

A Greenwashing Index *

Elise Gourier

ESSEC Business School and CEPR

Hélène Mathurin

ESSEC Business School

February 2024

Abstract

We construct a news-implied index of greenwashing. Our index reveals that greenwashing has become particularly prominent in the past five years. Its increase was driven by skepticism towards the financial sector, specifically ESG funds, ESG ratings and green bonds. We show that greenwashing impacts investors' behavior. Unexpected increases in the greenwashing index are followed by decreases of flows into funds advertised as sustainable, both for retail and institutional investors. They furthermore bias the estimation of stocks' beta on climate risk, distorting the estimated climate risk premium. When accounting for greenwashing, the climate risk premium becomes small and statistically insignificant.

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, Laurence Daures, Andras Fulop, José-Miguel Gaspar, Alan de Genaro (discussant), Christian Julliard, Alessandro Melone, Ludovic Phalippou, Alberto Quaini, Roberto Reno and Gustavo Schwenkler (discussant) for feedback and suggestions on our project. We also thank participants to seminars at SKEMA, Vienna Graduate School of Finance and participants to the 5th International Workshop in Financial Econometrics for comments. This research has been supported by a JCJC ANR grant and a grant from the "Institut Europlace de Finance and Labex Louis Bachelier".

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

Measuring climate risks is not straightforward. A recent literature, building on the paper of [Engle, Giglio, Kelly, Lee, and Stroebele \(2020\)](#), uses indices based on the volume of news on climate change as proxies for climate risks. The underlying assumption is that when there are events causing a rise in climate risks, these events are likely to be covered by the media. Greenwashing may distort these indices, both downwards and upwards. Firms that greenwash may seem less exposed to climate risks than they actually are, causing an underestimation of these risks. But articles focusing on greenwashing do not bring any relevant information on climate risks per se, wrongly inflating climate risk indices.

In this paper, we build an index which measures the fraction of climate-related news articles mentioning firms’ greenwashing. This index can be interpreted as a measure of attention to greenwashing. Similar to news-implied indices of climate risk, it can also 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, firms 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

To build our greenwashing index, 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 less than 10 articles before 2007, and in less 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

¹New 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.

²It 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>.

the words or combination of words that characterize articles related to greenwashing, even if they do not contain the word itself. Second, the dataset 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 that are 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 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, climate risk-related articles are well identified using this approach. However, it does not allow identifying precisely 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.

Articles classified as greenwashing-related represent a small part, on average 0.07%, of all selected articles. Whereas 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, attention to greenwashing has therefore grown to a magnitude of the same order as that of attention to climate risk.

Our greenwashing index counts the greenwashing-related articles as a fraction of climate risk-related articles. 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 repeated criticisms of the 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 attention to greenwashing. 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 affect investors' choices? Does it impair the measurement of climate risk?

The answer to both questions is yes. To study the effect of greenwashing on investors' choices, we examine the response of market demand in green funds 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 these funds. Consistent with this, we find, using time series and panel regressions, that both retail and institutional investors strongly react to unanticipated shocks in the greenwashing index by decreasing their flows into green funds: a one standard deviation unexpected increase in the greenwashing index causes institutional investors to decrease their investments in green funds by 8% the following month, and retail investors by 10% within two months. Their reaction however differs along two dimensions. First, the timing is different: The reaction of institutional investors is short-lived, with a rebound of inflows in green funds two months after the greenwashing news release. Second, they target different funds. Retail investors' mostly target funds advertised as green in their name *and* with a five-globe Morningstar rating. Institutional investors, in turn, target funds also advertised as green but with lower globe rating. This difference suggests that retail investors lose confidence in Morningstar ratings when reading about greenwashing, while institutional investors do not.

We further show that greenwashing impairs the measurement of climate risk and of the climate risk premium. The recent literature came to the consensus that the climate risk premium has been time-varying, and positive albeit small over the past 20 years, except following the financial crisis. As a side-product of our two-step classification of greenwashing-related articles, we build a climate risk index which exhibits high correlation with the indices in the literature, e.g., [Engle et al. \(2020\)](#). In line with the literature, we find that innovations to this index carry a time-varying premium that is positive and statistically significant most of the time, when abstracting from greenwashing. However, when accounting for greenwashing, results change drastically. We allow the beta of stock returns on climate innovations to load on the greenwashing shocks, to reflect the effect of greenwashing on the measurement of firms' sensitivity to climate risk. We estimate the resulting climate risk premium, using the method of [Gagliardini, Ossola, and Scaillet \(2016\)](#), and find that the premium for climate risk becomes smaller (by an order of 3), and loses its statistical significance. We conclude that greenwashing has an effect on the beta of stocks on climate risk, and that this effect absorbs the climate risk premium.

Finally, we investigate whether greenwashing is itself related to the cross-section of returns. As our index is mostly driven by the financial sector, we need to be careful about the possibility that greenwashing shocks may be a weak factor, and only affect part of the stocks. This may result in an artificially inflated risk premium. Our tests suggest that greenwashing is not a spurious factor, and had an impact on the cross-section of excess realized returns. This impact was sizable before the financial crisis and over the past five years.

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. They find that only the climate policy factor is priced in the cross-section of U.S. stocks.

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. Their measures are hence likely to be affected by 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 media’s attention to greenwashing.

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 ESG terminology.

With the increasing popularity of the environmental, social and governance (ESG) framework, the definition of greenwashing has in the past years often been broadened to misleading claims on ESG issues. In this paper, we use the traditional definition, which focuses on environmental claims.

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. The U.K. was probably the first country to implement these 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., a new Corporate Sustainability Reporting Directive (CSRD) entered into force in the beginning of 2023, and introduces 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. 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 years 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 ordinance 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 proposed in 2022 new rules on climate disclosure that are aligned with the TCFD framework. But the proposal remains to be finalized, and this last step has been delayed until the Spring 2024.

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.

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 has not been adopted yet. Recently, the European Commission published its proposed Directive on Green Claims on March 22, 2023.³ 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 environmental claims. The proposal also requires that all environmental claims and labels be independently verified and certified before being publicized. Furthermore, 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 first progress report was released on May 31, 2023, by the [European Securities and Markets Authority \(2023\)](#). The Final Report is expected to be published in May 2024, and will include possible changes to the EU regulatory framework. Finally, 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.⁴

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,

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

⁴See <https://www.europarl.europa.eu/news/en/press-room/20230918IPR05412/eu-to-ban-greenwashing-and-improve-consumer-information-on-product-durability> for details.

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.⁵

4 Measuring Attention to Greenwashing

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 investors’ attention to greenwashing. 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 attention to greenwashing. 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 2022.⁶ We filter out articles that belong to the opinion section of the Wall Street Journal as well as pieces that we deem irrelevant.⁷ 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.

⁵See, <https://www.sec.gov/news/press-release/2023-188>.

⁶Since 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.

⁷The 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.

4.2 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 notion of greenwashing can be blurry, we encouraged the students to consult us for cases they had doubts about. This led us to take a stand on whether some particular articles should be considered greenwashing-related or not. We decided to include the following articles in the set of greenwashing-related articles: 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' 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). Importantly, we did not include articles that did not mention investors, or customers, being misled or confused by these concepts.

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.

4.2.1 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 Stroebl (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} .⁸ 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,

⁸Using 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} .

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 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.⁹

⁹The 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

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).

4.2.2 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"), "greenwashing" (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.

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.

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).

4.3 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 Stroebl \(2020\)](#) 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 articles¹⁰, 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 than 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 articles. Indeed, whereas a fair number of bigrams are specific to climate risk (e.g., global warming, greenhouse gas, Kyoto protocol 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.

Finally, we perform the same two-step classification using the OpenAI API of GPT-3.5 Turbo. Table 2 shows that GPT brings a marginal improvement in performance in the first step, and a substantial improvement in the second step. This improvement likely results from the fact that Large Language Models are able to fully account for context, at the expense of interpretability.

¹⁰The 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 Skiadopoulou (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 repeated 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

Attention to greenwashing 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 performs 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.¹¹

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.¹²

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

¹¹Wikidata SQL queries can be performed on <https://query.wikidata.org/>. The detail of the query is provided in Appendix C.

¹²The 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 model¹³, 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",

¹³The 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 6.1, 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 these 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 Affect Investors' Behaviors?

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 green funds. In this section, we study whether investor demand for funds reacts to unexpected greenwashing news. On the one hand, investors with preferences including a non-pecuniary taste for asset greenness may update their expected utility from investing in these funds, after reading that they may not be as green as they thought. On the other hand, a return-chasing investor may anticipate a decrease in green funds' returns following the release of greenwashing news. Both effects would modify investors' beliefs about the sustainability-performance tradeoff (Gantchev et al. (2024)) and induce decreased flows into green funds. Under similar arguments, climate risk-related news may incentivize investors to invest more in these funds, as part of a hedging strategy or of a return-chasing strategy. The question of the impact of greenwashing and climate shocks 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 2022,¹⁴ and focus on the evolution of their monthly flows. Flows in dollars are defined as the sum of flows across share classes, divided by fund size.¹⁵ We filter out funds with size less than a million dollars. Both relative flows and fund 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 a Morningstar sustainability rating. Table 5 summarizes the sample. The final dataset contains 7,613 funds.

We first study the relation between average fund flows and greenwashing in the time series, for various groups of funds. Specifically, we regress average fund flows (normalized by fund sizes) on lagged shocks of the greenwashing and climate indices. To construct greenwashing and climate shocks, we compute the innovations of an AR(1) model applied to each index. The two resulting time series of shocks are almost perfectly uncorrelated,

¹⁴Morningstar 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.

¹⁵Following Chevalier and Ellison (1997) and Sirri and Tufano (1998), net fund flows 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 fund size as in Hartzmark and Sussman (2019).

with a correlation of -0.04. We denote them by GS_t (greenwashing shocks) and CS_t (climate shocks). We control for variables which have been shown to be important for flows: the average 1-month, 1-year and 2-year lagged returns, and the average lagged fund size. The regressions run as follows:

$$\text{Fund flows}_t = a_0 + a_1 \text{Fund flows}_{t-1} + \sum_{h=1}^2 b_h GS_{t-h} + \sum_{l=1}^2 c_l CS_{t-l} + \text{Controls}_{t-1} + \epsilon_t \quad (1)$$

To rule out reverse causality, all variables on the left hand-side are lagged, and hence known at time $t - 1$. Table 6, Panel A, reports the results of regression (1).

The first column considers all funds in the sample, and shows a negative and statistically significant association between past greenwashing shocks (at time $t - 2$) and current flows. Furthermore, there is a positive association between past climate shocks (both at time $t - 1$ and $t - 2$) and current flows. These results indicate that an increase in climate risk has an overall positive impact on fund flows, while an increase in greenwashing has a negative impact on these flows. These results hold in aggregate across all funds, and do not distinguish between green funds and other funds.

The next step is to test whether green funds are driving these variations in flows. It is well known that mutual fund investors favor sustainability-oriented funds, see, e.g., [Bialkowski and Starks \(2016\)](#), [Hartzmark and Sussman \(2019\)](#) and [Pàstor and Vorsatz \(2020\)](#). Several proxies of fund greenness are available to investors. The first is the name of the fund: many funds are advertised as sustainable in their names. 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 such funds. The second proxy of fund greenness is the Morningstar globe rating. Since 2016, Morningstar has been giving monthly sustainability ratings to funds, based on the underlying holdings. An aggregated score is calculated for each fund, 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%). There is no one-to-one relation between the two proxies for funds' sustainability: nearly 40% of the funds advertised as green 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. Both the name of the fund, if containing a green word, and a high Morningstar globe rating, likely increase the salience of a fund. Furthermore, both proxies have flaws. While the name of the fund may or may not really

reflect the fund's really sustainability level¹⁶, the Morningstar rating is given by inspection of each fund's holdings, and aggregation of their Sustainalytics ESG ratings. ESG ratings were, in Section 5.4, among the main topics of greenwashing-related articles. Therefore it is not clear whether investors consider the Morningstar rating more reliable than the fund's name.

The second column of Table 6 highlights a negative link between greenwashing shocks at time $t-1$ and flows into funds advertised as green. Although the coefficient is not statistically significant, it contrasts with the lack of link between greenwashing shocks at time $t-1$ and flows into funds that are *not* advertised as green (column 3). The negative link between greenwashing shocks at time $t-2$ and flows, however, is the same for funds advertised as green and other funds. Columns 4 and 5 do not highlight any difference in reaction between low Morningstar rated funds, and high rated funds.

Our results, so far, indicate that investors' reaction to greenwashing news tends to target funds advertised as green more than other funds. Similarly, following a climate risk shock of one standard deviation, flows into self-labelled green funds increase by 30.6% of standard deviation. This increase is significant at the 1% level, and three times as large as the increase in flows into funds that are not advertised as green. There is, here again, no difference in the reaction of funds with low and high Morningstar rating.

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. In Panel B, we repeat the same regressions, but distinguish between funds targeting retail investors and funds targeting institutional investors. Panel B of Table 6 reveals some interesting differences between the reactions of retail and institutional investors. First, the negative reaction of flows into self-labelled green funds is mostly driven by institutional investors. A one standard deviation greenwashing shock leads flows into self-labelled green funds to decrease by 9.4% of standard deviation. This reaction is large and statistically significant at the 1% level. In contrast, the reaction of retail investors is small and not statistically significant. Second, both retail and institutional investors react to climate risk shocks with large inflows into self-labelled green funds, but the reaction of retail investors is twice as large as the one of institutional investors.

We further investigate these differences between retail and institutional investors by distinguishing between funds with Morningstar globe rating from one to four, and five-globe

¹⁶The 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.

funds. Table 7 shows that retail investors do react to greenwashing news, but specifically target funds that are both advertised green *and* have a five-globe rating. The reaction is large and significant, with a decrease in fund inflows by 12.6% one month after the news, and by 6.6% two months after. In contrast, institutional investors lower their inflows into all funds that are advertised as green, irrespective of the Morningstar rating. Their reaction is large but short-lived, with nearly no effect two months after the news. These results are in line with retail investors relying on both the Morningstar rating and on the fund name to form their opinion on a fund's sustainability. Funds with either a low Morningstar rating or a name unrelated to sustainability, would not be believed to be truly sustainable. Greenwashing news would thus affect flows into the intersection of funds that are considered sustainable, i.e., funds that are both advertised as green and with a five-globe rating. Institutional investors, in contrast, seem to put more emphasis on funds' name when evaluating sustainability: greenwashing news lead to lower investments in funds that are advertised as green, irrespective of their globe rating.

Both types of investors also react differently to climate shocks. Retail investors increase their inflows into all funds advertised green, but the increase is nearly twice as large for five-globe funds. In line with the previous results, this indicates that retail investors pay attention to both the name of the fund and the Morningstar rating to identify green funds. Institutional investors target funds with low globe ratings, but advertised green, and funds with high globe rating not advertised green.

In order to understand whether our results are robust to the inclusion of fund-level variables, we run panel regressions of flows on a dummy that is 1 if the fund is advertised as green, and on this dummy interacted with lagged shocks of the greenwashing index. Specifically, we start from the following regression:

$$\begin{aligned} \text{Fund flows}_{i,t} = & a_0 + a_1 \text{Fund flows}_{i,t-1} + a_2 \times \text{Advertised green}_i \\ & + \sum_{h=1}^2 b_h GS_{t-h} \times \text{Advertised green}_i \\ & + \text{Controls}_{i,t-1} + \text{Category/year-fixed effect.} \end{aligned} \quad (2)$$

Table 8 reports the results of this regression, run for retail and institutional investors separately (in Panels A and B, respectively). We control for lagged flows, lagged returns over a month, a year and two years, lagged fund size, lagged star rating and fund age. Standard errors are clustered by Morningstar category - year fixed effect.

The results confirm those of the time series regressions: Retail and institutional investors

both react to greenwashing news by investing less into funds that are advertised as green, but retail investors react one month later. Institutional investors react in the month following the news, but two months after, they get back to buying these funds, so that the effect is offset.

We further examine specifications including more lags of the greenwashing shocks, as well as climate shocks, and dummies indicating whether funds have five Morningstar globes:

$$\begin{aligned}
 \text{Fund flows}_{i,t} = & a_0 + a_1 \text{Fund flows}_{i,t-1} + a_2 \times \text{Advertised green}_i \\
 & + \sum_{h=1}^3 b_h G S_{t-h} \times \text{Advertised green}_i \\
 & + \sum_{h=1}^2 c_h G S_{t-h} \times \text{Advertised green}_i \times 5 \text{ globes}_{i,t} \\
 & + \sum_{l=1}^2 d_l C S_{t-l} \times \text{Advertised green}_i + \text{Controls}_{i,t-1} \\
 & + \text{Category/year-fixed effect.}
 \end{aligned} \tag{3}$$

Adding these variables highlights a second major difference between retail investors and institutional investors, in line with the results of the time series regressions: Retail investors do react one month after the greenwashing news, but they target funds that are both advertised as green and have a five-globe rating. The effect is large and statistically significant: They decrease their inflows into these funds by 10.6% (columns 3 and 4). One month later, the decrease in inflows remains, and extends to all funds advertised as green, irrespective of their Morningstar rating (columns 1 to 4). In contrast, institutional investors target funds advertised as green but with low Morningstar rating, with a decrease in inflows of around 8% (columns 7 and 8). Five-globe funds are spared, although the positive coefficient in front of the crossed dummies is not statistically significant. As highlighted in regression (2), the reaction is short-lived, with a rebound of inflows into all funds advertised green two month after the news release.

To sum up, we find, using both time series and panel regressions, that unexpected greenwashing shocks are followed by a large and significant decrease in flows into funds that are advertised as sustainable. We uncover interesting differences in the reactions of retail and institutional investors: Both the target and the timing of their reactions differ.

6.2 Does Greenwashing Impair the Measurement of Climate Risk Premiums?

In this section, we study the impact of unexpected shocks to our indices of climate risk and greenwashing on stock returns. We construct these shocks similarly to before, after standardizing the climate risk and greenwashing indices to have zero mean and an annual volatility of 10%.

We obtain two factors, which are almost uncorrelated (correlation of 4%), and exhibit low correlation with other standard factors. All correlations are displayed in Table 9.

It is not clear whether attention to climate risk should command a risk premium. On the one hand, it could be that firms that are more sensitive to unexpected shocks to climate risk are riskier: they could be more affected by the regulations, technology changes or physical events that are mentioned in the news. In this case, they should bear a positive risk premium. On the other hand, these could also be the firms that are the most pro-active developing new technologies and policies to address climate change. In this case, investors may tolerate additional risks for non-pecuniary reasons, see, e.g., [Pastor et al. \(2021\)](#); [Pedersen et al. \(2021\)](#) and [Zerbib \(2022\)](#).

[Sautner, van Lent, Vilkov, and Zhang \(2023b\)](#) use Fama MacBeth regressions ([Fama and MacBeth \(1973\)](#)) to calculate the risk premium associated with high exposure to climate risk. They build their exposure measure from transcripts of quarterly earnings calls. They find evidence for a small unconditional risk premium. Over their time series, however, they find that the risk premium is positive between 2006 and 2009, and most of the time after 2015. They attribute the negative premium between 2009 and 2015 to the economy's recovery following the financial crisis. [Sautner, van Lent, Vilkov, and Zhang \(2023a\)](#) use the extension of the Fama MacBeth regression framework developed by [Gagliardini, Ossola, and Scaillet \(2016\)](#) to improve their study of the evolution of the risk premium over time. They find evidence for a premium that is positive at almost all times.

In presence of greenwashing, the positive risk premium can now be explained by a firm's exposure to climate risk, but also by the additional uncertainty induced by the public's suspicion that the firm is greenwashing. We attempt to disentangle the two effects.

Formally, we consider an augmented 4-factor model, which includes shocks to our climate index, CS_t as well as the three Fama French factors and the momentum factor of [Carhart \(1997\)](#):

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_{i,t}^{Clim} CS_t + \beta_{i,t}^{Mkt} Mkt_t + \beta_{i,t}^{SMB} SMB_t + \beta_{i,t}^{HML} HML_t + \beta_{i,t}^{Mom} Mom_t + \epsilon_{i,t} \quad (4)$$

Similarly to [Sautner, van Lent, Vilkov, and Zhang \(2023a\)](#), we work in a small- T sample, for which the method of [Gagliardini, Ossola, and Scaillet \(2016\)](#) is well adapted. Excess returns follow a linear conditional factor model, with time-varying factor loadings and risk premiums. The factor loadings are modelled as linear functions of lagged instruments, some of them being common to all assets (Z_{t-1}) and some asset-specific ($Z_{i,t-1}$):

$$\beta_{i,t} = \beta_c Z_{t-1} + \beta_a Z_{i,t-1}. \quad (5)$$

Following [Gagliardini, Ossola, and Scaillet \(2016\)](#), we use the term spread and the default spread as common instruments, and the log of the book-to-market ratio as asset-specific instrument.

In order to measure whether greenwashing impairs the measurement of the risk premium associated to climate risk, we estimate two specifications of this model. The first specification abstracts from greenwashing, whereas the second specification adds the shocks to our greenwashing index to the common instrument vector Z_{t-1} . If greenwashing does not impair the measurement of the risk premium for climate risk, then it will receive a loading of zero in the second specification, and the risk premium will be the same as the one in the first specification. If, in turn, greenwashing increases the returns in a predictable way, then the beta on climate shocks will load on the greenwashing (lagged) shocks, potentially capturing part of the risk premium.

Panels A and B of [Figure 10](#) display the time-varying risk premium of climate risk in the two model specifications. In Panel A, greenwashing is not taken into account in the model. In line with the results of [Sautner, van Lent, Vilkov, and Zhang \(2023b\)](#) and [Sautner, van Lent, Vilkov, and Zhang \(2023a\)](#), we find that the conditional risk premium is positive most of the time, reaching 4% in 2007 and 2020. It is only negative during a short time period following the financial crisis. It is statistically significant between 2005 and 2008, and then most of the time after 2013. Unconditionally, it is also sizable, more than 2% per annum. In Panel B, greenwashing is part of the instruments, and the resulting premium looks very different. It crosses 0 many times, and no longer goes above 3%. The unconditional level decreases from 2.2% per annum to 0.9% per annum. Importantly, it is no longer statistically significant, at any time during the period considered.

Greenwashing, used as an instrument, therefore absorbs most of the climate risk premium that was uncovered in the first specification. The resulting risk premium for climate risk is close to zero and insignificant.

Panels C and D report the results when using an augmented six-factor model, adding the investment and profitability factors of Fama and French (2015). The results are similar to those of the four-factor model: not accounting for greenwashing gives the illusion of a sizable climate risk premium, whereas using greenwashing as an instrument reveals that it is in fact oscillating around zero and statistically insignificant.

To sum up, our results confirm that greenwashing does impair the measurement of climate risk premium, making it look inflated.

6.3 Is Greenwashing a Risk Factor Itself?

Section 6.2 shows that greenwashing impairs the measurement of the climate risk premium, and that using greenwashing shocks as common instrument makes this risk premium disappear.

In this Section, we go one step further and investigate whether greenwashing can itself be considered a factor. We augment the model by adding shocks to the greenwashing index, GS_t :

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_{i,t}^{Clim} CS_t + \beta_{i,t}^{Greenw} GS_t + \beta_{i,t}^{Mkt} Mkt_t + \beta_{i,t}^{SMB} SMB_t + \beta_{i,t}^{HML} HML_t + \beta_{i,t}^{Mom} Mom_t + \epsilon_{i,t}. \quad (6)$$

Lagged shocks to the greenwashing index are still used as common instruments.

Even though Section 4 has highlighted that many industries are concerned by greenwashing, the main result was that greenwashing concerns had been mainly driven by the financial sector since 2018. Before 2018, greenwashing was small in the news. It is therefore possible that greenwashing, as a factor, be a weak factor, i.e., that it has small covariance with the returns of part (or all) of the assets. Weak factors have been shown to lead to spurious high risk premiums, Bryzgalova (2015). It is hence important to restrict the test assets that are used in the second pass of the Fama MacBeth regressions, to those that have a high enough covariance (in absolute value) with the weak factor, Giglio, Xiu, and Zhang (2023).

Figure 11 displays the scree plot of the eigenvalues of the matrix $\beta_t^T \beta_t$, where $\beta_t = [\beta_t^{Clim}, \beta_t^{Greenw}, \beta_t^{Mkt}, \beta_t^{SMB}, \beta_t^{HML}, \beta_t^{Mom}]$ with the notations of equation 6. In the extreme case where one of the β 's, for example β^{Greenw} , is zero for all assets, the resulting matrix

$\beta_t^T \beta_t$ will be singular, and will have one or more eigenvalues at 0. The distance of the eigenvalues to zero is therefore an indicator of whether a model contains weak factors. In our case, all eigenvalues are far from zero, for all the dates considered. We are therefore not in a situation where we have one (or more) factor with no impact on any asset. The first eigenvalue (which correspond to the market factor) is substantially higher than all the others. The two smallest eigenvalues correspond to the climate and to the greenwashing shocks. These two factors may therefore be weaker than the other factors.

In order to address the concern of the climate and greenwashing factors possibly being weak, we restrict the test assets used in the second pass of the regression, in the spirit of Giglio, Xiu, and Zhang (2023). In their setup, risk premiums are constant, and they are working with a long time series of data. Therefore, they choose to keep $\lfloor qn \rfloor$ test assets with the highest β 's, where q is chosen by splitting the time sample and running cross-validation. As we work with a short time period and time-varying premiums, we cannot apply cross-validation to choose the optimal number of test assets to keep. We decide to rank the β^{Clim} 's and β^{Greenw} 's by their absolute value, and to discard the stocks with one of these β 's in the bottom 20%, each month. The cross-sectional regression is therefore run only on test assets with β 's sufficiently far from zero.

Figure 12, Panels A and B, display the estimated risk premiums for climate risk and greenwashing, when all individual stocks are included as test assets. In line with the results of Section 6.2, the premium for climate risk is small, slightly positive on average but insignificant throughout the time series. The premium for greenwashing is much more sizable, with peaks at 4% p.a., and positive throughout the time series except after the financial crisis in 2009. It is statistically significant from 2005 to 2008, in 2013 and 2015 punctually, from 2016 to 2019 and in 2020 punctually. Overall, our results on the greenwashing premium are therefore much more convincing than those on the climate risk premium.

This large premium could however be an artefact of greenwashing shocks being a weak factor. Panels C and D display the same premiums, but assets whose absolute beta on climate or greenwashing is in the bottom 20% have been discarded from the second pass of the regression. If climate risk, or greenwashing is a weak factor, the corresponding premium should disappear. Panel C shows that restricting the test assets yields a risk premium for climate risk that does not change much and is still never statistically significant. The premium for greenwashing, in Panel D, looks very similar to the one in Panel B, albeit with a slightly smaller magnitude.

To sum up, restricting the test assets to account for the possibility of weak factors indicates a small positive premium for climate risk, and a sizable positive premium for

greenwashing. These premiums reveal a possible weak link between climate risk and returns, and stronger link between greenwashing and returns.

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.
- Bialkowski, J. and L. T. Starks (2016). Sri funds: Investor demand, exogenous shocks and esg profiles. 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.
- Bryzgalova, S. (2015). Spurious Factors in Linear Asset Pricing Models. Working paper.
- Cao, S., Y. Kim, A. Wang, and H. Xiao (2020). From Words to Syntax: Identifying Context-specific Information in Textual Analysis. Working paper.
- Carhart, M. (1997). On persistence in mutual fund performance. *Journal of Finance* 52(1), 57–82.
- 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. Stroebe (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.
- Fama, E. F. and J. D. MacBeth (1973). Risk return and equilibrium: Empirical tests. *Journal of Political Economics* 81(3), 607–636.
- 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.
- Gagliardini, P., E. Ossola, and O. Scaillet (2016). Time-varying risk premium in large cross-sectional equity data sets. *Econometrica* 84, 985–1046.

- 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.
- Giglio, S., D. Xiu, and D. Zhang (2023). Test assets and weak factors. Working Paper.
- 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.
- 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.
- 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.
- 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.
- Pàstor, L. and M. B. Vorsatz (2020). Mutual fund performance and flows during the covid-19 crisis. *Review of Asset Pricing Studies* 10(4), 791–833.
- Pastor, L. S., R. F. Stambaugh, and L. A. Taylor (2021). Sustainable investing in equilibrium. *Journal of Financial Economics* 142(2), 550–571.
- Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021). Responsible investing: the ESG-efficient frontier. *Journal of Financial Economics* 142(2), 572–597.
- 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.
- Zerbib, O. D. (2022). A sustainable capital asset pricing model (s-capm): Evidence from green investing and sin stock exclusion. *Review of Finance* 26, 1345–1388.

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 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 4.2. 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 4.2. 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 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: Time-Varying Risk Premium for Climate Risk

This figure shows the risk premium, in percentage per annum, associated to climate risk, calculated using the method of [Gagliardini, Ossola, and Scaillet \(2016\)](#). The confidence bounds at the 95% level are displayed. In Panels A and B, the 4-factor model of equation 4 is used. In Panels C and D, a 6-factor model is used which augments model 4 with the investment and profitability factors of [Fama and French \(2015\)](#). In Panels A and C, greenwashing is absent from the factor model. In Panels B and D, shocks to the greenwashing index are added as common instruments following equation 5.




Figure 11: Scree plot

This figure shows the log eigenvalues of the matrix $\beta_t^T \beta_t$, for a chosen point in time t , where $\beta_t = [\beta_t^{Clim}, \beta_t^{Greenw}, \beta_t^{Mkt}, \beta_t^{SMB}, \beta_t^{HML}, \beta_t^{Mom}]$ with the notations of equation 6. β_t is an $n \times 6$ matrix, where n is the number of stocks. The six lines correspond to six points in time: January 2002, January 2006, January 2010, January 2014, January 2018 and January 2022.




Figure 12: Time-Varying Risk Premium for Climate Risk and Greenwashing

This figure shows the risk premium, in percentage per annum, associated to climate risk (Panels A and C) and greenwashing (Panels B and D), calculated using the method of [Gagliardini, Ossola, and Scaillet \(2016\)](#). The confidence bounds at the 95% level are displayed. In Panels A and B, all individual stocks are used. In Panels C and D, the stocks whose first-pass absolute beta is in the lowest 20% are discarded from the second pass. The factor model used is the one in equation (6).

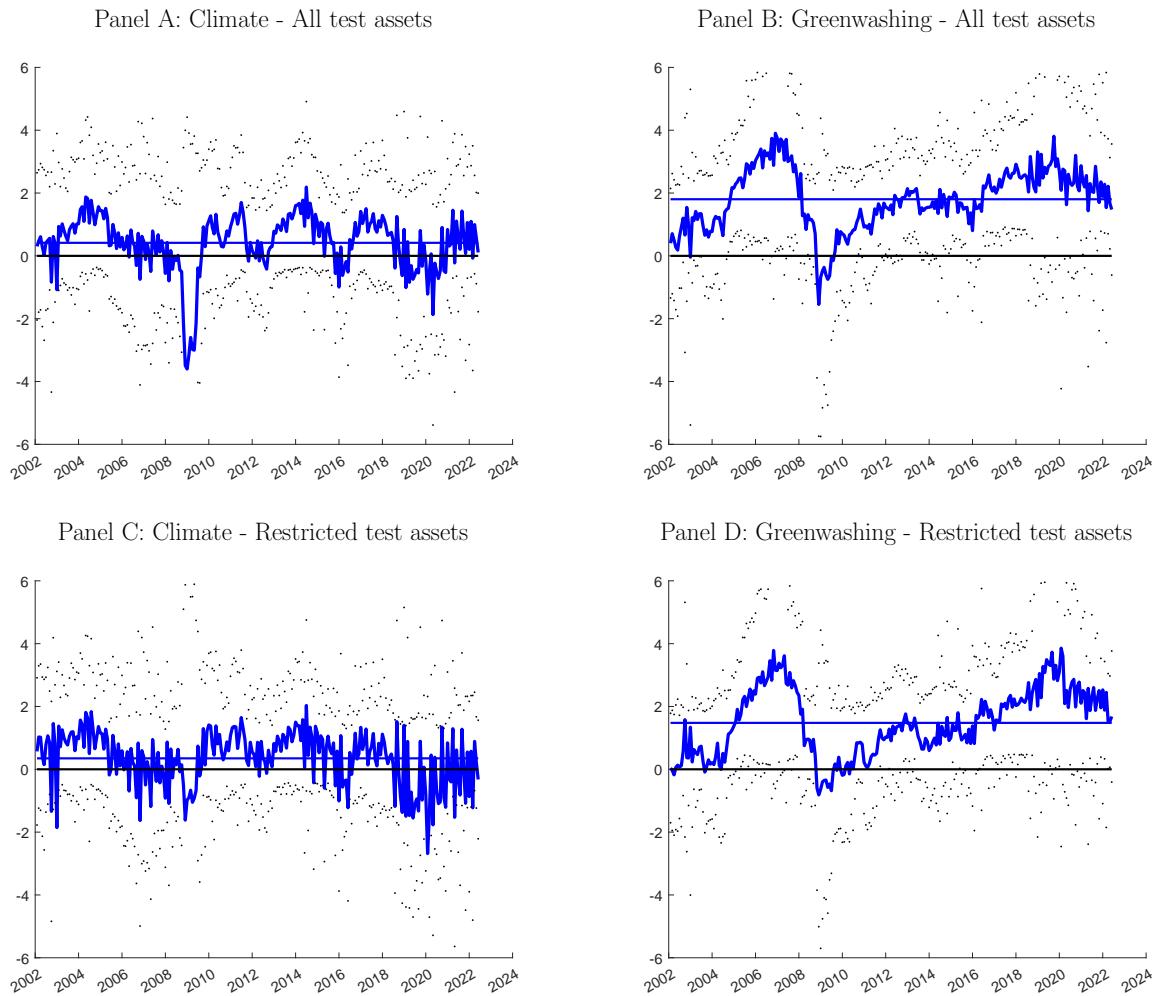


Table 1: Description of Sets Used in the Classification

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 4.2.

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 4.2.

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
	GPT	93.1	92.7	93.1	92.9	92.5
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
	GPT	94.7	91.9	63.2	84.2	90.0

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: 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 6: Time Series Regressions of Flows on Greenwashing Shocks

This table reports the results of regressions of average flows (relative to fund sizes) on lagged shocks of the greenwashing (GS) and climate (CS) indices, following equation (1). The columns report the results for different groups of funds: 1) all funds (column 1), funds advertised as green or not (columns 2 and 3), funds with Morningstar globe rating of 1 to 4 globes (column 4) and five-globe funds (column 5). 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.

Panel A: Funds advertised as green versus not						
	Overall	Ad. green	Not ad. green	1G to 4G	5G	
GS _{t-1}	0.003 (0.017)	-0.026 (0.036)	0.003 (0.015)	-0.000 (0.016)	0.020 (0.027)	
GS _{t-2}	-0.032** (0.016)	-0.033 (0.032)	-0.032** (0.016)	-0.032* (0.017)	-0.033*** (0.013)	
CS _{t-1}	0.113** (0.055)	0.306*** (0.099)	0.100* (0.054)	0.115** (0.058)	0.081 (0.052)	
CS _{t-2}	0.064 (0.063)	0.182* (0.097)	0.055 (0.063)	0.055 (0.062)	0.123 (0.100)	
Controls	yes	yes	yes	yes	yes	
Adj. R ²	0.858	0.900	0.839	0.844	0.898	
N funds	7613	392	7221	-	-	
Obs.	46	46	46	46	46	

Panel B: Funds Targeting Retail or Institutional Investors						
	Retail			Institutional		
	Overall	Ad. green	Not ad. green	Overall	Ad. green	Not ad. green
GS _{t-1}	0.003 (0.019)	-0.013 (0.044)	0.002 (0.018)	0.003 (0.009)	-0.094*** (0.029)	0.008 (0.009)
GS _{t-2}	-0.037** (0.018)	-0.082 (0.053)	-0.035** (0.018)	-0.024 (0.020)	0.041 (0.045)	-0.025 (0.020)
CS _{t-1}	0.146*** (0.053)	0.339*** (0.100)	0.135** (0.054)	0.023 (0.064)	0.174* (0.093)	0.015 (0.062)
CS _{t-2}	0.075 (0.070)	0.189 (0.129)	0.068 (0.070)	0.022 (0.065)	0.097 (0.091)	0.014 (0.065)
Controls	yes	yes	yes	yes	yes	yes
Adj. R ²	0.900	0.893	0.886	0.244	0.706	0.151
N funds	5353	265	5088	2151	115	2036
Obs.	46	46	46	46	46	46

Table 7: Time Series Regressions of Flows on Greenwashing Shocks (cont'd)

This table reports the results of regressions of average flows (relative to fund sizes) on lagged shocks of the greenwashing (GS) and climate (CS) indices, following equation (1). Panel A focuses on funds targeting retail investors, and Panel B on funds targeting institutional investors. In both panels, columns 1 and 2 restrict the sample to funds that are advertised green while columns 3 and 4 focus on funds that are not advertised green. Columns 1 and 3 restrict the sample to funds with one to four Morningstar globe ratings, and columns 2 and 4 to five-globe funds. 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.

Panel A: Funds advertised to retail investors				
Globes	Advertised green		Not advertised green	
	1 to 4	5	1 to 4	5
GS _{t-1}	0.009 (0.075)	-0.126*** (0.038)	-0.002 (0.018)	0.039** (0.019)
GS _{t-2}	-0.086 (0.072)	-0.066*** (0.022)	-0.033* (0.018)	-0.024 (0.038)
CS _{t-1}	0.272** (0.106)	0.469** (0.207)	0.146*** (0.053)	-0.064 (0.084)
CS _{t-2}	0.023 (0.155)	0.603*** (0.205)	0.062 (0.069)	0.024 (0.131)
Controls	yes	yes	yes	yes
Adj. R2	0.814	0.935	0.883	0.809
Obs.	46	46	46	46

Panel B: Funds advertised to institutional investors				
Globes	Advertised green		Not advertised green	
	1 to 4	5	1 to 4	5
GS _{t-1}	-0.112*** (0.038)	-0.109*** (0.027)	0.005 (0.008)	0.030 (0.020)
GS _{t-2}	0.064 (0.050)	-0.003 (0.056)	-0.029 (0.021)	-0.013 (0.015)
CS _{t-1}	0.229* (0.134)	0.072 (0.159)	-0.004 (0.060)	0.239*** (0.073)
CS _{t-2}	0.135 (0.155)	0.097 (0.115)	0.017 (0.063)	-0.074 (0.088)
Controls	yes	yes	yes	yes
Adj. R2	0.692	0.455	0.107	0.740
Obs.	46	46	46	46

Table 8: 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 five-globe Morningstar rating ('5G').

Panel A: Funds advertised to retail investors					
	Flows				
	(1)	(2)	(3)	(4)	(5)
Ad. green	0.005***	0.005***	0.005***	0.005***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
GS _{t-1} × Ad. green	0.027	0.019	0.060	0.060	0.052
	(0.030)	(0.031)	(0.038)	(0.038)	(0.039)
GS _{t-1} × Ad. green × 5G			-0.106**	-0.105**	-0.106**
			(0.051)	(0.051)	(0.051)
GS _{t-2} × Ad. green	-0.052	-0.051	-0.052*	-0.052*	-0.050*
	(0.035)	(0.035)	(0.030)	(0.031)	(0.030)
GS _{t-3} × Ad. green				-0.005	
				(0.032)	
CS _{t-1} × Ad. green		0.227			0.227**
		(0.153)			(0.104)
CS _{t-2} × Ad. green		0.057			0.056
		(0.120)			(0.100)
cat-year-FE	yes	yes	yes	yes	yes
Controls	yes	yes	yes	yes	yes
R ²	0.0669	0.0670	0.0669	0.0669	0.0670
Observations	148,919	148,919	148,919	148,919	148,919

Panel B: Funds advertised to institutional investors					
	Flows				
	(1)	(2)	(3)	(4)	(5)
Ad. green	0.004***	0.004***	0.004***	0.004***	0.004***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
GS _{t-1} × Ad. green	-0.050	-0.044	-0.082***	-0.082***	-0.076**
	(0.035)	(0.036)	(0.031)	(0.031)	(0.032)
GS _{t-1} × Ad. green × 5G			0.080	0.078	0.079
			(0.056)	(0.056)	(0.056)
GS _{t-2} × Ad. green	0.052	0.045	0.053*	0.051*	0.046
	(0.062)	(0.061)	(0.031)	(0.031)	(0.031)
GS _{t-3} × Ad. green				-0.020	
				(0.027)	
CS _{t-1} × Ad. green		-0.148			-0.148
		(0.112)			(0.097)
CS _{t-2} × Ad. green		0.112			0.110
		(0.124)			(0.090)
cat-year-FE	yes	yes	yes	yes	yes
Controls	yes	yes	yes	yes	yes
R ²	0.0669	0.0670	0.0669	0.0669	0.0670
Observations	148,919	148,919	148,919	148,919	148,919

Table 9: 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

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


```
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