particular type of investment, industry or geographic region, or those that suggested certain tax treatment. Its scope has now been extended to funds with "particular characteristics", including ESG terms.

flows. Share classes that are 5S attract more capital, whether they are sustainable or not. The effect is statistically significant and holds across retail and institutional share classes. Second, the effect more than doubles if the fund is advertised as green, and targetting institutional investors. The increase is small and insignificant for retail investors. These two results shed light on what investors generally pay attention to: (i) Performance is an important determinant of flows for all types of investors, and (i) A green name matters for institutional investors, much less for retail investors. Finally, investors value star ratings more than they value globe ratings. Share classes of funds rated 5G but not 5S do not receive more inflows than other funds. This result is in line with the findings of Gantchev et al. (2024).

We next run regressions with fund-fixed effects, including the greenwashing index as trend. Fund fixed effects capture the popularity of some funds, which may benefit from better marketing. The results are in the fourth to sixth columns. They indicate that increases in greenwashing salience are followed by retail investors investing less into all funds, whether they are rated and labelled sustainable or not. This slowdown in inflows occurs with a lag, it is statistically significant three months after the news. The results are not significant for institutional investors.

Finally, we run regressions with time- and fund-fixed effects. The results of these regressions are in columns seven to nine. They reveal a different reaction of retail and institutional investors to increases in greenwashing salience. While retail investors do not target specific funds, and reduce their inflows in all funds, institutional investors invest less in funds that are 5S, 5G and self-labelled as green. This effect is strong, statistically significant, and persists three months after the increase in greenwashing salience.

We conclude that greenwashing salience has a significant impact on fund flows, which however differs for retail and institutional investors. Our results indicate that following an increase in greenwashing salience, retail investors lose trust in all funds, and decrease their inflows irrespective of the rating and name of funds. In contrast, institutional investors target funds that have a high star and globe rating, and self-label as green. These funds are presumably the most expensive funds. The impact of greenwashing on flows is important as fund flows have been shown to have a strong effect on returns, Van der Beck (2023). As past returns also have a strong effect on flows, both flows and returns can be affected through a spiral effect. More broadly, if our results extend to ESG investing as a whole, greenwashing can impair the positive effects of ESG investing documented in Lakkis (2022).

References

Amenc, N., F. Goltz, and V. Liu (2021). Doing good or feeling good? detecting greenwashing in climate investing. Technical report. EDHEC Research Chair Publication.

Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2022). Climate change concerns and the performance of green vs, brown stocks. *Management Science*, 1–26.

Ash, E., D. L. Chen, and A. Ornaghi (2021). Gender attitudes in the judiciary: Evidence from u.s. circuit courts. QAPEC Discussion Papers 08.

Atta-Darkua, V., S. Glossner, P. Krueger, and P. Matos (2023). Decarbonizing institutional investor portfolios: Helping to green the planet or just greening your portfolio? Working Paper.

Bazillier, R. and J. Vauday (2009). The Greenwashing Machine : is CSR more than Communication. Working Paper.

Bingler, J., M. Kraus, M. Leippold, and N. Webersinke (2022a). Cheap talk and cherry-picking: What climatebert has to say on corporate climate risk disclosures. *Finance Research Letters 47*.

Bingler, J., M. Kraus, M. Leippold, and N. Webersinke (2022b). How cheap talk in climate disclosures relates to climate initiatives, corporate emissions, and reputation risk. SFI Working Paper.

Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent dirichlet allocation. *Journal of Machine Learning Research 3*, 993–1022.

Brandon, R. G., S. Glossner, P. Krueger, P. Matos, and T. Steffen (2022). Do responsible investors invest responsibly? *Review of Finance 26*, 1389–1432.

Cao, S., Y. Kim, A. Wang, and H. Xiao (2020). From Words to Syntax: Identifying Context-specific Information in Textual Analysis. Working paper.

Cartellier, F., P. Tankov, and O. Zerbib (2024). Can Investors Curb Greenwashing? Working paper.

Chang, J., J. Boyd-Graber, S. Gerrish, C. Wang, and D. M. Blei (2009). Reading tea leaves: How humans interpret topic models. In *Neural Information Processing Systems*, pp. 288–296.

Chevalier, J. and G. Ellison (1997). Risk taking by mutual funds as a response to incentives. *Journal of Political Economy 105*(6), 1167–1200.

Dahl, R. (2010). Green washing: Do you know what you're buying? *Environmental Health Perspectives 118*(6).

de Freitas Netto, S. V., M. F. Falcão Sobral, A. R. Bezerra Ribeiro, and G. R. da Luz Soares (2020). Concepts and forms of greenwashing: a systematic review. *Environmental Sciences Europe 32*.

Deegan, C. (2002). Introduction: The legitimising effect of social and environmental disclosures – a theoretical foundation. *Accounting, auditing, & accountability 15*, 282–311.

Delmas, M. A. and V. C. Burbano (2011). The drivers of greenwashing. *California Management Review 54*(1).

Dumitrescu, A., J. Gil-Bazo, and F. Zhou (2023). Defining greenwashing. Working Paper.

Engle, R. F., S. Giglio, B. Kelly, H. Lee, and J. Stroebel (2020). Hedging climate change news. *Review of Financial Studies 33*(3), 1184–1216.

Eshima, S., K. Imai, and T. Sasaki (2023). Keyword assisted topic models. *American Journal of Political Science*.

European Securities and Markets Authority (2023, May). Progress Report on Greenwashing. Response to the European Commission's request for input on "greenwashing risks and the supervision of sustainable finance policies". Technical report.

Faccini, R., R. Matin, and G. Skiadopoulos (2023). Dissecting Climate Risks: Are they Reflected in Stock Prices? Forthcoming: Journal of Banking and Finance.

Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics 116*(1), 1–22.

Fisch, J. E. (2022). Purpose proposals. *University of Chicago Business Law Review*, 113–155.

Flagstad, I., Å. L. Hauge, and S. Å. K. Johnsen (2022). Certification dissonance: Contradictions between environmental values and certification scheme requirements in small-scale companies. *Journal of Cleaner Production 358*.

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

Furlow, N. E. (2022). Greenwashing in the new millenium. *Journal of Applied Business & Economics 10*, 22–25.

Gantchev, N., M. Giannetti, and R. Li (2024). Sustainability or performance? ratings and fund managers' incentives. Working Paper.

Gibson, D. (2009). Awash in green: A critical perspective on environmental advertising. *Tulane Environmental Law Journal 22*(2), 423–440.

Hartzmark, S. M. and A. B. Sussman (2019). Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows. *Journal of Finance 74*(6), 2789–2837.

Heath, D., D. Macciocchi, R. Michaely, and M. C. Ringgenberg (2023). Does Socially Responsible Investing Change Firm Behavior? *Review of Finance 27*(6), 2057–2083.

Ilhan, E., P. Krueger, Z. Sautner, and L. T. Starks (2023). Climate risk disclosure and institutional investors. *Review of Financial Studies 36*, 2617–2650.

Inderst, R. and M. Opp (2024). Greenwashing and Retail investors: The case for a Taxonomy for ESG Investments?*. Working Paper.

Kacperczyk, M. T. and J.-L. Peydró (2022). Carbon emissions and the bank-lending channel. Working Paper.

Kim, E.-H. and T. P. Lyon (2015). Greenwash vs. Brownwash: Exaggeration and Undue Modesty in Corporate Sustainability Disclosure. *Organization Science 26*(3), 705–723.

King, G., P. Lam, and M. E. Roberts (2017). Computer-assisted keyword and document set discovery from unstructured text. *American Journal of Political Science 61*(4), 971–988.

Lakkis, E. (2022). Real effects of esg investing. Working Paper.

Lane, E. L. (2010). Consumer protection in the eco-mark era: A preliminary survey and assessment of anti-greenwashing activity and eco-mark enforcement. *Marshall Rev. Intell. Prop. L. 742*.

Laufer, W. S. (2003). Social Accountability and Corporate Greenwashing. *Journal of Business Ethics 43*(3), 253–261.

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

Lyon, T. P. and J. W. Maxwell (2011). Greenwash: Corporate Environmental Disclosure under Threat of Audit. *Journal of Economics and Management Strategy 20*(1), 3–41.

Lyon, T. P. and A. W. Montgomery (2015). The Means and End of Greenwash. *Organization & Environment 28*(2), 223–249.

Marquis, C., M. W. Toffel, and Y. Zhou (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science 27*(2), 483–504.

Mitchell, L. and W. Ramey (2011). Look how green i am! an individual-level explanation for greenwashing. *Journal of Applied Business and Economics 12*(6), 40–45.

Morstatter, F. and H. Liu (2016). A novel measure for coherence in statistical topic models. In *ACL Anthology*, pp. 543–548.

Newman, D., E. V. Bonilla, and W. Buntine (2011). Improving topic coherence with regularized topic models. In *Neural Information Processing Systems*, pp. 496–504.

Parise, G. and M. Rubin (2023). Green window dressing.

Pizzetti, M., L. Gatti, and P. Seele (2021). Firms Talk, Suppliers Walk: Analyzing the Locus of Greenwashing in the Blame Game and Introducing 'Vicarious Greenwashing'. *Journal of Business Ethics 170*, 21–38.

Reuters (2023). Fund groups warn over eu legal definition of 'greenwashing'. 16 January.

Sautner, Z., L. van Lent, G. Vilkov, and R. Zhang (2023a). Firm-Level Climate Change Exposure. *Journal of Finance 78*, 1449–1498.

Sautner, Z., L. van Lent, G. Vilkov, and R. Zhang (2023b). Pricing Climate Change Exposure. *Management Science*, 1–22.

Sirri, E. R. and P. Tufano (1998). Costly search and mutual fund flows. *Journal of Finance 5*(53), 1589–1622.

Solomon, D. H., E. F. Soltes, and D. Sosyura (2014). Winners in the Spotlight: Media Coverage of Fund Holdings as a Driver of Flows. *Journal of Financial Economics 113*(1), 53–72.

Szabo, S. and J. Webster (2021). Perceived Greenwashing: The Effects of Green Marketing on Environmental and Product Perceptions. *Journal of Business Ethics 171*, 719–739.

van Binsbergen, J. H., S. Bryzgalova, M. Mukhopadhyay, and V. Sharma (2023). (almost) 200 years of news-based economic sentiment. Working Paper.

Van der Beck, P. (2023). Flow-driven esg returns. SFI Working Paper.

Van Tulder, R. and A. Van der Zwart (2006). *International Business-Society Management: Linking Corporate Responsibility and Globalization*. Routledge.

Vos, J. (2009). Actions speak louder than words: Greenwashing in corporate america. *Notre Dame Journal of Law, Ethics and Public Policy 23*.

Figure 1: Evolution of the number of articles

This figure displays the daily number of paper-based Wall Street Journal articles from January 1986 to June 2022. In total, the database contains 860,427 articles.

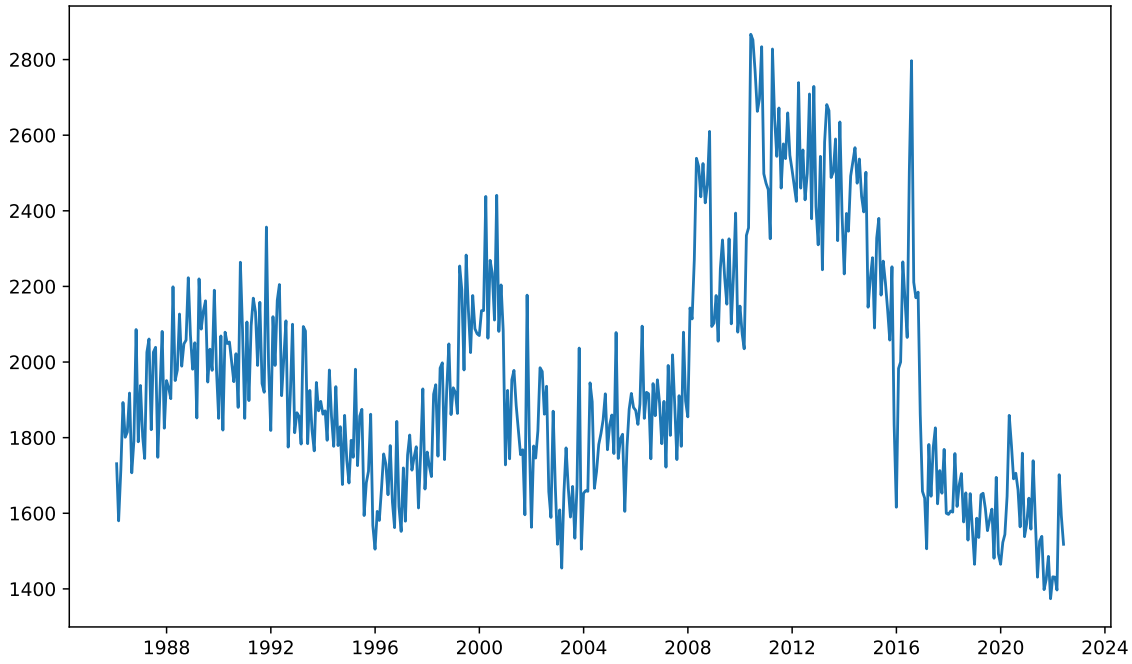


Figure 2: Word Clouds

This figure displays the word cloud of n-grams used in our climate dictionary $C = C_0 \cup C_S$. The construction of C_0 and C_S are described in Section 4.3. The size of each n-gram depends on its term frequency.

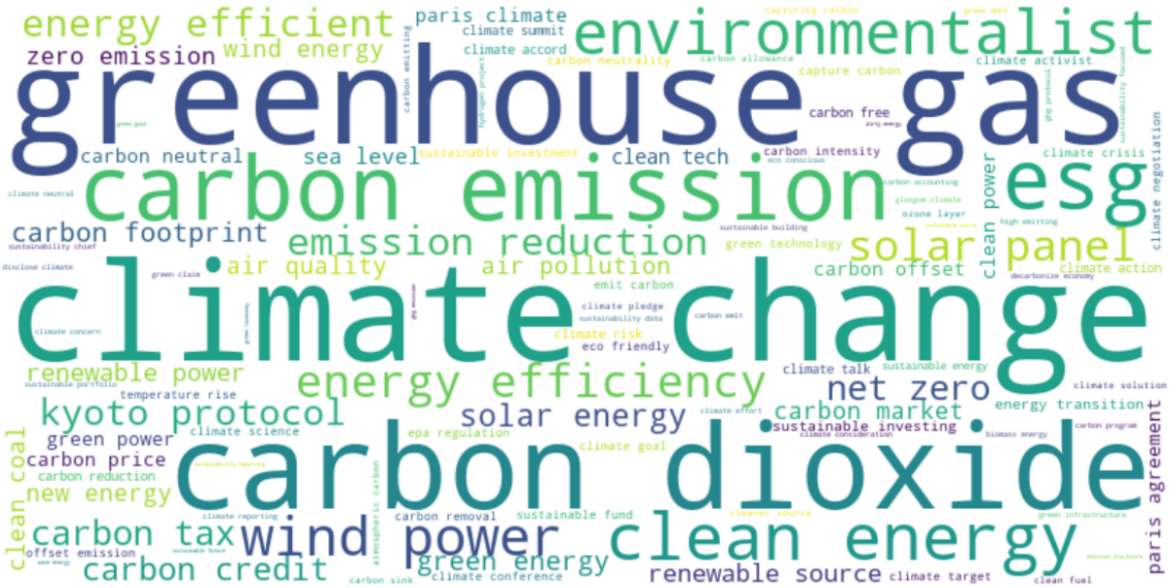


Figure 3: Climate Risk Index

This figure displays, in blue, our monthly index of climate risk, and in red, our monthly index of greenwashing. Indices are built following the procedure outlined in Section 4.3. Both indices are volume indices, obtained by counting the number of climate risk-related articles (respectively, greenwashing-related articles) each month, as fractions of the total number of Wall Street Journal articles. Annotations, in green, mark climate-related events.

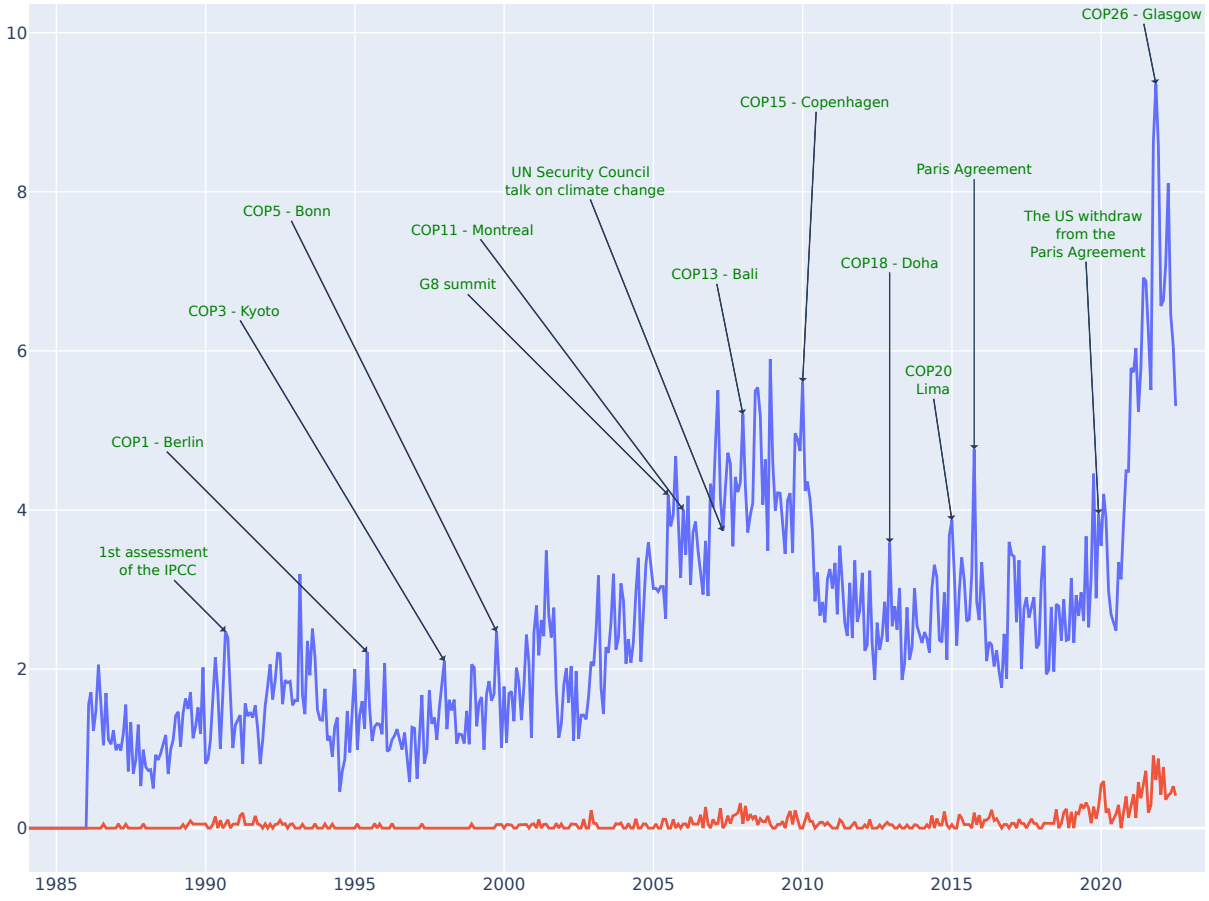


Figure 4: Greenwashing Index

This figure displays the quarterly greenwashing index, built following the procedure outlined in Section 4.3. This index is obtained by counting the number of greenwashing-related articles each quarter, as a fraction of the number of climate risk-related articles. Annotations, in green, mark accusations or charges of greenwashing that were made against given companies.

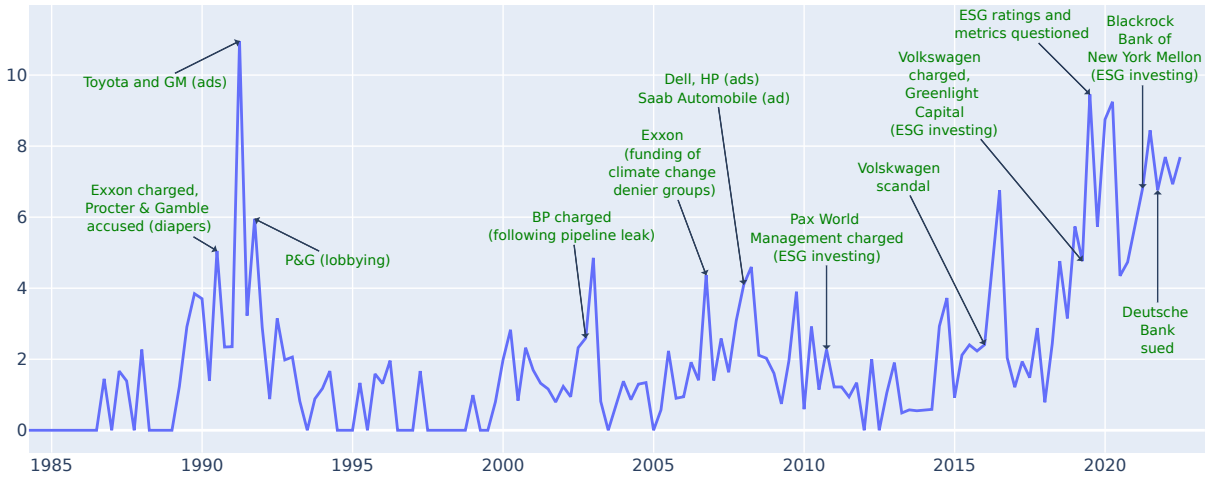


Figure 5: Companies Mentioned in Greenwashing-Related Articles

This figure shows the companies that are mentioned in the set of greenwashing-related articles. The method is restricted to companies traded on the NYSE, the NASDAQ or the AMEX. The size of words depends on the number of times companies are mentioned.

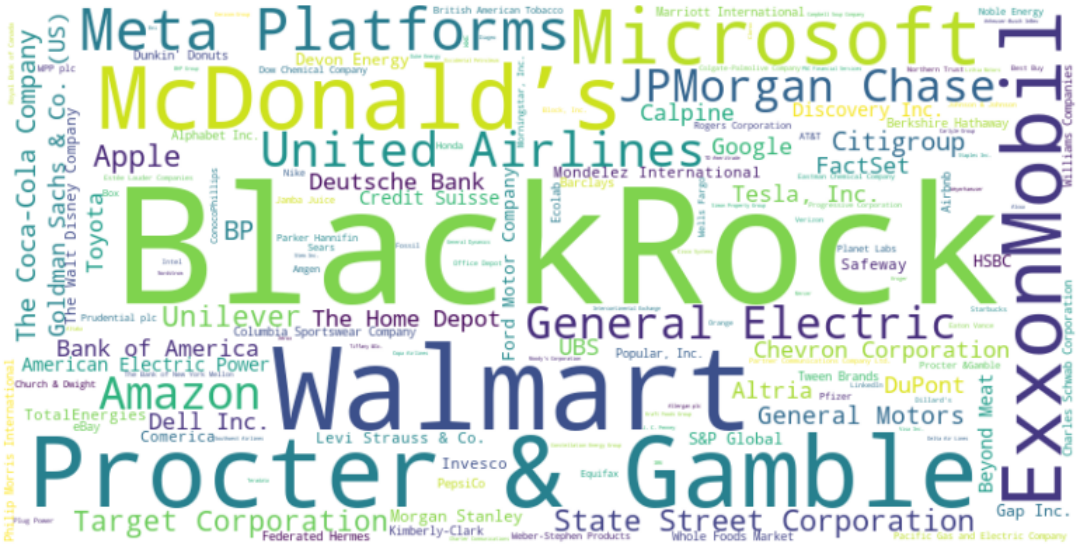


Figure 6: Industries Mentioned in Greenwashing-Related Articles

This figure shows the industries of mentioned companies over time, using the Fama-French classification of SIC codes into 17 industries. The method is restricted to companies traded on the NYSE, the NASDAQ or the AMEX.

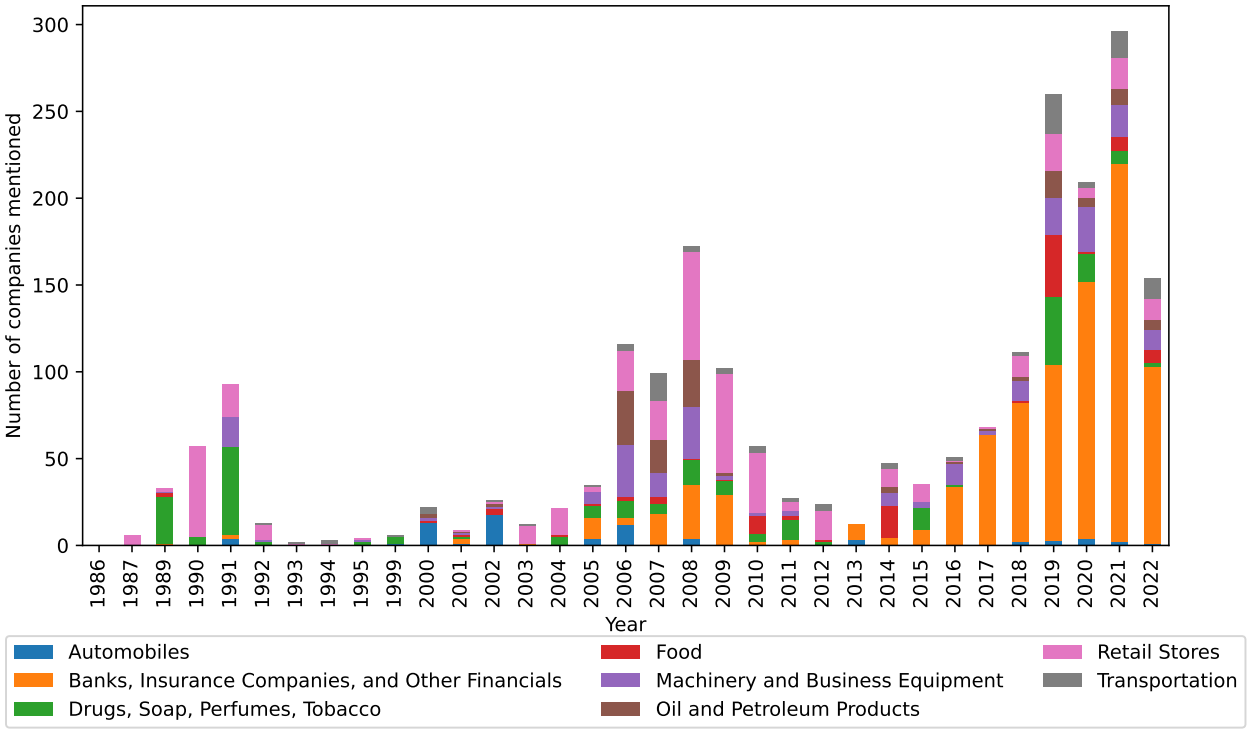


Figure 7: Topic Analysis: Industries

This figure shows the yearly topic proportion estimated by the keyATM algorithm described in Section 5.4, for three industry topics: the Financial, Energy and Construction sectors.

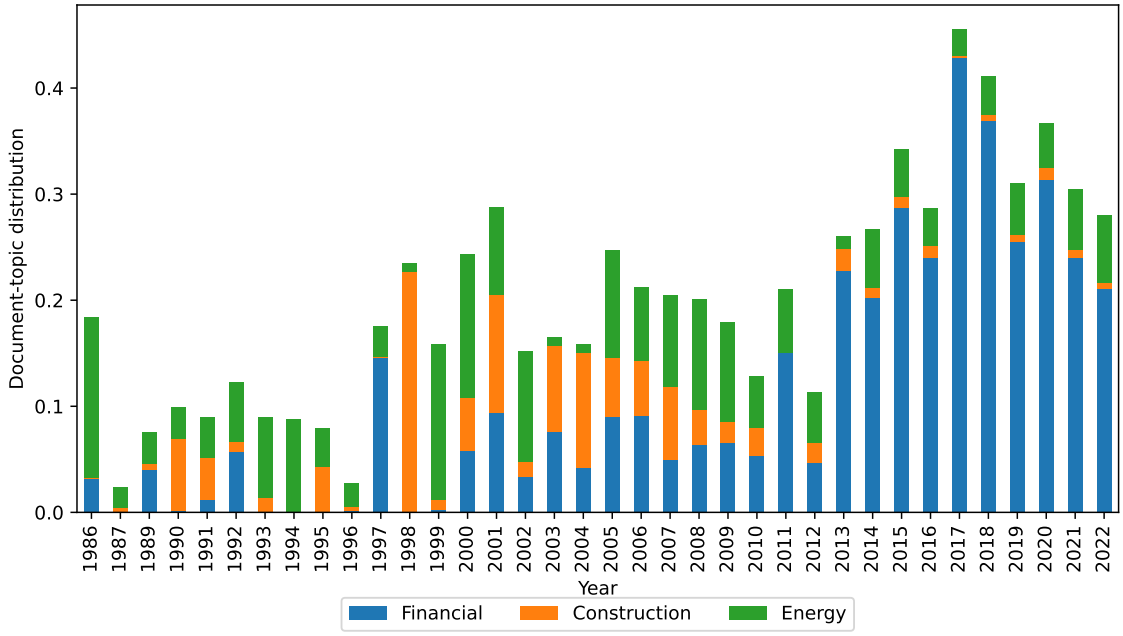
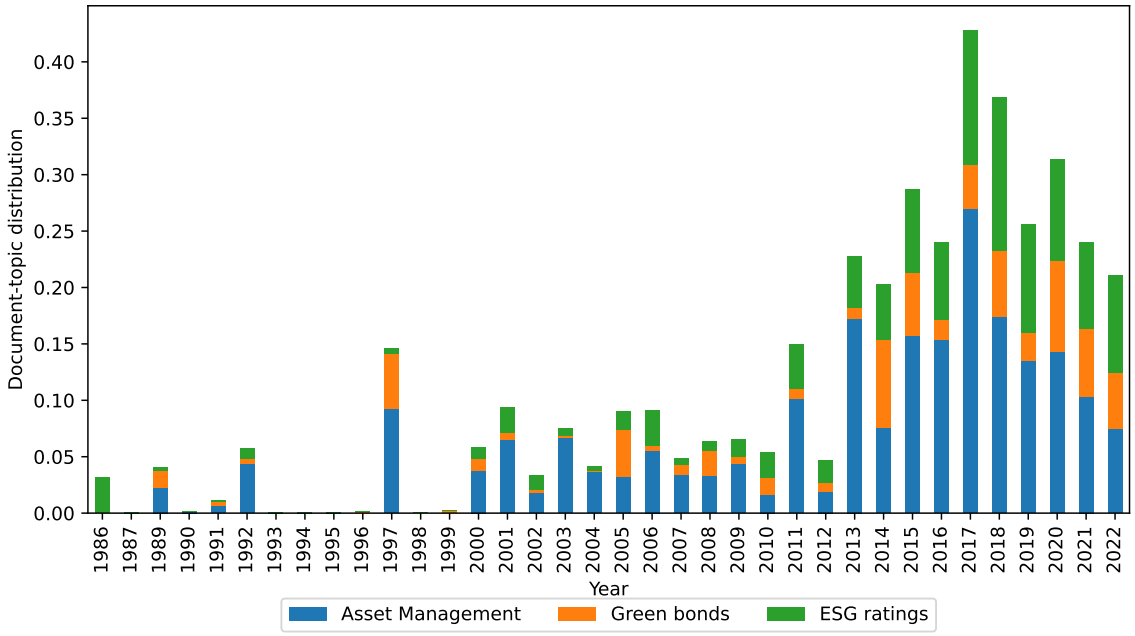


Figure 8: Topic Analysis: Sectors

This figure shows the yearly topic proportion estimated by the keyATM algorithm described in Section 5.4, for the three topics related to the Financial sector: Asset Management, ESG ratings and Green bonds (Panel A), and for the two topics related to the Energy sector: Fossil fuels and Alternative energies (Panel B).

Panel A: Financial sector



Panel B: Energy sector

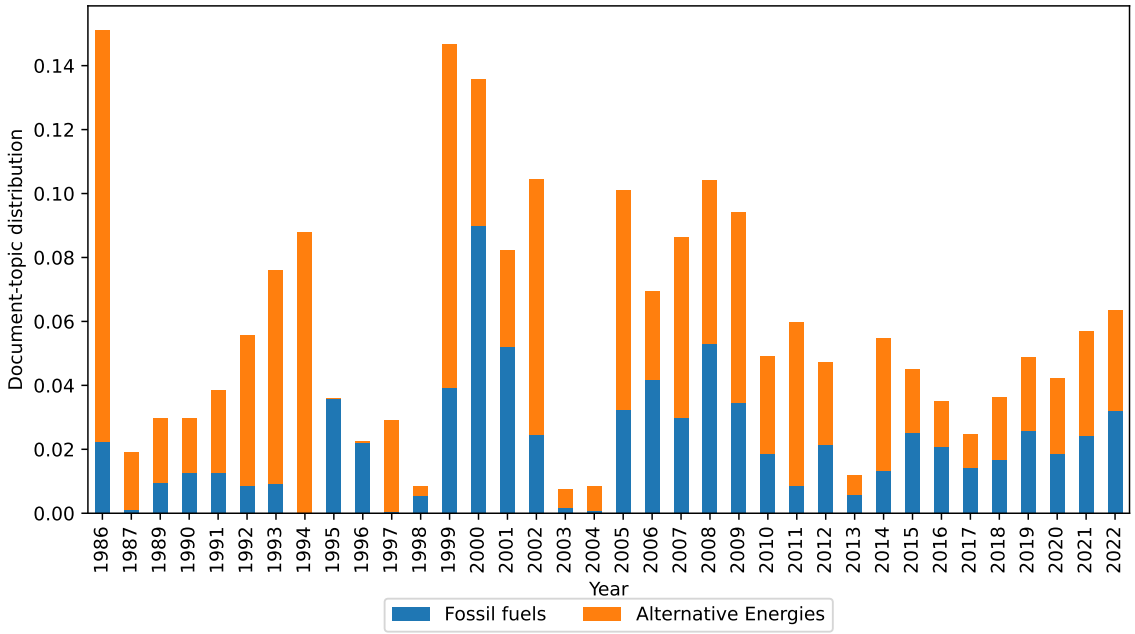
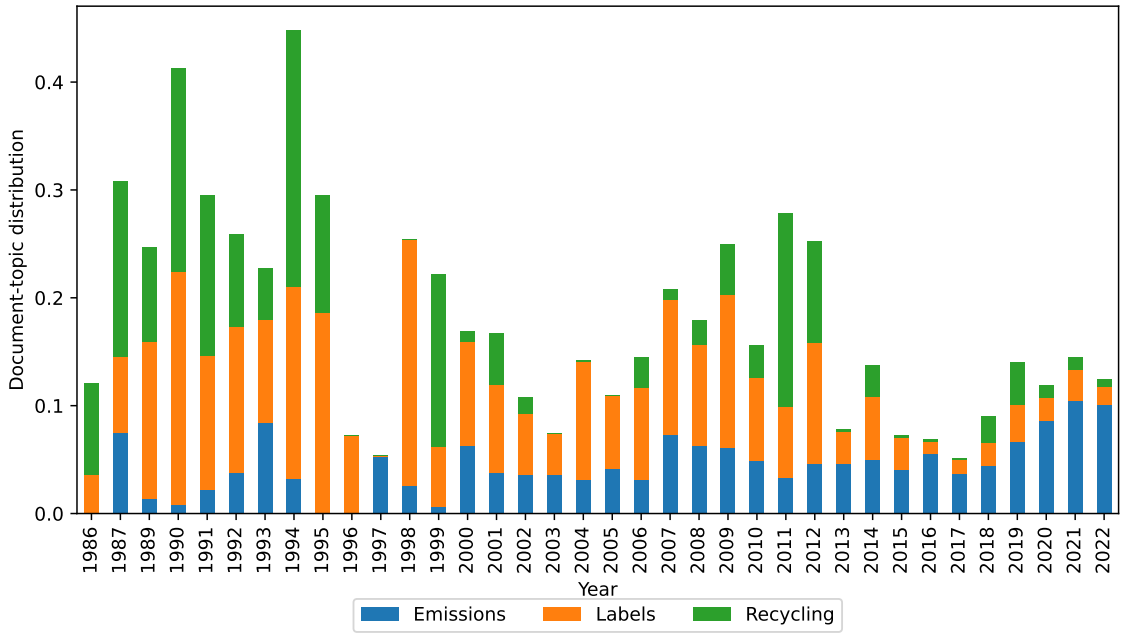


Figure 9: Topic Analysis: Transversal Topics

This figure shows the components of greenwashing that are not specific to a particular sector, and the evolution of their relative importance over time. The three topics displayed in Panel A are related to companies' green actions: Emissions, Labels and Recycling. The three topics displayed in Panel B are related to companies' public relations: Disclosure, Shareholder activism and Law suits.

Panel A: Companies' actions



Panel B: Companies' public relations

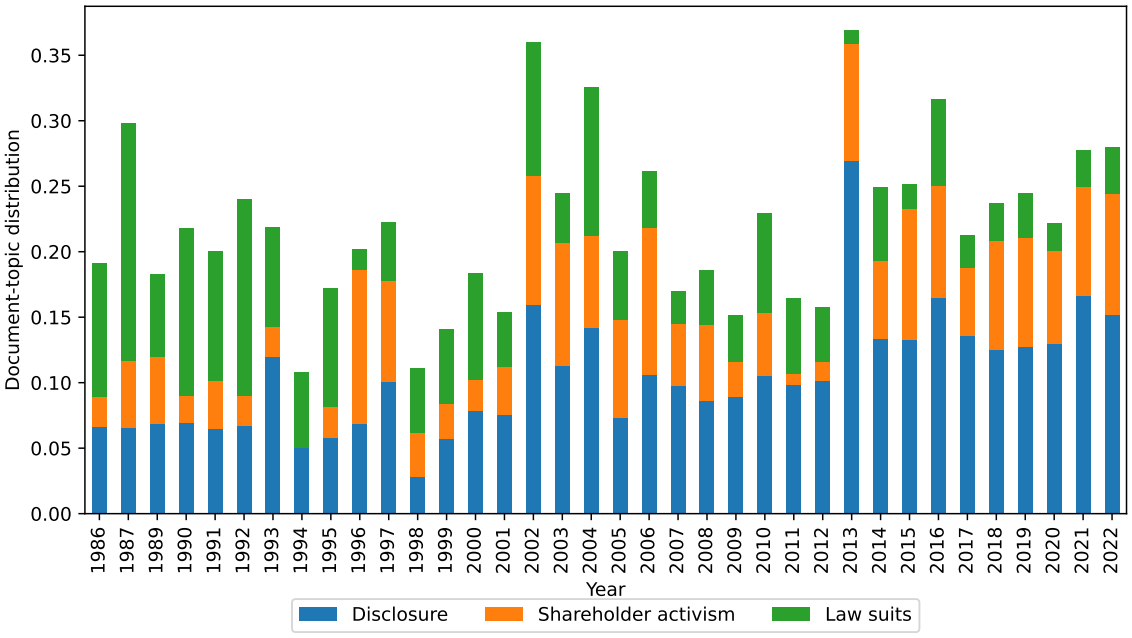


Figure 10: Anti-ESG Sentiment Index

This figure displays the quarterly Anti-ESG index, built following the procedure outlined in Section 6.1. This index is obtained by counting the fraction of climate-related opinion pieces each quarter that are negatively toned. The dashed line is the 1-year moving average of the anti-ESG sentiment.

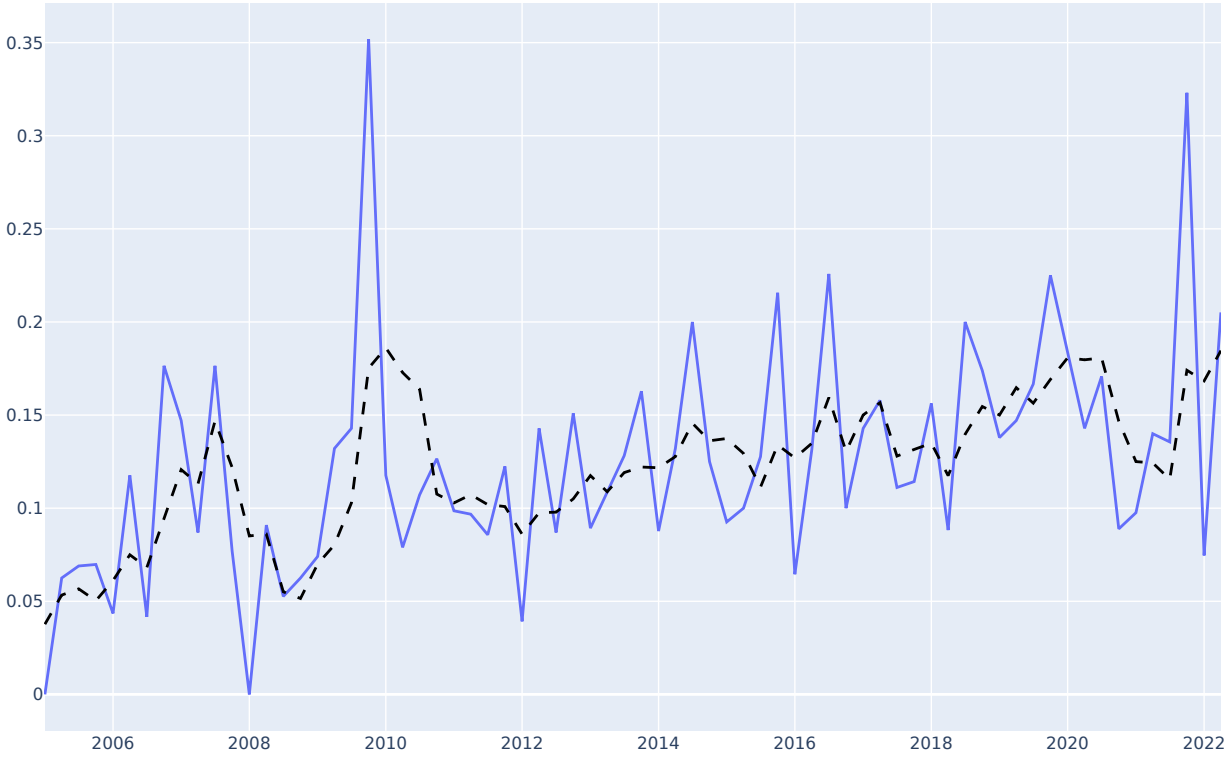


Table 1: Description of Sets Used in the Classification

This table summarizes the training, validation and testing sets used in the two steps of our classification algorithm. The column named '1. Climate' contains the volume of articles used in each set in the first step (identification of climate risk-related articles), and the column named '2. Greenwashing' contains the numbers corresponding to the second step (identification of greenwashing-related articles). The classification algorithm is described in Section 4.3.

Step		1. Climate		2. Greenwashing	
Training set					
	# articles	857,548		20,846	
	Min year	1986		1986	
	Max year	2021		2021	
Validation set		1	0	1	0
	# articles	437	302	30	407
	Min year	1986	1986	1990	1986
	Max year	2021	2021	2021	2021
Testing set		1	0	1	0
	# articles	438	302	30	408
	Min year	1986	1986	1987	1986
	Max year	2021	2021	2021	2021

Table 2: Performance Metrics of Classification Algorithms

This table summarizes the performance of the three types of algorithms tested for the classification of climate-related articles (step 1) and greenwashing-related articles (step 2). The steps of these algorithms are described in Section 4.3.

Step		Accuracy	Precision	Recall	F-score	AUC
1. Climate	**This paper, BoW**	**90.8**	**89.7**	**95.4**	**92.5**	**89.8**
	This paper, Word2Vec	89.1	86.8	96.1	91.2	87.5
	Engle et al. (2020)	79.6	77.2	92.9	84.3	76.7
	Sautner et al. (2023a)	79.7	77.5	92.5	84.3	76.9
2. Greenwashing	**This paper, Word2Vec**	**97.4**	**73.9**	**56.7**	**69.7**	**77.9**
	This paper, BoW	96.8	71.4	33.3	58.1	66.4
	Engle et al. (2020)	95.5	42.9	30.0	39.5	64.2
	Sautner et al. (2023a)	95.4	40.9	30.0	38.1	64.1

Table 3: Correlation between our Climate index and Alternative Indices

This table displays the correlation matrix between our monthly climate risk index, and alternative indices. The alternative indices considered are the monthly indices built by Ardia, Bluteau, Boudt, and Inghelbrecht (2022) (ABBI), Engle, Giglio, Kelly, Lee, and Stroebel (2020) (EGKLS) and the four indices of Faccini, Matin, and Skiadopoulos (2023): the U.S. climate policy (FMS-policy), the International summits (FMS-summits), the Global warming (FMS-glowarm) and the Natural disasters (FMS-natdisast) index.

	This paper	ABBI	EGKLS	FMS-policy	FMS-summits	FMS-glowarm	FMS-natdisast
This paper	1						
ABBI	**0.51**	1					
EGKLS	**0.64**	0.58	1				
FMS-policy	**0.64**	0.36	0.46	1			
FMS-summits	**0.66**	0.00	0.34	0.55	1		
FMS-glowarm	**0.58**	0.41	0.56	0.57	0.70	1	
FMS-natdisast	**0.55**	0.28	0.40	0.47	0.61	0.72	1

Table 4: Topic Analysis

This table indicates the words identified as the most representative of each topic identified by the keyATM method, described in Section 5.4. The topic are ordered by topic proportion: the most important topic comes first.

Topic	%	Most characteristic words
Disclosure	12.4%	sustainability, group, firm, report, investor, data, plan, impact, disclosure, financial
Asset Management	8.4%	fund, investment, investor, asset, etf, investing, manager, portfolio, sustainable, management
Shareholder activism	6.8%	corporate, shareholder, public, executive, issue, board, proposal, vote, social, activist
Emissions	6.5%	carbon, climate, emission, change, global, world, business, target, gas, greenhouse
Labels	6.1%	product, consumer, green, store, marketing, sale, label, environmentally, food, brand
ESG Ratings	5.4%	esg, stock, index, firm, risk, investor, score, rating, governance, morningstar
Law suits	4.5%	claim, state, general, official, case, action, federal, law, spokesman, agency
Green bonds	3.5%	bond, green, bank, project, finance, market, credit, investor, debt, issuer
Alternative energies	3.3%	energy, power, price, plant, clean, market, technology, solar, alternative
Recycling	3.3%	plastic, bag, paper, recycling, recycled, packaging, bottle, waste, diaper cup
Fossil fuels	2.6%	industry, oil, fuel, exxon, car, gas, fossil, coal, shell, global
Construction	2.2%	group, program, building, real, energy, forest, certification, standard, council, green

Table 5: Impact of Greenwashing on Anti-ESG Sentiment

This table reports the results of regressions of shocks in the anti-ESG sentiment index on lagged shocks in the greenwashing index, following equation (1). The regression was performed on time series starting in 2000 in column (1) and in 2010 in column (2). Newey-West standard errors are reported in parentheses, with 12 lags. $*$, $**$ and $***$ indicate statistical significance at the 1%, 5%, and 10% level, respectively. The number of observations reported corresponds to the number of months in the sample.

	Since 2005	Since 2010
GS_{t-1}	0.299	0.429*
	(0.248)	(0.236)
GS_{t-2}	0.649***	0.578***
	(0.171)	(0.205)
GS_{t-3}	0.560**	0.513*
	(0.244)	(0.282)
Observations	266	150
R^2	0.036	0.061

Table 6: Summary Statistics

This table presents summary statistics for the funds obtained from Morningstar, from August 2018 to June 2022. Fund size is given in million dollars. Flows are relative to the size of funds. Panel A summarizes the data after filtering out funds with size less than a million dollars. Panel B summarizes the data after winsorization, and filtering out funds that did not receive a Morningstar rating.

Panel A: Raw dataset

	Min	Max	Mean	Median	Std. dev.
Flows	-7.52	4130.47	0.05	0.00	6.48
Size	1	1373298	2931	355	17587

Panel B: Final dataset

	Min	Max	Mean	Median	Std. dev.
Flows	-0.19	0.51	0.01	0.00	0.08
Size	1	43730	2301	346	6225

Table 8: Correlations Between Factors

This table reports the correlation between our greenwashing factor, our climate risk factor and the other factors of the Fama and French (2015) model. The greenwashing and climate factors are built by taking innovations to AR(1) processes fitted to the respective indices.

	Greenwashing (GS)	Climate risk (CS)
Greenwashing (GS)	1	-0.04
Climate risk (CS)	-0.040	1
Market	-0.094	0.012
Size (SMB)	-0.076	0.005
Value (HML)	-0.024	-0.026
Profitability (RMW)	-0.112	-0.008
Investment (CMA)	0.082	0.008
Momentum (Mom)	-0.038	0.040

Table 7: Panel Regressions

This table reports the results of the panel regression of fund flows (relative to size) on lagged shocks of the greenwashing and climate risk indices, crossed with dummies that are 1 if the fund is advertised as ESG ('Ad. Green'), and if the fund has received a 5-globe Morningstar rating ('5G'), or a 5-star rating ('5S').

	Time Fixed Effects			Fund Fixed Effect			Time and Fund Fixed Effects		
	All	Retail	Institutional	All	Retail	Institutional	All	Retail	Institutional
$D_{5G\setminus 5S}$	0.001	0.001	0.001	-0.003***	-0.002**	-0.003**	-0.003***	-0.002**	-0.003**
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$D_{5S\setminus 5G}$	0.017***	0.015***	0.020***	0.012***	0.011***	0.014***	0.012***	0.010***	0.013***
	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$D_{5G\cap 5S}$	0.020***	0.020***	0.020***	0.010***	0.012***	0.008***	0.009***	0.011***	0.006**
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)
Ad. green × $D_{5G\cap 5S}$	0.021***	0.010	0.027***	0.019***	0.004	0.028***	0.021***	0.006	0.030***
	(0.004)	(0.006)	(0.005)	(0.004)	(0.006)	(0.005)	(0.004)	(0.006)	(0.005)
GS_{t-1}				0.002	0.004	-0.002			
				(0.012)	(0.012)	(0.016)			
GS_{t-2}				-0.004	-0.007	0.003			
				(0.010)	(0.011)	(0.012)			
GS_{t-3}				-0.026**	-0.028**	-0.019			
				(0.012)	(0.013)	(0.014)			
GS_{t-1}× $D_{5S\setminus 5G}$	-0.011	-0.009	-0.016	-0.006	-0.007	-0.007	-0.003	-0.004	-0.002
	(0.021)	(0.017)	(0.032)	(0.014)	(0.012)	(0.024)	(0.013)	(0.011)	(0.023)
GS_{t-1}× $D_{5G\setminus 5S}$	-0.011	-0.004	-0.028	-0.012	-0.006	-0.027	-0.010	-0.003	-0.026
	(0.013)	(0.012)	(0.024)	(0.010)	(0.010)	(0.020)	(0.011)	(0.010)	(0.020)
GS_{t-1}× $D_{5G\cap 5S}$	0.024	0.035	0.005	0.028	0.040	0.007	0.030	0.041	0.009
	(0.025)	(0.027)	(0.037)	(0.025)	(0.029)	(0.033)	(0.024)	(0.028)	(0.033)
GS_{t-1}× Ad. green × D_{5G5S}	-0.157***	-0.029	-0.238***	-0.134***	-0.000	-0.204***	-0.134***	0.002	-0.204***
	(0.044)	(0.070)	(0.067)	(0.038)	(0.065)	(0.060)	(0.041)	(0.066)	(0.064)
GS_{t-2}× $D_{5S\setminus 5G}$	-0.051**	-0.041*	-0.070*	-0.039***	-0.027**	-0.062***	-0.037***	-0.025**	-0.058***
	(0.026)	(0.021)	(0.036)	(0.015)	(0.013)	(0.021)	(0.013)	(0.013)	(0.020)
GS_{t-2}× $D_{5G\setminus 5S}$	-0.008	-0.006	-0.010	0.002	0.006	-0.007	0.003	0.007	-0.007
	(0.010)	(0.012)	(0.015)	(0.007)	(0.007)	(0.016)	(0.008)	(0.008)	(0.016)
GS_{t-2}× $D_{5G\cap 5S}$	-0.031	-0.035	-0.034	-0.022	-0.023	-0.033	-0.019	-0.020	-0.029
	(0.037)	(0.039)	(0.051)	(0.031)	(0.038)	(0.040)	(0.029)	(0.037)	(0.039)
GS_{t-2}× Ad. green × D_{5G5S}	-0.200***	-0.036	-0.324***	-0.159***	0.029	-0.273***	-0.172***	0.012	-0.280***
	(0.046)	(0.081)	(0.071)	(0.040)	(0.071)	(0.069)	(0.037)	(0.068)	(0.069)
GS_{t-3}× $D_{5S\setminus 5G}$	-0.059**	-0.053***	-0.069**	-0.031*	-0.024	-0.044**	-0.030*	-0.025	-0.041*
	(0.024)	(0.020)	(0.034)	(0.016)	(0.016)	(0.022)	(0.016)	(0.016)	(0.022)
GS_{t-3} × $D_{5G\setminus 5S}$	-0.034**	-0.037***	-0.028	-0.006	-0.006	-0.006	-0.008	-0.008	-0.007
	(0.013)	(0.010)	(0.027)	(0.006)	(0.008)	(0.018)	(0.007)	(0.009)	(0.018)
GS_{t-3} × $D_{5G\cap 5S}$	-0.051*	-0.099**	0.017	-0.026	-0.068*	0.026	-0.026	-0.068*	0.025
	(0.026)	(0.043)	(0.048)	(0.024)	(0.040)	(0.039)	(0.023)	(0.039)	(0.040)
GS_{t-3}× Ad. green × $D_{5G\cap 5S}$	-0.153***	-0.060	-0.258***	-0.117**	-0.007	-0.219***	-0.126***	-0.019	-0.220***
	(0.057)	(0.094)	(0.050)	(0.047)	(0.085)	(0.050)	(0.048)	(0.083)	(0.053)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	543044	386687	155015	543044	386687	155015	543044	386687	155015
R^2	0.030	0.027	0.036	0.022	0.020	0.026	0.021	0.019	0.025

ONLINE APPENDIX

A Selection of articles for the validation and testing sets

Given that the high fraction of articles in the Wall Street Journal are not climate risk-related, and the under-representation of the first twenty years of the database, sampling articles at random would not have been a feasible option. Therefore, the articles assigned to students were selected at random only during the first round of manual labelling, which included about a third of the total amount of articles labelled. For the remaining rounds of labelling, articles were selected such that they satisfy two conditions: 1) their similarity with a climate dictionary be higher than a given threshold, and 2) each decade be represented.

We follow Engle, Giglio, Kelly, Lee, and Stroebel (2020) to build a climate dictionary. We use the same set of selected documents that are known to be related to climate risk, e.g., authoritative documents from the IPCC. See Engle, Giglio, Kelly, Lee, and Stroebel (2020) for a detailed list of these documents. For each decade, we randomly select articles of the Wall Street Journal that were published in that decade, and compare the contents of these articles to the climate dictionary. The comparison is made by computing the cosine similarity. Articles with a cosine similarity below the threshold are discarded. We choose, on purpose, a low threshold to ensure that the probability of disqualifying relevant articles is low.

B Alternative Methods for Greenwashing Classification

We benchmark our method to build the greenwashing index to two alternative methods, namely those of Engle, Giglio, Kelly, Lee, and Stroebel (2020) and Sautner, van Lent, Vilkov, and Zhang (2023a) , adapted to our purpose.

B.1 Benchmark of Engle, Giglio, Kelly, Lee, and Stroebel (2020)

Engle, Giglio, Kelly, Lee, and Stroebel (2020) build a daily climate risk index from Wall Street Journal articles, by comparing the vocabulary in each daily edition of the Wall Street Journal to a reference climate-related vocabulary. The closer the vocabulary of the Wall Street Journal is to this reference vocabulary on a given day, the larger the index. The reference vocabulary is built as described in Appendix A.

To use the method of Engle, Giglio, Kelly, Lee, and Stroebel (2020) in the greenwashing-related classification, we first need to select a set of documents that we know are related to greenwashing. Tables A1 and A2 list the articles that we have selected. Figure A1 shows in a wordcloud the resulting greenwashing vocabulary, obtained from this corpus of documents.

Wall Street Journal articles are classified as greenwashing-related if their cosine similarity with this vocabulary is larger than a chosen threshold. The optimal threshold is optimized using the validation set.

Table A1: Non-academic reference documents on greenwashing for classification as in Engle, Giglio, Kelly, Lee, and Stroebel (2020)

This table lists the non-academic documents used in the implementation of the method of Engle, Giglio, Kelly, Lee, and Stroebel (2020) in the step 2 of our algorithm, i.e., to identify greenwashing-related articles.

Source	Title	Year
Laws, directives and press releases		
European Commission	Directive 2005/29/EC of the European Parliament and of the Council of 11 May 2005 concerning unfair business-to-consumer commercial practices in the internal market, and related documents	2005, 2013, 2016
European Commission	Press release - Screening of websites for 'greenwashing': half of the green claims lack evidence	2021
Organizations' (governmental, international, NGOs...) reports		
BSR	Understanding and Preventing Greenwash: A Business Guide	2009
OECD	Environmental Claims: Findings and Conclusions of the OECD Committee on Consumer Policy	2011
Underwriters Laboratories	Neither boastful nor bashful: Making effective sustainability claims	2016
Centre for Biological Diversity	Utility Greenwashing in Websites and Investor Reports	2019
Deloitte	Greenwashing or Measurable Results?	2019
SCM Direct	Greenwashing - Misclassification and mis-selling of ethical investments	2019
New Money	The New Money Guide to Greenwashing	2020
Good Energy	Renewable Energy Tariffs: The Problem of Greenwashing	2020
Shift Inside	Green Lies: Exploring Consumer Perceptions of Greenwashing	2020
Greenpeace	Words vs Actions: The truth behind fossil fuel advertising	2021
Ernst & Young	Greenwashing won't wash - The new sustainability imperative	2021
Allen & Overy	Greenwashing - key risks and issues in financial services	2021
Energy and Environment Alliance	Greenwashing - Legal guidance for the hospitality sector on environmental advertising claims	2021
Eversheds Sutherland	Greenwashing - A survey of recent litigation	2021
Generation Climate Europe	Greenwashing in the Fashion Industry - Policy Paper	2021
Geneva Center for Business and Human Rights	The Great Green Machine Part 1: Back to the Roots of Sustainability	2021
ShareAction	Combatting Greenwashing: TCFD Reporting	2021
Natural Resources Defense Council (NRDC)	Recycling Lies: "Chemical Recycling" of Plastic Is Just Greenwashing Incineration	2022
Climate Social Science Network	Climate-Washing Litigation: Legal Liability for Misleading Climate Communications	2022
Changing Markets	Licence to Greenwash	2022
Carbon Market Watch	Regulating Corporate Green Claims and Greenwashing	2022

Table A2: Academic reference documents on greenwashing for classification as in Engle, Giglio, Kelly, Lee, and Stroebel (2020)

This table lists the academic documents used in the implementation of the method of Engle, Giglio, Kelly, Lee, and Stroebel (2020) in the step 2 of our algorithm, i.e., to identify greenwashing-related articles.

Source	Title	Year
Academic papers and books		
Laufer (2003)	Social Accountability and Corporate Greenwashing	2003
Van Tulder and Van der Zwart (2006)	Sustainability Challenge # 14: Window Dressing	2006
Gibson (2009)	Awash in Green: A Critical Perspective on Environmental Advertising	2009
Vos (2009)	Actions Speak Louder than Words: Greenwashing in Corporate America	2009
Bazillier and Vauday (2009)	The Greenwashing Machine : is CSR more than Communication	2009
Furlow (2022)	Greenwashing in the New Millenium	2010
Lane (2010)	Consumer Protection in the Eco-Mark Era: A Preliminary Survey and Assessment of Anti-Greenwashing Activity and Eco-mark Enforcement	2010
Dahl (2010)	Green washing: Do you know what you're buying?	2010
Lyon and Maxwell (2011)	Greenwash: Corporate Environmental Disclosure under Threat of Audit	2011
Delmas and Burbano (2011)	The Drivers of Greenwashing	2011
Mitchell and Ramey (2011)	Look How Green I am! An Individual-level Explanation for Greenwashing	2011
Kim and Lyon (2015)	Greenwash vs. Brownwash: Exaggeration and Undue Modesty in Corporate Sustainability Disclosure	2015
Lyon and Montgomery (2015)	The Means and End of Greenwash	2015
Marquis, Toffel, and Zhou (2016)	Scrutiny, Norms, and Selective Disclosure: A Global Study of Greenwashing	2016
Kanda R.	Window Dressing In Financial Practices	2016
de Freitas Netto et al. (2020)	Concepts and forms of greenwashing: a systematic review	2020
Amenc, Goltz, and Liu (2021)	Doing Good or Feeling Good? Detecting Greenwashing in Climate Investing	2021
Pizzetti, Gatti, and Seele (2021)	Firms Talk, Suppliers Walk: Analyzing the Locus of Greenwashing in the Blame Game and Introducing 'Vicarious Greenwashing'	2021
Szabo and Webster (2021)	Perceived Greenwashing: The Effects of Green Marketing on Environmental and Product Perceptions	2021
Flagstad, Hauge, and Johnsen (2022)	Certification dissonance: Contradictions between environmental values and certification scheme requirements in small-scale companies	2022

Figure A1: Greenwashing Dictionary

This figure displays the n-grams that compose the greenwashing dictionary in our implementation of the method of Engle, Giglio, Kelly, Lee, and Stroebel (2020). These n-grams are obtained from a set of reference documents related to greenwashing, listed in Tables A1 and A2. The size of each n-gram depends on its term frequency.



B.2 Benchmark of Sautner, van Lent, Vilkov, and Zhang (2023a)

Sautner, van Lent, Vilkov, and Zhang (2023a) build a measure of intensity of climate risk in earning call transcripts by computing the frequency of words that belong to a climate-change dictionary, denoted by $CCExposure$. Their dictionary is not specified ex-ante, but is built, instead, in an iterative manner using machine learning methods. We calculate the measure $CCExposure$ for each article of our training set. When this measure is above a threshold, the article is classified as climate-related. The threshold is chosen by maximizing the F1-score on the validation set.

For the second step, we repeat the same procedure using the F0.5-score as optimization criterium. The greenwashing dictionary is obtained as described in Section B.1.

C Entity Linking

The following query was made on Wikidata, to obtain, for each company traded on the NYSE, the NASDAQ or the AMEX, the list of all the names that can be used to refer to this company (aliases):

```
SELECT DISTINCT ?id ?altLabel ?idLabel ?exchangesLabel ?ticker
WHERE {
SERVICE wikibase:label {
bd:serviceParam wikibase:language "[AUTO_LANGUAGE],en". }
?id p:P414 ?exchange.
VALUES ?exchanges { wd:Q13677 wd:Q82059 wd:Q846626}
?exchange ps:P414 ?exchanges;
pq:P249 ?ticker.
OPTIONAL { ?id skos:altLabel ?altLabel .
FILTER (lang(?altLabel) = "en") }
}
```

D Topic Analysis using the LDA method

Table A3: LDA results

This table indicates the words identified as the most representative of each topic identified by the LDA method.

Topic	Most characteristic words
1	**fund**, esg, investment, plastic, green, industry, paper, group, product, energy
2	**fund**, green, investor, investment, esg, bond, energy, etf, carbon, firm
3	green, product, business, **fund**, bond, issue, investor, shareholder, social, group
4	**fund**, esg, plastic, investor, bottle, group, product, investment, firm, business
5	green, **fund**, investor, esg, change, business, firm, group, investment, energy
6	product, **fund**, investor, green, stock, group, etf, esg, energy, plastic
7	energy, green, investor, bond, business, sustainability, social, climate, group, people
8	esg, **fund**, energy, repor, investor, investment, issue, group, change, social
9	**fund**, green, product, esg, plastic, bag, investor, bond, etf, group
10	product, energy, **fund**, green, investor, group, issue, consumer, industry, based
11	**fund**, green, investor, energy, product, investment, investing, group, asset, etf
12	esg, **fund**, investment, bag, product, business, manager, public, group
13	esg, **fund**, investment, emission, carbon, investor, business, asset, firm, group
14	bag, plastic, product, investor, **fund**, consumer, social, use, group, industry
15	product, green, **fund**, bond, group, sustainable, paper, consumer, claim, energy
16	**fund**, esg, investor, etf, product, asset, investment, market, manager, green
17	esg, **fund**, group, plastic, green, investment, product, consumer, investor, cup
18	**fund**, investor, esg, shareholder, manager, asset, vote, investment, social, group
19	product, bag, energy, green, group, investor, business, **fund**, consumer, change
20	**fund**, energy, investor, climate, green, carbon, esg, plastic, emission, group