A Greenwashing Index *

Elise Gourier
ESSEC Business School and CEPR

Hélène Mathurin
ESSEC Business School

October 2025

Abstract

We construct a novel, news-based index that captures the salience of greenwashing in media coverage. The index reveals several episodes of heightened attention to greenwashing, with the most recent surge driven by growing skepticism toward the financial sector—particularly regarding sustainable assets. We show that this increase in greenwashing salience tends to follow a period of heightened confusion about how to measure corporate sustainability and evaluate the risk and returns of sustainable investments. Furthermore, this peak in greenwashing salience is followed by a rise in negative sentiment toward ESG-related issues. Finally, we examine mutual fund flow responses to changes in greenwashing salience and document differences in investor behavior: institutional investors increase allocations to funds with high sustainability ratings—consistent with elevated signaling concerns—while retail investors respond primarily to past fund performance.

JEL Classification: G11, G12, Q51

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

---

*Email: elise.gourier@essec.edu, helene.mathurin@essec.edu. We would like to thank Svetlana Bryzgalova, Alan de Genaro, Javier Gil-Bazo, Samuel Hartzmark, Christian Julliard, Elisabeth Kempf, Johannes Klausmann, Stefano Lovo, Nadya Malenko, Alessandro Melone, Ludovic Phalippou, Alberto Quaini, Zacharias Sautner (discussant), Gustavo Schwenkler (discussant), Laura Starks and Koen van den Bosch (discussant) for feedback and suggestions on our project. We also thank the participants of seminars at ESCP, the University of Liverpool, SKEMA, Vienna Graduate School of Finance, HEC-HKUST Sustainable Finance Webinar and participants to the 2025 EFA Conference, the Sciences-Po HEC Summer Workshop in Sustainable Finance, NEOMA Conference on Sustainable Finance, the Spring Finance Workshop, the 17th Financial Risks International Forum, the 8th ESSEC Asset Pricing Breakfast, and the 5th International Workshop in Financial Econometrics for comments. This research has been supported by the JCJC ANR grant ANR-22-CE26-0010-01 ("GreenLies") and a grant from the "Institut Europlace de Finance and Labex Louis Bachelier". Finally, we would like to thank the research assistants who helped with this project.

1 Introduction

The mounting pressure from activists, investors and government agencies to tackle climate change has led some firms to provide misleading information about the environmental soundness of their practices. This phenomenon is often referred to as 'greenwashing'. Greenwashing plausibly hinders the measurement of climate risks and the actions to reduce these risks. It has thus recently been the focus of regulators and policy makers, and the subject of a heated policy debate.[1] In this paper, we build an index of greenwashing salience. We show that greenwashing has become increasingly salient since 2018, to reach a peak in 2021. This surge was driven by suspicions in the financial sector. Furthermore, we show that this episode of greenwashing was accompanied by heightened media-reported confusion over ESG issues—particularly regarding the measurement of corporate sustainability and the performance evaluation of sustainable assets. It was subsequently followed by a rise in anti-ESG sentiment, which remains elevated to date. Finally, we document the reaction of retail and institutional investors to increases in greenwashing salience.

Our greenwashing index measures the fraction of news articles alluding to firms' greenwashing. We apply an advanced Natural Language Processing algorithm to the history of paper-based Wall Street Journal articles between January 1986 and June 2025 (nearly one million articles). The algorithm is trained to identify the articles that allude to greenwashing. Identifying these articles is challenging for several reasons. First, the word "greenwashing" became widely used only recently: it was used in fewer than 10 articles before 2007 and in fewer than 35 articles before 2020. In contrast, it was used in nearly 50 articles between January 2020 and June 2025. Our algorithm is able to learn the words or combinations of words that characterize articles related to greenwashing, even if they do not contain the word itself. Second, the data set is large and unlabelled. Third, it is highly imbalanced: the fraction of all Wall Street Journal articles that are greenwashing-related is small. To make these issues less critical, we proceed in two steps. First, we identify articles related to firms' sustainability actions and claims, and second, we pinpoint, among these articles, those that mention greenwashing (or equivalent terms). For each step, we use a Large Language Model (LLMs) to enlarge a manually labelled set of articles, and train standard classifiers to learn the patterns of these articles. Our manually labelled database guarantees that we can validate the output of the LLM. Furthermore, the use of a combination of LLM and standard machine learning algorithm ensures that our method is low-cost, easily scalable

[1] New rules have recently been proposed by the Securities and Exchange Commission in the U.S., and by the European Parliament and Council, to reduce greenwashing. We provide details on the actions and regulations adopted in the U.S. and in Europe over time in Section 2.3.

and interpretable. The algorithm yields as output a classification rule of all articles.

The out-of-sample performance of our algorithm critically depends on how words are embedded into a numerical vector. The simplest approach is to use a bag-of-words representation, which counts the frequency of each word or group of words in an article. We find that in the first step, articles related to firms' sustainability are well identified using this approach. However, it does not allow for the precise identification of greenwashing articles in the second step. Indeed, there is no list of words (besides the word "greenwashing" itself) that uniquely characterizes greenwashing articles, making this step more challenging. We use OpenAI embeddings, which account for the context of words in their numerical representation, and show that they provide 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 yields the best trade-off between parsimony and performance.

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

Articles classified as greenwashing articles represent a small part, on average 0.2%, of all articles. Although one could see this number as an indication that greenwashing does not matter, it is important to note that this fraction has reached peaks near 1% twice in the past ten years. In comparison, there is on average 4.1% of Wall Street Journal articles that are related to firm sustainability. Greenwashing has thus accounted for a substantial amount of articles on firm sustainability in the recent years.

We find that there are three main waves of greenwashing salience in the media. The first wave is from 1990 to 1992, the second from 2015 to 2017 and the third from 2018 to 2022. The first wave 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). The second peak was triggered by the Volkswagen scandal in 2015. The last wave, from 2018, coincided with greenwashing accusations towards investment firms (e.g., Blackrock, Deutsche Bank). In order to understand better the nature of the media articles about greenwashing, we list the firms that are mentioned in these articles. These firms include names in the oil & gas industry such as Exxon Mobil and Chevron, automobile companies,

[2] It is possible that greenwashing 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.

and investment firms (Blackrock, JPMorgan Chase, Goldman Sachs etc.). The latter firms are more present in greenwashing articles published after 2018. We further perform topic analysis on the set of greenwashing articles, to gain a more granular understanding of the contents of these 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 topics that relate to the retail, automobile, energy and financial sectors. Unlike other topics, the Financial sector topic has been growing since 2018, and clearly drives the most recent peak in greenwashing salience. Overall, it accounts for 6% of the text in greenwashing articles, but this ratio reaches nearly 20% in 2021. It is composed of two subtopics: Asset Management (specifically, ESG funds), and Green Bonds. The first topic dominates the second, and is the one that drove up the greenwashing index in 2021. Another noteworthy topic is the Social topic, which has been growing from 5% of greenwashing text contents in 2020 to 15% in 2025.

Given the importance of the financial sector in the last greenwashing wave, we next study the relation between this greenwashing wave, and two concepts that have often been associated with greenwashing in the media: the confusion around sustainability and the ESG backlash. Two sources of confusion have been widely reported. The first one relates to the metrics used to quantify firms' sustainability, and the second to the expected performance of sustainable assets. These two sources of confusion are linked, and have been the subject of heated debates both in academia and in industry. We build an ESG confusion index which captures this debate in the news, and find that it exhibits similar time series properties as a greenwashing-in-finance subindex, which only includes greenwashing articles that are related to the financial sector. Specifically, a regression of unexpected changes in the greenwashing-in-finance subindex on the ESG confusion index reveals that greenwashing tends to follow ESG confusion, and confirms that the two concepts are positively correlated.

Similarly, we build an index of negative ESG sentiment, to capture the ESG backlash reported in the news. This sentiment index is built from opinion pieces of the Wall Street Journal. Opinion pieces were excluded from the construction of the greenwashing index, but contain useful information on the general opinion towards specific topics. Our sentiment index counts the fraction of firm sustainability articles that are negative about firms' ESG actions. We find that there is a large and statistically significant increase in anti-ESG sentiment following the 2021 peak in greenwashing salience. This anti-ESG sentiment remains high to date, whereas ESG confusion and greenwashing salience have remained low since mid-2022. This substitution of greenwashing by the ESG backlash suggests that greenwashing may not have disappeared despite being less salient, but may instead be covered up by the ESG backlash.

Finally, we study investors' reaction to increases in greenwashing salience. Specifically, we examine the response of market demand in mutual funds and their flow-performance sensitivity, as a function of greenwashing salience and across funds with varying sustainability levels. Our analysis starts in August 2018 and ends in June 2025. We find, using panel regressions, that retail and institutional investors react differently to heightered greenwashing salience. Institutional investors react to increased greenwashing salience by investing more in funds with the highest Morningstar sustainability rating, and less in funds with the lowest rating. This result is in line with the need to signal attention to values, Starks (2023), in times of high greenwashing salience. For these investors, the flow-performance sensitivity does not change significantly in times of high greenwashing salience. Retail investors, in turn, respond to increases in greenwashing salience by investing more in low-rated funds with high past returns. These differences indicate that institutional investors may focus more on reputation and signalling in times of high greenwashing salience, while retail investors focus more on performance.

We build on a literature that uses text analysis to quantify climate change 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.

Instead of building a broad climate change risk index, we focus on companies' actions, and on how they integrate environmental and social considerations into their corporate culture and financial decision-making. Identifying media articles that mention these firms' actions is for us an intermediary step in the construction of the greenwashing index.

All the papers previously cited abstract from the existence of greenwashing when measuring climate risk. Cartellier, Tankov, and Zerbib (2024) show that in equilibrium, firms face a trade-off between improving their fundamental environmental value (through, e.g. emission reductions) and improving their environmental communications (using, e.g., greenwashing). Inderst and Opp (2024) study, in equilibrium, the conditions under which a taxonomy of ESG investments would prevent greenwashing.

Several recent papers have tried to identify greenwashing firms. Bingler, Kraus, Leippold, and Webersinke (2022a) implement a BERT model that is trained on climate resources, to determine whether companies' TCFD disclosures are mostly cheap talk or not. In a parallel paper, the same authors (Bingler, Kraus, Leippold, and Webersinke (2022b)) analyze companies' annual reports. Finally, recent papers have tried to highlight greenwashing in the financial sector. Brandon, Glossner, Krueger, Matos, and Steffen (2022) find that signatories of the Principles for Responsible Investment outside the United States have better ESG ratings than non-signatories, but that this no longer holds in the United States. Kacperczyk and Peydró (2022) find evidence of banks' greenwashing. Parise and Rubin (2023) show that the timing of purchases and sales of sustainable assets by mutual funds is chosen to manipulate sustainability ratings. Heath, Macciocchi, Michaely, and Ringgenberg (2023) find that socially responsible investment funds do not follow on their promise to impact. Along the same lines, Atta-Darkua, Glossner, Krueger, and Matos (2023) raise doubts about the effectiveness of investor-led initiatives in reducing corporate emissions and helping an all-economy transition to "green the planet". Dumitrescu, Gil-Bazo, and Zhou (2023) propose a definition of greenwashing for mutual funds, and find that according to their definition, a third of self-labelled ESG funds are greenwashing. Recently, Cao et al. (2025) use Large Language Models to identify mutual funds engaging in greenwashing, and find that these funds charge higher fees while attracting more flows from investors.

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

Finally, we contribute to a large literature analyzing how media coverage affects financial markets, including Tetlock (2007), Dougal et al. (2012), and Carlin et al. (2014), and recently Baker et al. (2025), among others.

2 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.

2.1 Definition

The term greenwashing was originally coined by environmentalist Jay Westerveld in 1986 in an essay criticizing the hotel industry's superficial eco-friendly messaging—specifically the common practice of asking guests to reuse towels under the pretense of environmental concern. Since then, the term has come to denote a broader phenomenon: the act of misleading stakeholders by portraying a product, service, or organization as more environmentally sustainable than it truly is. One early regulatory response came in the aftermath of the 1989 Exxon Valdez oil spill, when the U.S. Securities and Exchange Commission (SEC) charged Exxon with making materially false and misleading environmental claims on the environmental impact of the spill.

Over time, the definition of greenwashing has expanded beyond anecdotal corporate behavior to encompass systematic misrepresentation within financial markets and broader firm and product marketing. Yet despite the growing importance of ESG (Environmental, Social, and Governance) claims, there is still no universally accepted legal definition of greenwashing.

In response to increasing concerns about greenwashing, several regulators have issued working definitions aimed at specific sectors. In 2020, the E.U. Taxonomy Regulation defines greenwashing, in the context of sustainable finance, 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". In January 2023, the European Securities and Markets Authority (ESMA) 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 progress report on greenwashing released in June 2024 (European Securities and Markets Authority (2023)), the ESMA defines it as "a practice where sustainability-related statements, declarations, actions, or communications do not clearly and fairly reflect the underlying sustainability profile of an entity, a financial product or financial service. This practice may be misleading to consumers, investors, or other market participants." The report highlights that marketing materials, labels and voluntary reporting are most exposed to greenwashing risk, and identifies several ways that can be used for greenwashing, such as cherry-picking, omission, ambiguity, empty claims (including exaggeration), and misleading use of environmental, social and governance (ESG) terminology.

With the increasing popularity of the ESG framework, the definition of greenwashing has in the past years often been broadened to misleading claims on ESG issues. We do so in this paper as well.

2.2 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. 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 2.3 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' growing 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, leading to selective disclosure and opportunistic reporting.

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

2.3 Actions and Regulations

To combat greenwashing, governments and regulatory bodies have taken a series of steps to address greenwashing, especially through enhanced disclosure requirements and labeling rules.

The U.K. was probably the first country to implement disclosure requirements: premium listed companies and regulated financial firms (banks, insurers and asset managers) have been required by the Financial Conduct Authority (FCA) to disclose under the TCFD framework since 2021. The U.K. government has also amended the Companies Act of 2006 to implement the TCFD recommendations for large and listed companies, since 2022. Furthermore, the Sustainability Disclosure Requirements were implemented in May 2024, with naming and marketing rules applying from December 2024—including an anti-greenwashing rule for ESG-labelled products. However, in April 2025, the FCA paused its further rollout for wealth managers. Finally, since 2024, the Competition and Markets Authority has exercised new enforcement powers to pursue misleading environmental claims across business sectors, with fines up to 10% of global turnover, under the 2024 Competition & Consumers Act. In July 2025, the UK Treasury withdrew plans for a domestic green taxonomy.

In France, companies that have more than 500 employees and more than 100 million euros in turnover or balance sheet total have had to file a document known as Greenhouse Gas Emission Balance ("Bilan des Emissions de Gaz à Effet de Serre") every four year since 2010, as well as an annual Declaration of Extra-Financial Performance (DPEF) since 2017. These documents were partially in compliance with TCFD recommendations. An ordonnance was signed in December 2023 to strengthen these reporting requirements in compliance with the E.U. Taxonomy, the French Sustainability Reporting Standards (FSRS) and thereby the TCFD recommendations.

In the EU, a series of actions have been taken in the past few years to ensure better disclosure of firms. The Sustainable Finance Disclosure Regulation (SFDR) aimed to improve transparency in the market for sustainable investment products, by requiring investment firms to make detailed disclosures in relation to the products and services with environmental or social characteristics. Initial requirements became applicable in 2021. The Taxonomy

Regulation that started coming into force in 2022 further required firms to disclose how aligned their products or services were with what is considered "green" under the taxonomy. Furthermore, a Regulation on European green bonds was proposed in 2021. It was adopted in October 2023, and started being applied in October 2024. On March 22, 2023, the European Commission published its proposed Directive on Green Claims.[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. The European Commission however announced in June 2025 that it intended to withdraw the proposal.

In 2023, a new Corporate Sustainability Reporting Directive (CSRD) also entered into force. It introduced more detailed and standardized reporting requirements compared to the previous Non-Financial Reporting Directive (NFRD). It applies to all public interest entities operating in the E.U., such as listed companies, banks, and insurance companies. It covers around 49,000 companies and requires them to report on ESG topics using the European Sustainability Reporting Standards (ESRS), which are based on the TCFD recommendations. Moreover, in September 2023, the European Parliament and the European Council reached a provisional agreement on new rules to ban misleading advertisements and provide consumers with better product information.[4]

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 final report was released on June 4, 2024, by the European Securities and Markets Authority (2023).

As of November 2024, the European Union furthermore adopted the Green Bond Certification Framework, a set of rules for independent certification of carbon removal and green bond products, aimed at avoiding greenwashing in carbon credits and green bonds.

The EU also took steps to address greenwashing outside the financial sector, with the EU Empowering Consumers Directive (Directive 2024/825), effective since April 2024. This directive amends the Unfair Commercial Practices Directive, banning generic environmental claims (e.g., "eco-friendly," "green") unless substantiated by recognized evidence. Member

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

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

States must transpose it into national law by 2026.

In the US, until recently, there was no specific federal law that regulated greenwashing. There were several laws that could be used to address it, such as the Federal Trade Commission Act (FTC Act), which prohibits unfair or deceptive acts or practices in commerce. The FTC issued guidelines for environmental marketing claims, known as the Green Guides, which provide examples of how to make truthful and substantiated claims about the environmental attributes of products or services. The FTC can take enforcement actions against companies that violate the FTC Act or the Green Guides, such as issuing cease and desist orders, imposing civil penalties, or requiring corrective advertising. In 2021, for example, the FTC fined Walmart and Kohl's, a combined $5.5m for mislabeling rayon as "sustainable" bamboo. Additionally, some states, e.g., California, have enacted their own laws to prevent greenwashing. Finally, in September 2023, the U.S. Securities and Exchange Commission (SEC) amended the "Names Rule" initially adopted in 2001, which regulates the names of registered funds to ensure they are marketed to investors truthfully. The rule ensures that ESG funds have at least 80% of their assets be ESG assets, and improves transparency on the criteria chosen to select the assets.[5]

In 2024, the SEC adopted final rules requiring large public companies to disclose climate risks and GHG emissions (scopes 1 and 2), starting public filings in 2025. These rules were later scaled back, excluding mandatory scope 3 emissions. Also in 2024, the FTC's long-anticipated update to the Green Guides, which aimed to help define permissible environmental marketing, was delayed indefinitely.

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

3 Measuring Greenwashing Salience

Greenwashing is by nature challenging to measure. We do not claim to be able to identify individual incidents of greenwashing. Our goal, instead, is to measure overall greenwashing salience, proxied by the quantity of media articles alluding 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 greenwashing salience. We show that our algorithm provides a good trade-off between out-of-sample performance and parsimony.

3.1 Data

We use the history of daily paper-based Wall Street Journal articles from January 1986 to June 2025.[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 899,422 articles.

3.2 Algorithm

The challenge of dealing with our dataset is threefold: it is large, unlabelled, and most articles are about topics other than sustainability, so it is highly unbalanced. Furthermore, there may be many articles mentioning greenwashing without using the "greenwashing" word specifically.

We identify greenwashing articles in two steps, which correspond to the two components of greenwashing: to greenwash, firms need to make an action that is related to environmental or social responsibility, and this action needs to purposely mislead the consumer. Following this logic, we first identify articles that mention firms' actions related to climate and social risks. This set includes, for example, articles that describe firms' green products or disclosure, as well as articles that mention firms' environmental and social claims. We refer to these

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

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

articles as *sustainability* articles and filter out all the other articles. Second, we identify, among these sustainability articles, those that allude to greenwashing. Specifically, we define greenwashing as exaggerated, misleading, mislabelled, false or unsubstantiated companies' claims or marketing about their environmental and social actions or green products.

A set of manually labelled articles was created for the training of the algorithm, as well as in order to fine-tune its hyperparameters and to test its out-of-sample performance. We recruited around 40 research assistants, who each received the task to read 150 articles, and to evaluate whether these articles were sustainability and greenwashing articles. In Appendix A, we provide details on how these articles were selected, as well as on the instructions given to the assistants, and on the processing of their outputs.

A total of 1593 articles were read by at least two students. Among these articles, 1054 were classified as sustainability articles, and 138 as greenwashing articles. As the resulting sets of sustainability and greenwashing articles remain small, we use Large Language Models to extend these sets. In the first step, we use GPT 4o-mini to further classify 48500 randomly sampled articles. We followed an extensive prompt engineering process to achieve a high match between the manually labelled database and the GPT classification. With the chosen prompt, the output of GPT matches the manual classification for 94% of the articles. Furthermore, 95% of the climate-labelled articles were identified by GPT, and out of all articles that were climate-labelled by GPT, 90% were also climate-labelled by the assistants. The final labelled database, for the first step, contains 7111 sustainability articles.

In the second step, using GPT 4o-mini, we did not manage to achieve a high match with the manually labelled database. We thus used the more costly GPT-4 Turbo model to further classify 11000 articles selected randomly among those classified as sustainability articles by our algorithm in the first step. The output of GPT matched the manual classification for 90% of the articles. Furthermore, 62% of the greenwashing-labelled articles were identified by GPT, and out of all articles that were greenwashing-labelled by GPT, 70% were also greenwashing-labelled by the assistants. Our final labelled database for the second step contains 854 greenwashing articles. We describe the prompt engineering process for the two steps in Appendix B.

For both the sustainability and greenwashing classification steps, the training sets are built to achieve a balanced dataset, so as to ensure that the algorithm has enough articles from the minority class (sustainability articles in the first step, and greenwashing articles in the second step) to be able to learn the patterns of these articles. As the algorithm learns better if the topic to identify is discussed in a higher proportion of the article, in the first step, we further include in the training set solely sustainability articles for which GPT evaluates the sustainability-related content to be over 20% of the article. The resulting training set contains 1132 articles of each class. Due to the limited number of greenwashing articles in the second step, we cannot apply the same rule. The training set contains 454 articles of each class.

In the first step, as in the training set, we select in the validation and test sets articles for which GPT identified more than 20% of contents being sustainability related. In total, there are 1732 such articles in the manually labelled database. In the second step, we do not apply this filter. The validation and test sets are built using half of the remaining sustainability and greenwashing articles, and seven times as many articles of the non-sustainability and non-greenwashing articles. The ratio of seven is chosen so as to ensure comparability between the validation and the test sets, and the full database. A summary of the training, validation and test sets is provided in Table 1.

For each step, we train three classifiers to learn the differences between the articles in the two classes. This approach is similar to what is done by King, Lam, and Roberts (2017), and by Sautner, van Lent, Vilkov, and Zhang (2023a) at the level of the sentence, to identify climate mentions in earning calls. The 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 context of words. We also use transformer-based embeddings, which fully take into account the context of words in their numerical representation. Specifically, we use the OpenAI embeddings from the model "text-embedding-ada-002", which are trained on publicly available and licensed data (including large-scale internet text, academic articles, books, Wikipedia, code, and other diverse data sources) up to early 2022.

In the first step of the classification, we choose a Naive Bayes classifier, a random forest classifier and an eXtreme Gradient Boosting (XGBoost) classifier. These choices are natural given the nature of our dataset. The Naive Bayes classifier relies on the unrealistic assumption of conditional independence of the features, but it is computationally inexpensive and has been shown to perform well with large datasets. With bag-of-words embeddings, we use the Multinomial Naive Bayes classifier and with the OpenAI embeddings, the Gaussian Naive Bayes classifier. This change is due to the fact that the Multinomial Naive Bayes algorithm can only take as input an integer representation of documents and OpenAI provides a float representation instead. Tree methods are flexible classifiers, able to represent decision boundaries with any shape, but they are more computationally expensive, and can be prone to overfitting. In the second step of the classification, we replace the Naive Bayes classifier by a Support Vector classifier, which would have been too computationally expensive in the first step, but is subject to less assumptions and thus more flexible. These choices of classifiers are close to the choices of Sautner, van Lent, Vilkov, and Zhang (2023a).

We estimate the probability of each article in the full dataset to be a sustainability articles, and the probability of each sustainability article to be a greenwashing article. As each classifier assigns a probability to each article of being in each class, we aggregate these probabilities by averaging them. Each article is classified as sustainability, or greenwashing, if the resulting probability exceeds a given threshold, fine-tuned by maximizing the F1-score over the validation set of manually labelled articles.[8] We then refine the set of sustainability and greenwashing articles, by reclassifying articles with a probability between 50% and the classifier's threshold, using GPT, and the same prompt as the one that was used for the creation of the manually labelled set.

Over the 899,422 articles in the database, 36,843 are classified as sustainability articles, i.e., 4.1% of the database. 1948 are classified as greenwashing articles, i.e., 0.2% of the database, and 5.3% of the sustainability articles.

3.3 What do we capture?

Below are two extracts of news articles that we classify as greenwashing:

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

[8] The F1-score is the harmonic mean of the precision and recall, where the precision refers to the fraction of all the positive predicted (i.e. all articles classified as sustainability or greenwashing) that are true positives (i.e., rightly predicted), and the recall is the fraction of all positives (i.e., all sustainability or greenwashing 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 sustainability or greenwashing wrongly, that is, the volume of sustainability or greenwashing articles is over-estimated. A low recall means that few of the sustainability or greenwashing articles are classified as such, that is, the volume of sustainability or greenwashing articles is under-estimated. Optimizing the F1-score allows reaching a trade-off between these two situations. The accuracy, which measures the fraction of predictions which are right, is often the criterium that is maximized when fine-tuning a classifier's hyperparameters. As our dataset is highly unbalanced, and contains much more non-sustainability or non-greenwashing articles than sustainability or greenwashing articles, a high accuracy can be achieved by classifying all articles as sustainability or greenwashing. When working with such an unbalanced dataset, maximizing the F1-score is therefore a better method.

2021-09-27. In August, The Wall Street Journal reported that DWS struggled to define and implement an ESG strategy, at times painting a rosier-than-reality picture to investors, according to Ms. Fixler and internal emails and presentations.

What makes greenwashing difficult to identify in articles is the variety of terms that can be used to describe it. Furthermore, many of these terms, on their own, are not sufficient to characterize greenwashing, and need to be read as part of their context. Such importance of the context makes it natural to use Large Language Models, in complement to other traditional machine learning methods, in our classification.

3.4 Out-of-sample Performance

We evaluate the out-of-sample performance of each classification step using the manually labelled articles in the testing set. The first line of the two cells in Table 2 lists various measures of performance that each step of our algorithm achieves. All measures of performance are strikingly high for the first step of the algorithm, which identifies sustainability articles. Above 95% of all articles are classified in the right category (accuracy). About 83% of all articles classified as sustainability articles are indeed in this category (precision). Nearly 83% of all sustainability articles are well categorized (recall). These two statistics yield an F1-score (harmonic mean of the precision and recall) of 83%. 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 above 90%, confirming the high quality of the identification of sustainability articles.

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

We compare the performance of our algorithm to the ones of Engle, Giglio, Kelly, Lee, and Stroebel (2020) and Sautner, van Lent, Vilkov, and Zhang (2023a). More details on our implementation of these methods are provided in Appendix C. We are able to classify around 7% more articles in the right category than the algorithms of these two papers, which in our tests achieve comparable accuracy. In particular, our classification of climate-related articles is more precise (30% gain) and achieves a higher recall (30% gain compared to Sautner, van Lent, Vilkov, and Zhang (2023a), and 40% gain compared to Engle, Giglio, Kelly, Lee, and Stroebel (2020)).

The statistics of the second step of the algorithm are displayed in the bottom half of the table. Similar to the first step, our method displays a high accuracy of above 95%, which means that 95% of the articles are classified in the right category. As greenwashing articles only represent a small fraction of all articles[9], this number would however be achieved by classifying all articles as non-greenwashing. It is therefore reassuring to see that the precision is nearly 80%, meaning that very few articles classified as greenwashing are wrongly classified. The recall is however only of 52%, meaning that out of all greenwashing articles, our classifier is missing almost half of them, i.e., we underestimate greenwashing salience. This statistic shows that identifying greenwashing articles is much more challenging than identifying sustainability articles.

These statistics are obtained using OpenAI as embedding method. Using a simple bag-of-words approach indeed a significantly lower precision of 56%. Large Language Models therefore allow better capturing greenwashing concepts, so that less articles are wrongly classified. This finding is interesting to contrast to the first step of the algorithm, in which a bag-of-words approach yielded comparable results to OpenAI. This difference between the two steps sheds some light on why identifying greenwashing articles is more difficult than identifying sustainability artices. Indeed, whereas a fair number of bigrams are specific to firm's environmental and social actions (e.g., emission reduction, Kyoto protocole etc.), few words and bigrams are specific to greenwashing. Accounting for the context of words therefore brings value to the second step whereas it does not to the first step. This result is confirmed by the performance of the methods of Engle, Giglio, Kelly, Lee, and Stroebel (2020) and Sautner, van Lent, Vilkov, and Zhang (2023a). the former achieves a high recall of 66%, but its precision is only at 8%, implying that the number of greenwashing articles is overestimated. The latter, in contrast, achieves a high precision of 75%, but a recall inferior to 2%, indicating that it misses almost all of the greenwashing articles.

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

4 Greenwashing Index

Using the classification algorithm described in Section 3, we identify the articles of the Wall Street Journal that are sustainability and greenwashing articles. In this section, we analyze the evolution of the quantity of these articles, and describe the resulting greenwashing index. We investigate which firms, industries and topics are most related to greenwashing.

4.1 Firm Sustainability Index

Figure 2 displays the percentage of sustainability articles each month, out of all Wall Street Journal articles. This index is a proxy for the salience of firms' sustainability actions and claims. It has been increasing since the end of the nineties and exhibits peaks around climate-related events, such as the UN Climate Change conferences, and following calls for more disclosure rules. Interestingly, the peaks seem to be contemporaneous with financial 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), at the end of 2021. Sustainability articles then reached 10% of all articles published in the Wall Street Journal. The years from 2010 to 2020 show lower salience of firm sustainability, even though the index is higher than in the pre-2000 period.

The overall shape of our index of firm sustainability is related to the indices obtained by Engle, Giglio, Kelly, Lee, and Stroebel (2020); Ardia, Bluteau, Boudt, and Inghelbrecht (2022) and Faccini, Matin, and Skiadopoulos (2023). Table 3 displays the correlation matrix of these monthly indices, and shows that the index of Ardia, Bluteau, Boudt, and Inghelbrecht (2022) is the closest to ours, with a correlation of 70%. Faccini, Matin, and Skiadopoulos (2023) compute separate indices for policy risk, summit-related risk, risk of global warming and of natural disasters. Our index is the closest to their policy risk, with a correlation of nearly 60%. Beyond methodological differences in the construction of the index (different databases, algorithms and time periods), what we capture is not exactly the same as these indices. Specifically, we focus on firms' actions and claims. That is, we do not aim to capture either physical or transition risk, but rather the response of firms to these risks. As firms mostly respond to transition risks, it is not surprising that the index of Faccini, Matin, and Skiadopoulos (2023) that is the closest to ours is their climate policy index.

4.2 Greenwashing Index

Figure 3 displays our greenwashing index, built as the percentage of articles that are greenwashing articles, each month. The correlation between the firm sustainability index and the greenwashing index is of 32%. This correlation drops to 20% if we clean the firm sustainability index by removing from it all articles that are greenwashing articles. These correlations thus indicate that there is comovement between the salience of firm sustainability and that of greenwashing.

Greenwashing articles represent less than 1% of all Wall Street Journal articles, i.e., around one article out of twenty articles on firm sustainability. There are occasional peaks, which coincide with various companies being accused of greenwashing, and sometimes charged by the SEC. The first peak of greenwashing salience starts in 1990. Exxon was accused of trying to look greener than it actually was following the 1989 Exxon Valdez oil spill in Alaska. Around the same time, Procter & Gamble was criticized, first because of claims on disposable diapers, and second for lobbying against environmental measures while showing off their green credentials. Toyota and General Motors were criticized for releasing ads which implied that their cars were protecting the environment. In 1991, greenwashing articles reached nearly 0.8% of all 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, to reach a second peak. 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. Shortly after, in April 2008, Shell was identified as one of the worst offenders for overstated environmental marketing. Its ad campaign depicting a large oil refinery with flowers emanating from its chimneys, alongside the slogan "we use our waste CO2 to grow flowers", was formally ruled misleading by the UK's Advertising Standards Authority.

The next large peak was triggered by the 2015 Volkswagen emissions scandal, which shook the credibility of the entire car industry. Greenwashing articles then accounted for above 1% of Wall Street Journal articles. From 2018, a new greenwashing wave started and reached its peak with greenwashing accusations towards investment firms (e.g., Blackrock, Deutsche Bank). Since the end of 2022, the greenwashing index has remained low.

4.3 Greenwashing-related firms and industries

Greenwashing salience has gone up and down since the 1990s. We now investigate which firms have been driving these peaks. Specifically, we extract the names of companies that are mentioned in the greenwashing articles. 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 articles. The second step perfoms Entity Linking: it links the different ways to refer to the same company.

For simplicity, 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.[10] 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.

These two steps allow us to count how many times each company is mentioned in the set of greenwashing articles. Figure 4, Panel A, shows in a word cloud the companies that are mentioned. The size of each company's name increases with the number of mentions. The usual suspects are on the figure: Volkswagen, Ford and General Motors in the automobile industry, Chevron, Shell and ExxonMobil in the oil and gaz industry, among others.

Panel B of Figure 4 displays the word cloud of companies in greenwashing articles since 2018, i.e., since the beginning of the last peak in greenwashing. The word cloud now features several firms from the financial sector: Blackrock, JPMorgan Chase, Goldman Sachs, Bank of America, Deutsche Bank, Citigroup etc. Volkswagen is still mentioned frequently, as the emission cheating scandal had multiple developments, even after Volkswagen pleaded guilty. For example, Volkswagen agreed to a $9 billion settlement in Germany with institutional investors in May 2020.

4.4 Greenwashing Topics

As the definition in Section 2.1 indicates, companies can greenwash in many different ways. In order to better understand the nature of greenwashing over time, we perform a topic

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

analysis. We use an extension of the standard Latent Dirichlet Allocation (LDA) method introduced by Blei, Ng, and Jordan (2003), which allows, through human validation, to overcome some well-known issues of the LDA.

The LDA method is a probabilistic model, i.e., it posits a set of latent topics, multinomial distributions over words, and assumes that each document can be described as a mixture of these topics. It is an unsupervised learning method, which only uses unlabelled data. While it has been successfully used as a tool to explore the topics of a corpus of documents, it has also been shown to often inadvertently create multiple topics with similar content and combine different themes into a single topic (Chang, Boyd-Graber, Gerrish, Wang, and Blei (2009); Newman, Bonilla, and Buntine (2011); Morstatter and Liu (2016)). The keyword Assisted Topic Model (keyATM) introduced by Eshima, Imai, and Sasaki (2023) is a semi-supervised topic model[11], 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.

Starting from topics identified using a standard LDA approach and after several iterations, we obtain a list of 22 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 2.2. The topic is characterized by keywords such as "report", "standard" and "data". The second most important topic is the one that we labelled "Corporation". It contains keywords such as "executive", "board" and "investor", in line with ESG considerations having become a topic of importance at the high management level of firms. The next topic, labelled "Environmentalists", contains keywords showing links between activist groups, the public, government officials and firms. Further below but very related is the "Lawsuit" topic, in line with several companies having been charged and fined for greenwashing. The role of governments in anti-greenwashing actions is also featured in the topic analysis, with a "Politics" topic that contains keywords such that "state", "president", "administration" and "bill".

The topic analysis further shows that we do not only capture greenwashing related to green actions, but also social-washing. The "Social" topic accounts for 9% of the text and

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

characterized by keywords such as "people", "safety", and "workers".

About a third of the topics are specific to industries. The retail industry is most prominent with 10% of the text in greenwashing articles. It is followed by the automobile industry (6%), the energy sector, composed of a fossil fuels topic (3%) and a topic on alternative energies (3%) with keywords such as "plant", "solar", "electricity", and "dam". The financial industry amounts to 4% of the text in greenwashing articles. Two main financial topics are identified, namely asset management, with keywords such as "fund", "investment", "stock" and "asset", and bonds. The remaining industries that come out of the topic analysis are the construction, food and tobacco industries. They represent, respectively, 2%, 2% and 1% of the text in greenwashing articles.

One of the strengths of the keyATM method is that it provides as output a dynamic representation of the topics' importance over time. Figure 5 describes the evolution of the sector-based topics. Whereas some industries' presence in greenwashing articles remains rather constant over time -this is the case of the Retail and Energy sectors-, the Automobile and Financial industries have been clearly responsible for the last two waves of greenwashing, in 2015 and 2021. In fact, the Financial topics came to replace the Automobile topic, propelled by the Volkswagen scandal.

Figure 6 further decomposes the Financial and Energy sector topics, showing that the peak in the Financial topic was by and large due to an increase in greenwashing salience related to asset management. 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 5.3, some of these funds have received poor Morningstar globe ratings, indicating that their holdings are not aligned with Morningstar's sustainability criteria. The Bond topic, instead, grew around 2008, around the introduction of the first green bonds. It then remained stable over time. The concerns of greenwashing in green bonds were mostly due to the lack of standards in the definition of green projects and the absence of measures of their impact. In the Energy sector, the evolution of the two subtopics, "Fossil Fuels" and "Alternative Energies", shows a substitution effect in the salience of these topics. The Fossil Fuels topic was high until 2015, then died out. In contrast, the Alternative Energies topic increased around 2007, to remain relatively high until the recent years. This shift in concerns illustrates the skepticism associated to alternative energies and their implementation.

Figure 7 displays the evolution of the topics related to firms' actions and impact. Panel A reveals that the "Recycling" topic was predominant in the nineties, but has in the past twenty years almost disappeared from the greenwashing index. The Marketing topic has remained stable, consistent with the idea that firms greenwash by marketing their products as more sustainable than they are. The Labels topic has also remained stable, consistent with the many debates on labels in various sectors. Green labels have multiplied over the years, sparking criticism about their opacity and lack of regulation.[12]

Panel B of Figure 7 shows that the Emissions topic has grown in importance in the past 25 years. Reading through the main articles on emissions highlights the skepticism linked to emission measurements, and the lack of trust in these measurements. In contrast, the Health, Waste and Forests topics have decreased in salience. The most striking highlight of this graph, however, is the evolution of the Social topic. The topic importance remained close to constant until 2021, but then increased sharply in the past four years. Such increase, contrasted with the decrease in the traditional topics, indicates a shift in greenwashing concerns, from standard environmental topics to more social topics.

Figure 8, Panel A, shows the evolution of the parties involved in greenwashing, and in fighting greenwashing. As ESG considerations have become more and more integrated in firms' operations, the Corporation topic increased over time. In contrast, the role of environmentalists has become less salient. Panel B shows that the tools used to prevent greenwashing, i.e., lawsuits and disclosure rules, have remained essentially stable over time. This result indicates the shortcomings of the many disclosure frameworks proposed to address greenwashing, as detailed in Section 2.2.

[12]In the 1990s already, Proctor & Gamble misused the "biodegradable" label and composting symbols on diapers. As a result, attorneys general from 10 U.S. states forced Proctor & Gamble to stop using those claims. In 1991, S.C. Johnson used a self-made eco-label on their products, suggesting third-party environmental certification. In 2009, Kmart labelled their paper towels and tissues "biodegradable" and "eco-friendly" although they were standard tissue papers. The FTC consequently issued warnings as part of its crackdown on unsubstantiated environmental marketing claims. In 2007 and 2008, Dell claimed its computers would use 25% less energy by 2010, and both Dell and HP used green leaf imagery. Investigations found the claims were unverifiable at the time, and likely exaggerated. In the financial sector, self-labelled sustainable funds have raised many critics, leading the Names Rule (Rule 35d-1) to be amended by the SEC in 2023, to ensure more transparent labelling of green and ESG funds.

5 Greenwashing and Mutual Fund Flows

In this section, we focus on the latest years and attempt to understand the nature of the 2021 greenwashing peak and its link to the recent ESG backlash. As the peak in greenwashing was mainly driven by the financial sector, we build a sub-index of the greenwashing index, referred to as *greenwashing-in-finance* index, with greenwashing articles that are specifically about greenwashing in the financial industry. We show that greenwashing-in-finance is closely related to the confusion generated in the financial sector by the opacity of ratings, labels, and misleading claims on the risk-return profile of sustainable assets. We further create an index for the ESG backlash, and show that the backlash followed the peak in ESG confusion and greenwashing. Finally, we study the impact of greenwashing salience on mutual fund flows. Our analysis reveals interesting differences in how retail and institutional investors react to increases in greenwashing salience.

5.1 Greenwashing in finance and ESG confusion

In its January 17, 2025 article, the Financial Times underlines two possible sources of the recent ESG backlash: greenwashing, and the confusion around ESG concepts.[13] We first analyze the link between these two concepts.

In order to isolate the salience of greenwashing in the financial sector, we build a subindex of the greenwashing index which is solely composed of greenwashing articles about the financial industry. As the topic analysis gives us the importance of each topic in each article, we select all articles which have more than 5% of contents part of the Asset Management or the Bond topic. The resulting index is displayed in Figure 9, Panel A, and has a correlation of 52% with the greenwashing index. The high correlation is expected, as we have shown in Section 4.4 that the Financial topic increased in importance around 2021. Panel B compares greenwashing in and outside of finance. All greenwashing articles with less than 5% of finance-related contents are considered greenwashing-outside-finance. This graph confirms that greenwashing outside finance was dominant before 2021, but greenwashing in finance became dominant in 2021 and 2022. Between 2022 and 2025, the two types of greenwashing were comparable.

Similarly, we use our algorithm to build an index of confusion about what is sustainable, how to measure sustainability, and the risk-return profile of sustainable assets. The prompt used to build the training set used for this index is given in Appendix E.

[13] https://www.ft.com/content/14ee5968-79de-42bf-80a4-811531e80de7?

In total, we identify 186 articles about ESG confusion. The confusion index is displayed in Figure 10, together with the greenwashing-in-finance index, between 2018 and 2025. The two indices share 78 articles, and their time series exhibit similar patterns. Table 5 tests whether greenwashing tends to follow ESG confusion. As the two indices are persistent, we remove the autocorrelation by taking the innovations of an AR(1) process. The regression confirms that innovations in greenwashing tend to follow innovations in ESG confusion.

5.2 ESG Backlash

The media have unanimously reported an ESG backlash over the past few years. For example, in its December 4, 2023 article entitled "The real impact of the ESG backlash", the Financial Times gives examples of companies, such as Blackrock, which changed their communications over the years as a response to the anti-ESG sentiment. [14] In this section, we study whether increases in the greenwashing and ESG confusion indices are indeed followed by higher anti-ESG sentiment.

We build an anti-ESG sentiment index from the history of opinions of the Wall Street Journal. Since 1986, the Wall Street Journal has published nearly 80,000 opinion pieces. These pieces were filtered out in the construction of the index as by nature, they reflect sentiment rather than facts, and are targeting the general population rather than investors. We focus on the period 2018-2025, as greenwashing in finance was low before. Using GPT 4o-Turbo, we isolate the opinions that are about the adoption or effectiveness of ESG principles, and express a negative view on these principles. The prompt used is given in Appendix E.

The monthly index is displayed in Figure 11. It remains, on average, stable until 2021, but then jumps upward following the greenwashing peaks, until the end of the time series. This result does not prove any causality between the greenwashing peak and the ESG backlash, but is in line with a possible change in investors' behaviour following increases in greenwashing salience and ESG confusion.

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

5.3 Mutual Fund Flows

In this section we study whether the general anti-ESG sentiment that followed the 2021 greenwashing peaks translates into changes in how investors trade funds. Solomon et al. (2014) find that investors allocate more capital to funds whose holdings are more covered in the media. Fund flows, or at least some fund flows, are thus affected by media coverage. Our goal is to understand whether greenwashing salience affects flows, and the flow-performance sensitivity. This question is relevant as flows have been shown to have a strong effect on both funds' and their holdings' returns, see, e.g., Coval and Stafford (2007), and more recently in the context of green funds, 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 2025.[15] We focus on the evolution of their monthly flows. Flows in dollars are provided across share classes. The different share classes refer to the same portfolio, offered under different terms to different types of investors. Some of the share classes are flagged as targeting retail investors, the others are offered to institutional investors. On average, each fund has 2.7 share classes.

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

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

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

We filter out funds with size less than a million dollars. Both relative flows and share class sizes are skewed to the right. To avoid outliers, we winsorise them at the 99% level. We also winsorise relative flows at the 1% level. We further filter out funds which did not receive Morningstar sustainability and star ratings. Table 6 summarizes the sample. The final dataset contains 8,818 funds and 21,747 share classes.

Each share class in our sample has a Morningstar globe rating, updated on a monthly basis. Globe ratings assess the sustainability of a fund. An aggregated score is calculated for each fund by taking a weighted average of the ESG ratings of their holdings, obtained

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

from Sustainalytics. Funds are then ranked based on these scores. They are given five globes (four, three, two and one respectively) if they are in the top 10% of funds within their category (between 10% and 32.5%, between 32.5% and 67.5%, between 67.5% and 90%, and above 90%).

As in Solomon et al. (2014), it may be that some funds are more affected by greenwashing than others. In particular, we distinguish between share classes of funds that are sustainable, and other share classes. We use two proxies for fund sustainability: the globe rating given by Morningstar, and the name of the fund. We identify funds with an ESG name as funds whose name contains one or several words from a list of words commonly used in fund names and linked to sustainability. This list includes words such as "environment", "green", "sustainable", "ESG", "clean", "social" and "renewable". We find 392 funds with an ESG-related word. There is no one-to-one relation between having a 5-globe rating and an ESG name: nearly 40% of the funds advertised as sustainable have a time-average globe rating below 4, and about 10% of the funds that are not advertised as green have a time-average rating above 4.[16]

We run panel regressions of share class flows, as defined in equation (1), on lagged returns, fund sustainability and cross-effects, to study the flow-performance sensitivity in times of low and high greenwashing salience. We add the firm sustainability index as control.

The general regression is the following:

$$\begin{aligned}\text{Fund flows}_{i,t:t+1m} = {} & \overbrace{a_0 \text{ Performance}_{t-12m:t} + a_1 D_{\text{Fund sustainability},t-1m:t} + \text{Cross effects}_{t-1m:t}}^{\text{Determinants}_t} \\ & + GW_{t-1m:t} \times \text{Determinants}_t + \text{Firm sustainability}_{t-1m:t} \times \text{Determinants}_t \\ & + \text{Controls}_{i,t-1m:t} + FE. \end{aligned} \tag{2}$$

The coefficient a_0 measures the flow-performance sensitivity. Following Lou (2012), we use past returns over the previous year as main proxy for performance. Tests are also run with lagged monthly returns as an alternative measure for past performance. The coefficient a_1 captures the additional flows that are allocated to sustainable funds. Cross-effects capture the potential additional flow-performance sensitivity of sustainable funds. If the flows into sustainable funds are dominated by values investors, as defined by Starks (2023), these flows may be less sensitive to a lower performance of the funds that the flows of traditional funds.

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

The second line of the regression captures the impact of greenwashing salience. We cross all the terms in the first line with the lagged value of the greenwashing index. We use the following controls: flows over the past year, lagged size and age of the share class, and expense ratio.

Solomon et al. (2014) show that the reaction of fund flows to news depends on investors' sophistication. They find that investors that are less sophisticated react more to news than other investors. Therefore we run the regressions for all share classes, irrespective of which type of investor they target, and then separately for share classes targeting retail and institutional investors. We use category-year fixed effects, where the category is the one assigned by Morningstar to funds. We double-cluster standard errors by fund and time. Table 7 presents the results of the regression for equity funds. Columns 1 and 4 contain the results for all share classes, columns 2 and 5 correspond to retail investors and columns 3 and 6 to institutional investors. Columns 4 to 6 include the firm sustainability index as a control as well as cross-effects with it, whereas these effects are omitted in columns 1 to 3. The results reveal several interesting facts.

First, irrespective of the greenwashing index, flows increase with past performance as measured by the past 1-year returns. The relation is in line with the results of Lou (2012), and statistically significant in all regressions. Furthermore, globe ratings matter, but only for institutional investors: they invest more in 5-globe rated funds, and less in 1-globe rated funds. Retail investors invest more in both types of funds, but only when they have high past returns. This first finding highlights an important difference between institutional investors and retail investors: while the former assign weight on the globe rating of funds, the latter focus more on performance.

In times of high greenwashing salience, the flow-performance sensitivity tends to decrease, although the effect is not statistically significant. Furthermore, our previous result, in times of low greenwashing, is reinforced when greenwashing salience is high: institutional investors allocate more capital to 5-globe funds and less capital to 1-globe funds. The effect is an order of magnitude larger than in times of low greenwashing salience, and statistically significant. These results thus indicate that when greenwashing is salient, institutional investors prefer allocating their capital to funds externally certified as ESG-focused. Such behaviour is in line with investors assuming sustainable fund managers perform due diligence to filter out greenwashing firms. Alternatively, it could also be explained by social signalling and reputational concerns. When media scrutiny is high, institutional investors may seek to protect their reputation by visibly investing in certified green funds, avoiding backlash for supporting firms criticised for greenwashing.

These results do not hold for retail investors, who do not invest significantly more in 5-globe rated funds. Instead, they invest more capital into 1-globe funds with high past returns. For these investors, the flow-performance sensitivity of low-rated funds thus increases in times of high greenwashing salience.

Our findings hold whether we control for the salience of firm sustainability or not. Firm sustainability salience has no impact on the flow-performance sensitivity of funds. It has a weak impact on the trading of institutional investors. In times of high firm sustainability salience, they invest less capital in 5-globe funds, and more capital in 1-globe funds, in line with lower social signalling needs.

Table 8 shows that when using the fund's name as a proxy for its sustainability, results disappear. Institutional investors do not assign any weight to the name of funds, whether greenwashing salience is low or high. These results are in line with ESG ratings being either taken more seriously than fund names, or serving as better signalling tools. In contrast, this table shows weak evidence that retail investors care about fund names, and tend to invest more in funds labelled sustainable in times of high greenwashing salience.

In Appendix F, we show in Table A1 that our results still hold when using all funds (instead of equity funds only) in the regression. Table A2 shows that results on the flow-performance sensitivity no longer hold with lagged returns instead of past 1-year returns. But the shift of institutional investors' capital to 5-globe funds in times of high greenwashing salience remains. Furthermore, table A3 shows that using 5-star ratings as a proxy for funds' past performance preserves the results as well. When using the greenwashing-in-finance index instead of the greenwashing index in the regressions, coefficients have the same sign and magnitude but we lose statistical significance, see Table A4. This result indicates that it is overall greenwashing salience, and not solely greenwashing in the financial sector, which affects investors' behaviour.

All our results still hold when removing the greenwashing articles from the firm sustainability index, thereby decreasing the correlation between greenwashing and firm sustainability indices to 20%. See Table A5.

References

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.

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.

Baker, S. R., N. Bloom, S. J. Davis, and M. Sammon (2025). What triggers stock market jumps? University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2021-42.

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

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

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

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

Cao, S., Y. Chen, H. Xiao, and A. Zhang (2025). The Economics of Greenwashing Funds. Working paper.

Carlin, B. I., F. A. Longstaff, and K. Matoba (2014). Disagreement and asset prices. *Journal of Financial Economics 114*(2), 226–238.

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.

Coval, J. D. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics 86*(2), 479–512.

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).

Dougal, C., J. Engelberg, D. Garcia, and C. A. Parsons (2012). Journalists and the stock market. *Review of Financial Studies 25*(3), 639–679.

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

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

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

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

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

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

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

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

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

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

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

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.

Lou, D. (2012). A flow-based explanation for return predictability. *The Review of Financial Studies 25*(12), 3457–3489.

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.

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.

Starks, L. T. (2023). Presidential address: Sustainable finance and esg issues — value versus values. *Journal of Finance 78*(4), 1837–1872.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance 62*(3), 1139–1168.

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

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 2025. In total, the database contains 899,422 articles.

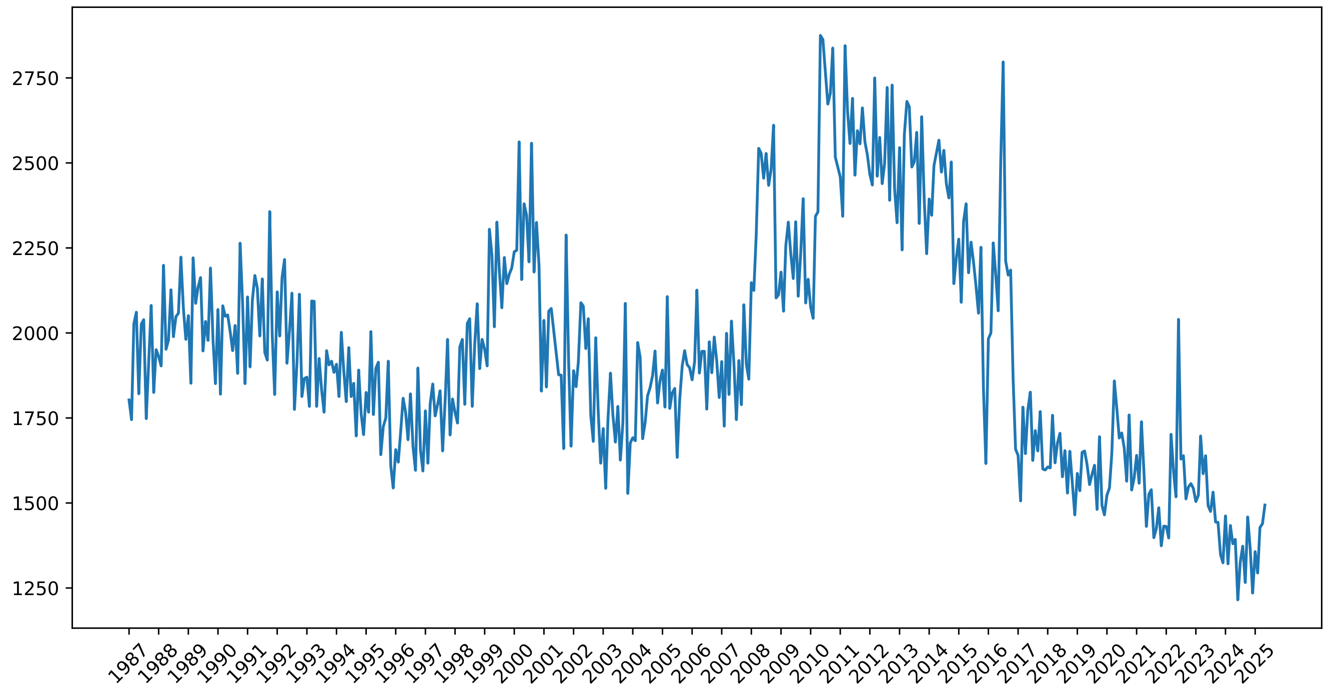

Figure 2: Firm Sustainability Index

This figure displays the monthly number of sustainability articles as a percentage of the total number of articles. The classification of sustainability articles follows the procedure outlined in Section 3.2. Annotations mark sustainability-related events.

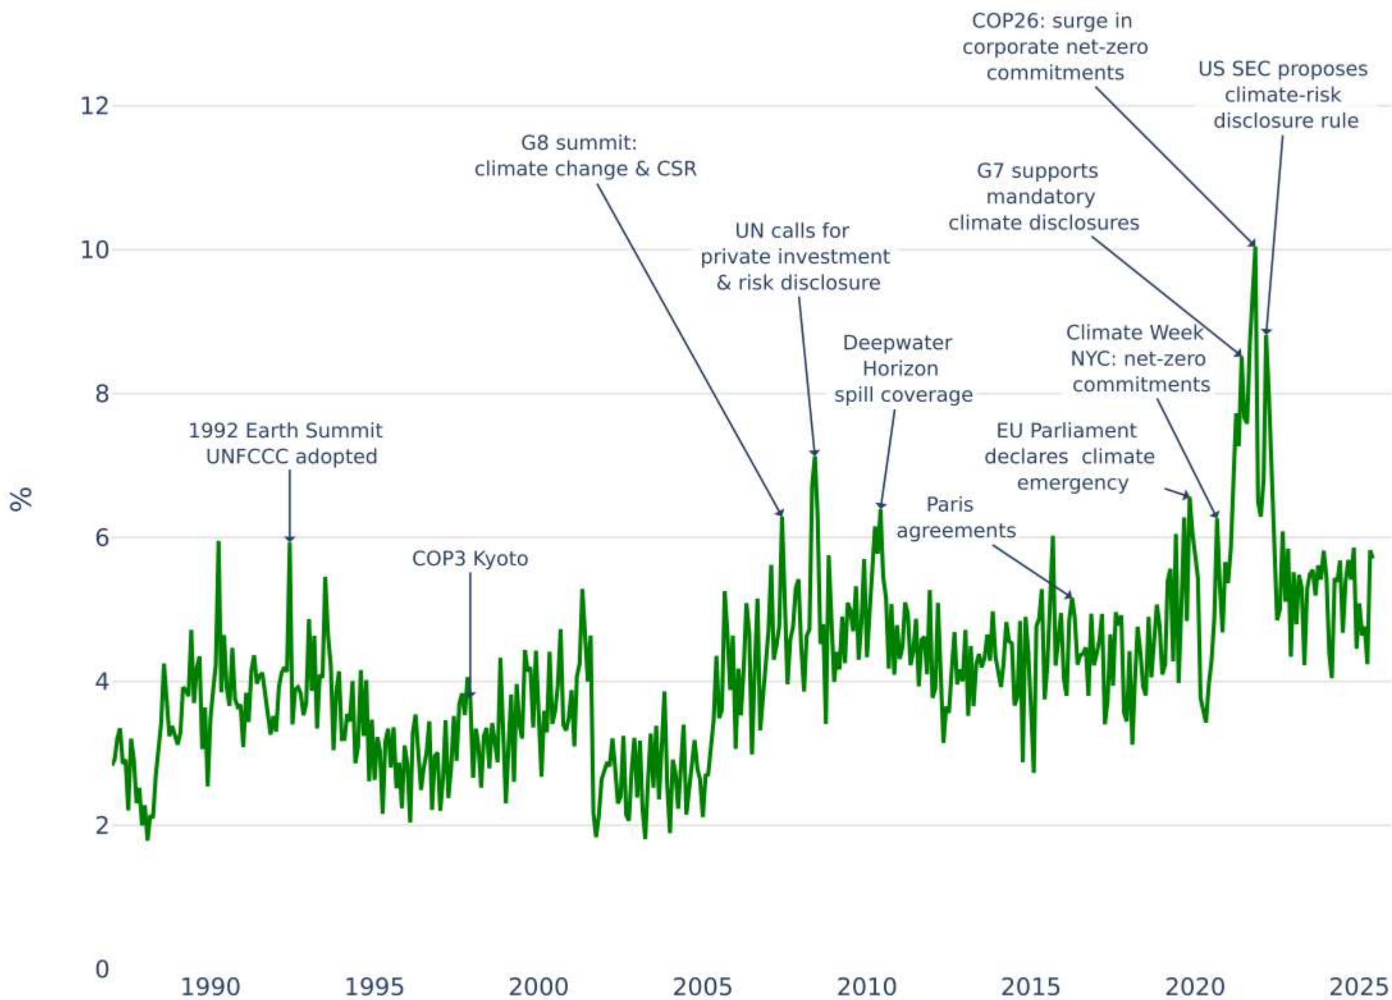

Figure 3: Greenwashing Index

This figure displays the monthly greenwashing index, built following the procedure outlined in Section 3.2. This index is obtained by counting the number of greenwashing articles each month, as a fraction of the total number of articles. Annotations mark accusations or charges of greenwashing that were made against given companies.

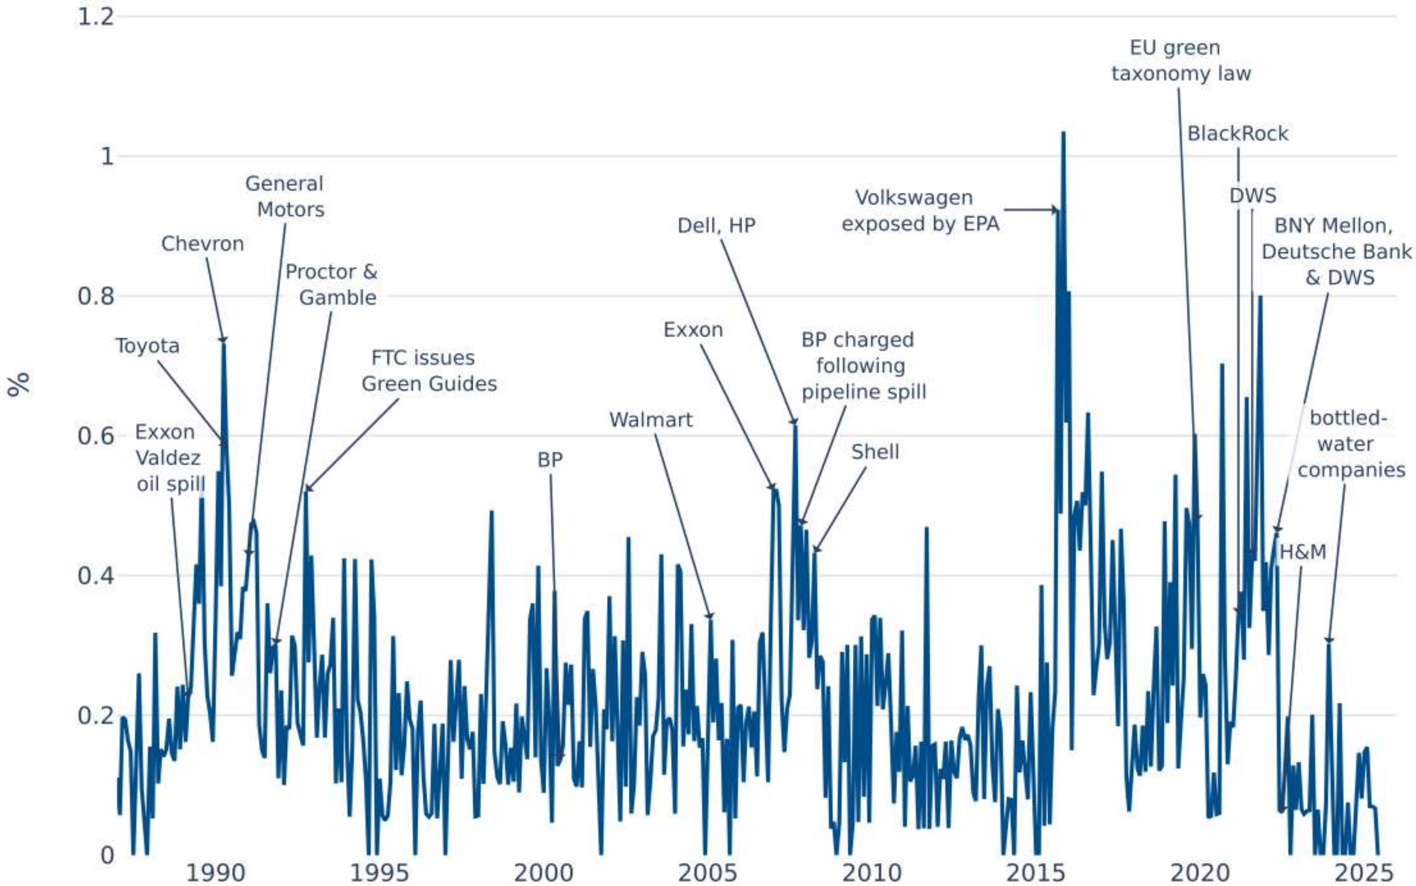

Figure 4: Companies Mentioned in Greenwashing Articles

This figure is a word cloud of the companies mentioned in greenwashing articles. Panel A corresponds to articles classified as greenwashing articles between 1986 and 2025. Panel B focuses on articles classified as greenwashing articles between 2018 and 2025. The method is restricted to companies traded on the NYSE, the NASDAQ or the AMEX. The size of words is proportional to the number of times companies are mentioned.

Panel A: 1986-2025

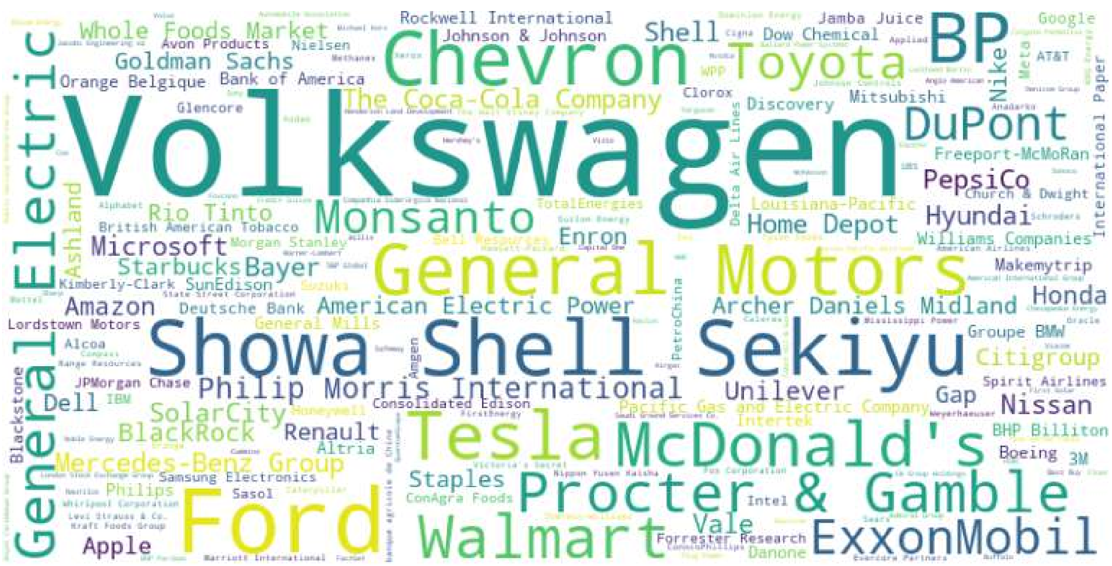

Panel B: 2018-2025

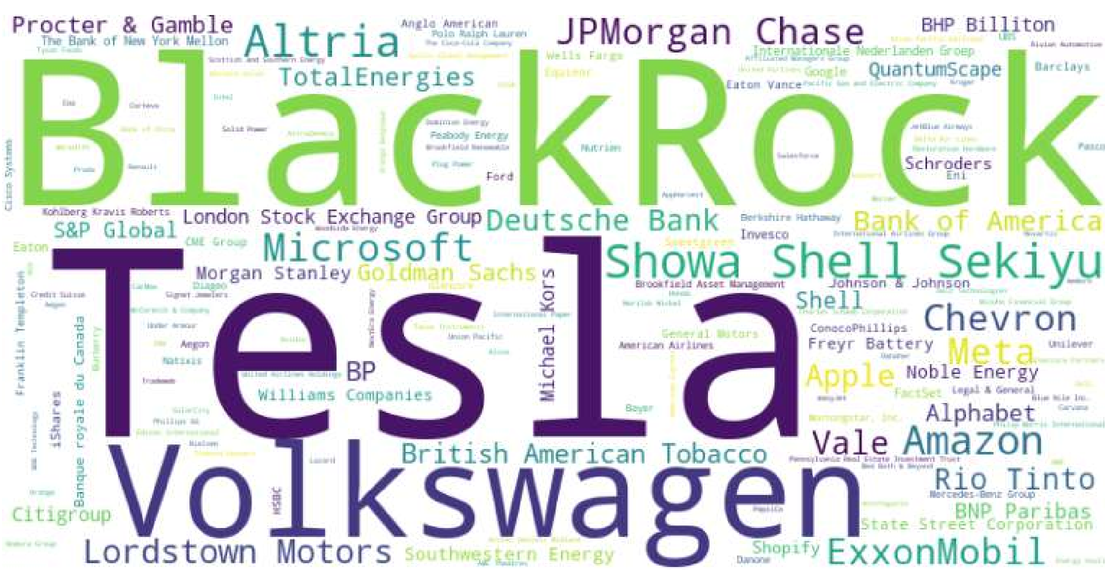

Figure 5: Topic Analysis: Industries

This figure shows the yearly topic proportion estimated by the keyATM algorithm described in Section 4.4, for industry topics, from These industries include the Financial, Retail, Automobile, Energy, Construction and Tobacco industries.

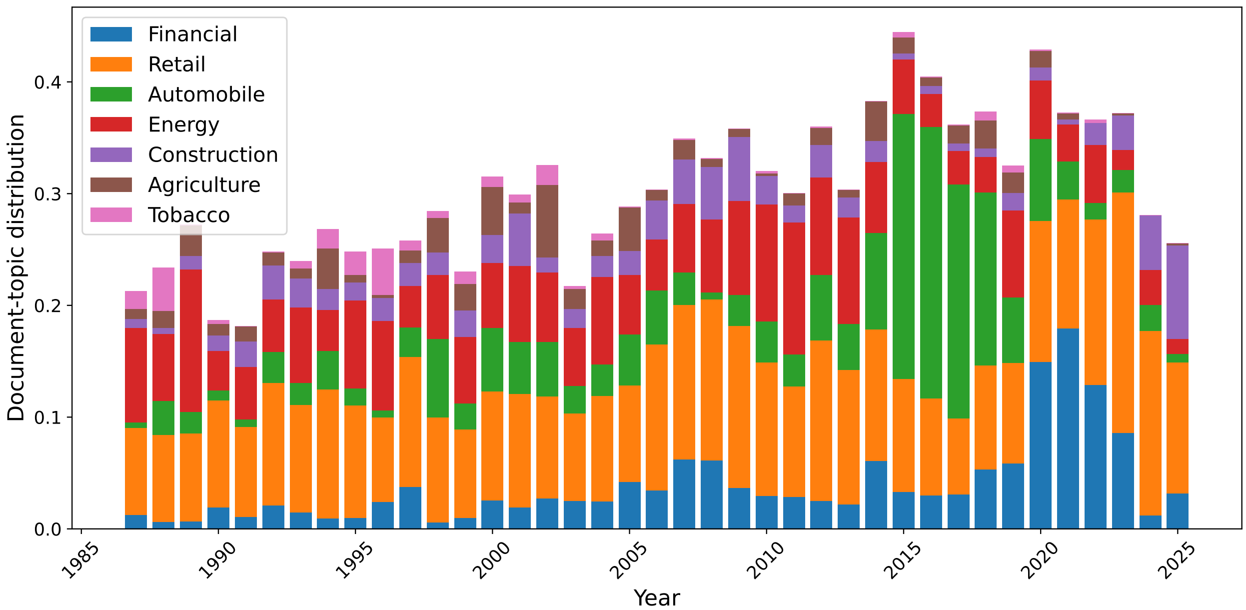

Figure 6: Zoom on Financial and Energy Sectors

This figure shows the yearly topic proportion estimated by the keyATM algorithm described in Section 4.4, for the two topics related to the Financial sector: Asset Management 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

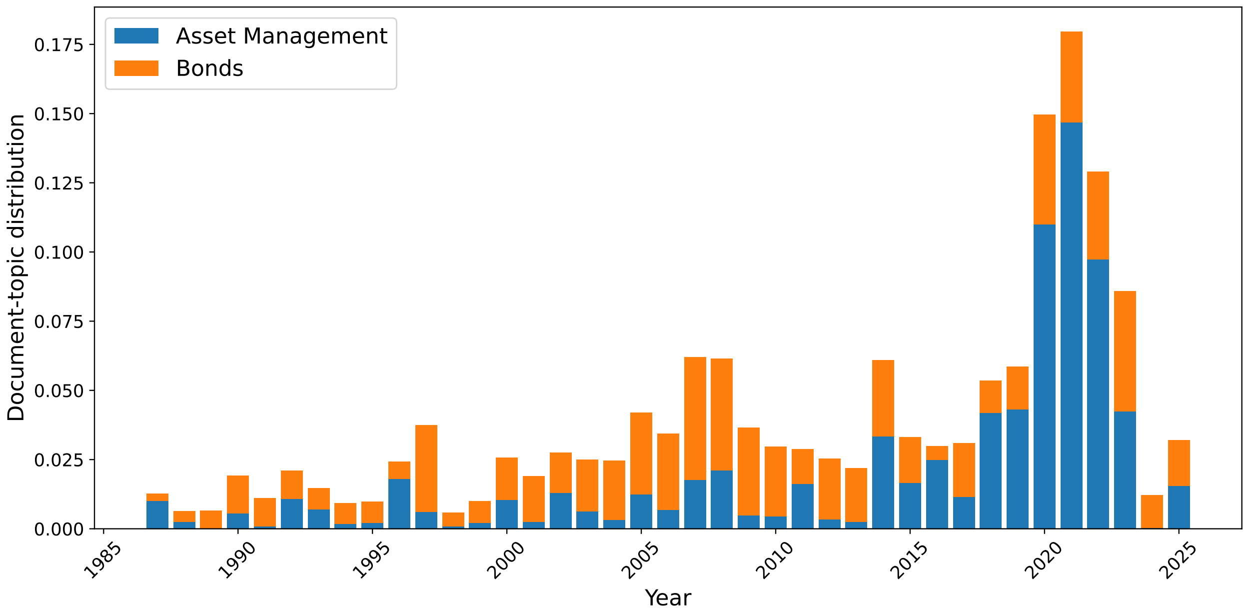

Panel B: Energy sector

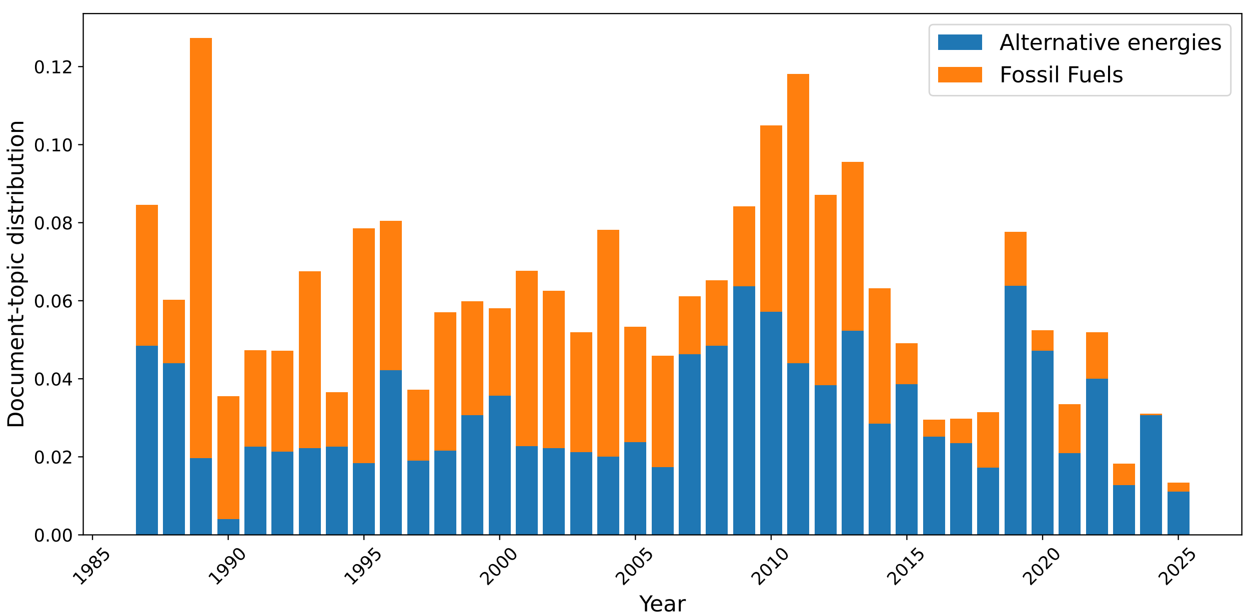

Figure 7: Topic Analysis: Firms' Actions and Impact

This figure shows greenwashing topics 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 firms' environmental and social actions: Labels, Marketing and Recycling. The five topics displayed in Panel B are related to firms' impact: Social, Emissions, Health, Waste and Forests.

Panel A: Firms' actions

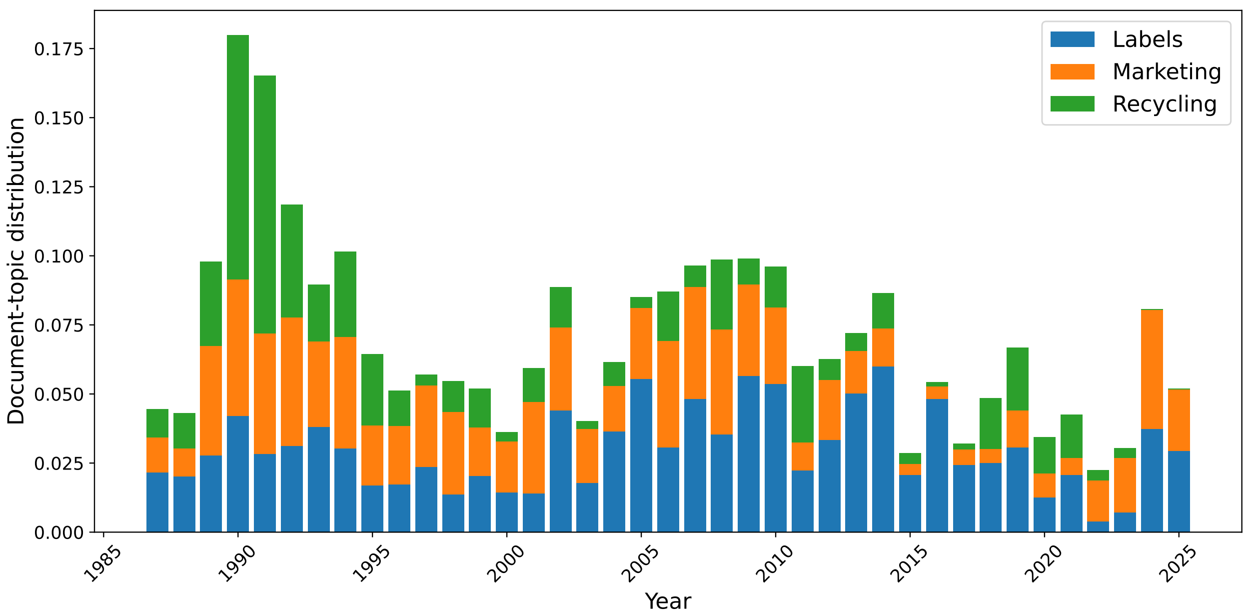

Panel B: Firms' impact

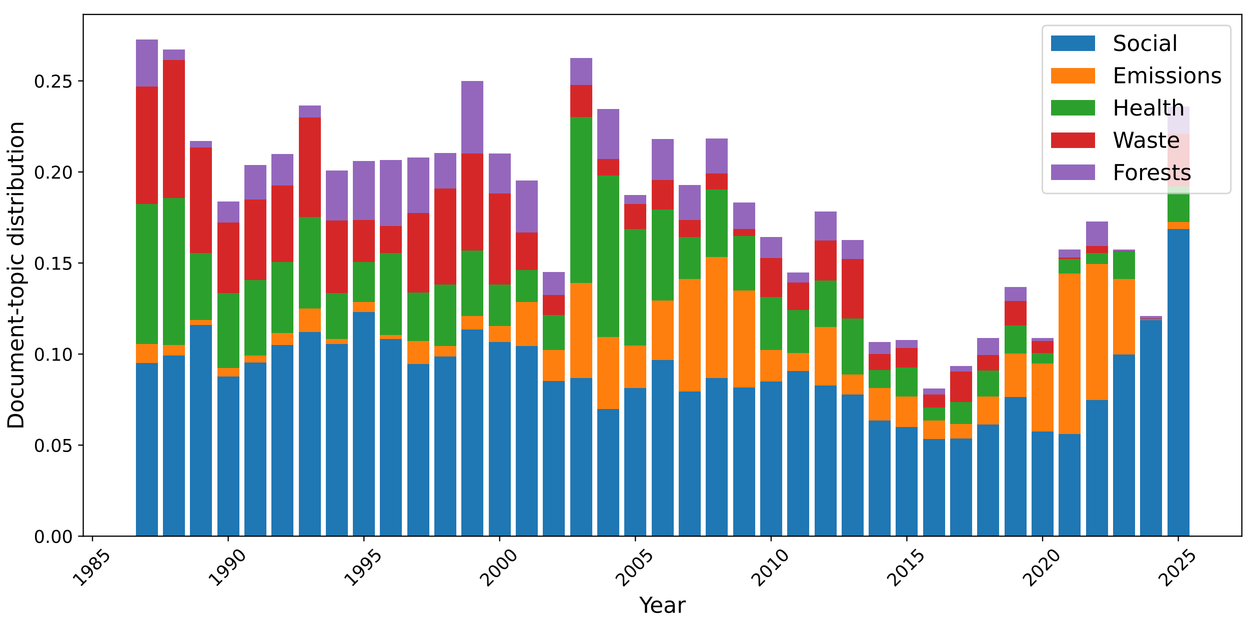

Figure 8: Topic Analysis: Greenwashing Actors and Tools

This figure shows the greenwashing topics related to parties involved, and tools used, over time. The three topics displayed in Panel A are parties involved in greenwashing: Corporations, Environmentalists and Politics. The two topics displayed in Panel B are related to tools that can be used to disincentivize greenwashing: Disclosure and Lawsuits.

Panel A: Parties involved

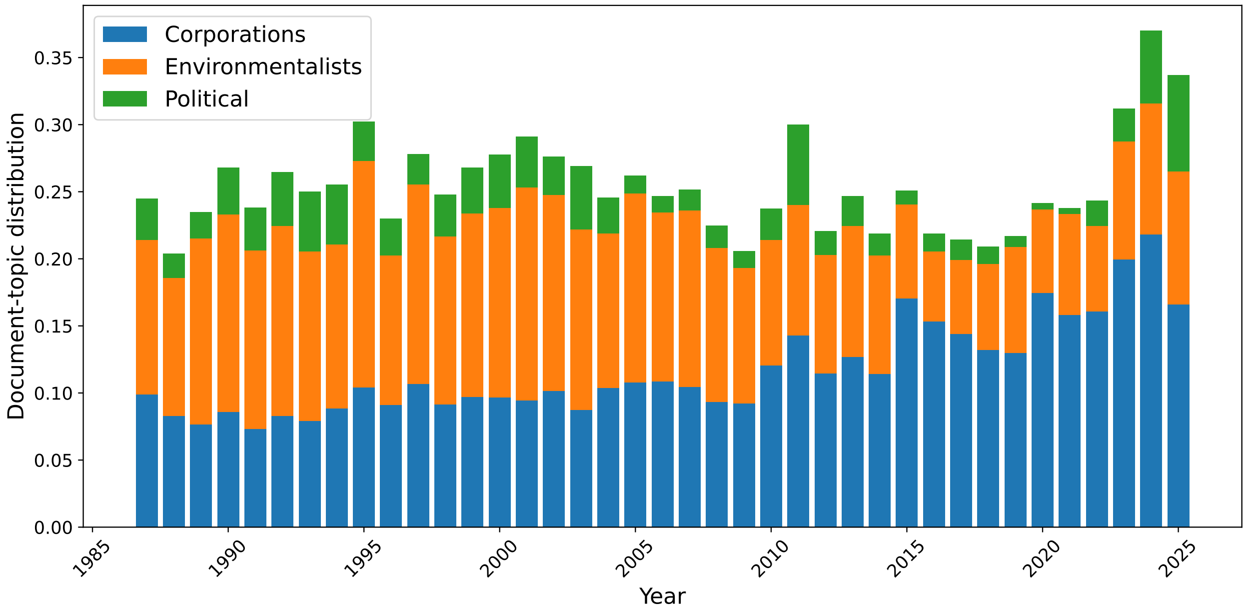

Panel B: Tools to prevent greenwashing

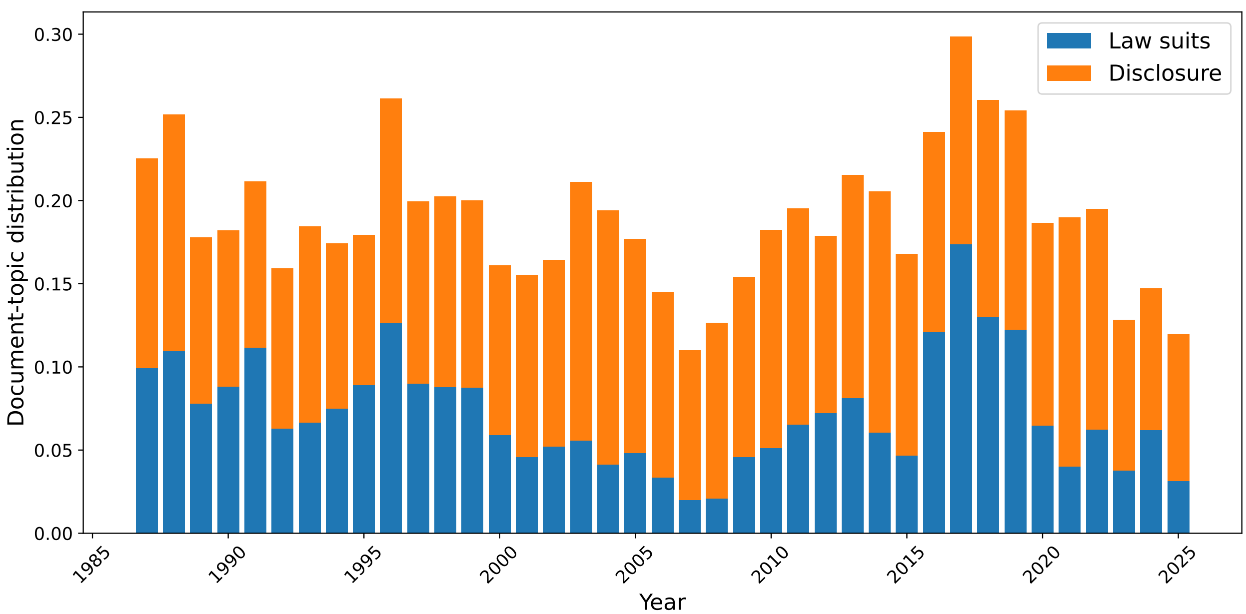

Figure 9: Greenwashing in Finance

This figure displays the greenwashing index and the greenwashing-in-finance subindex, from 2018 to 2025. The construction of the greenwashing-in-finance subindex is detailed in Section 5.

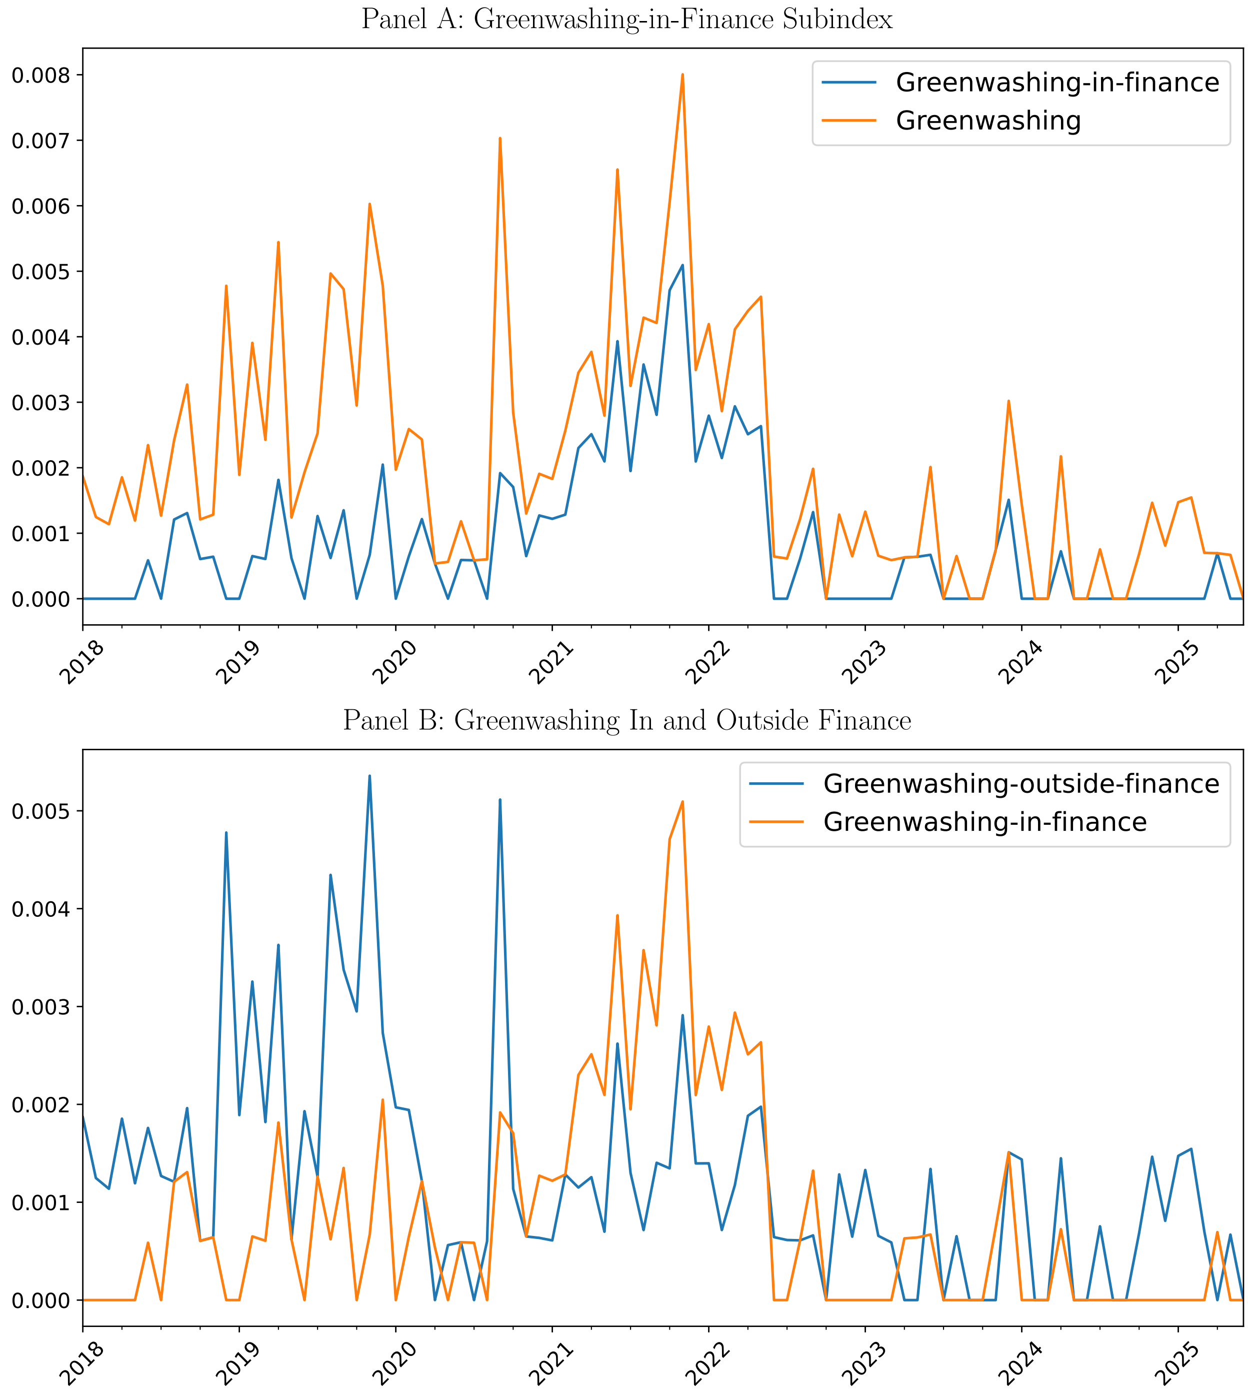

Figure 10: Greenwashing in Finance and ESG Confusion

This figure displays the monthly greenwashing-in-finance and ESG Confusion indices, built following the procedure outlined in Section 5.1.

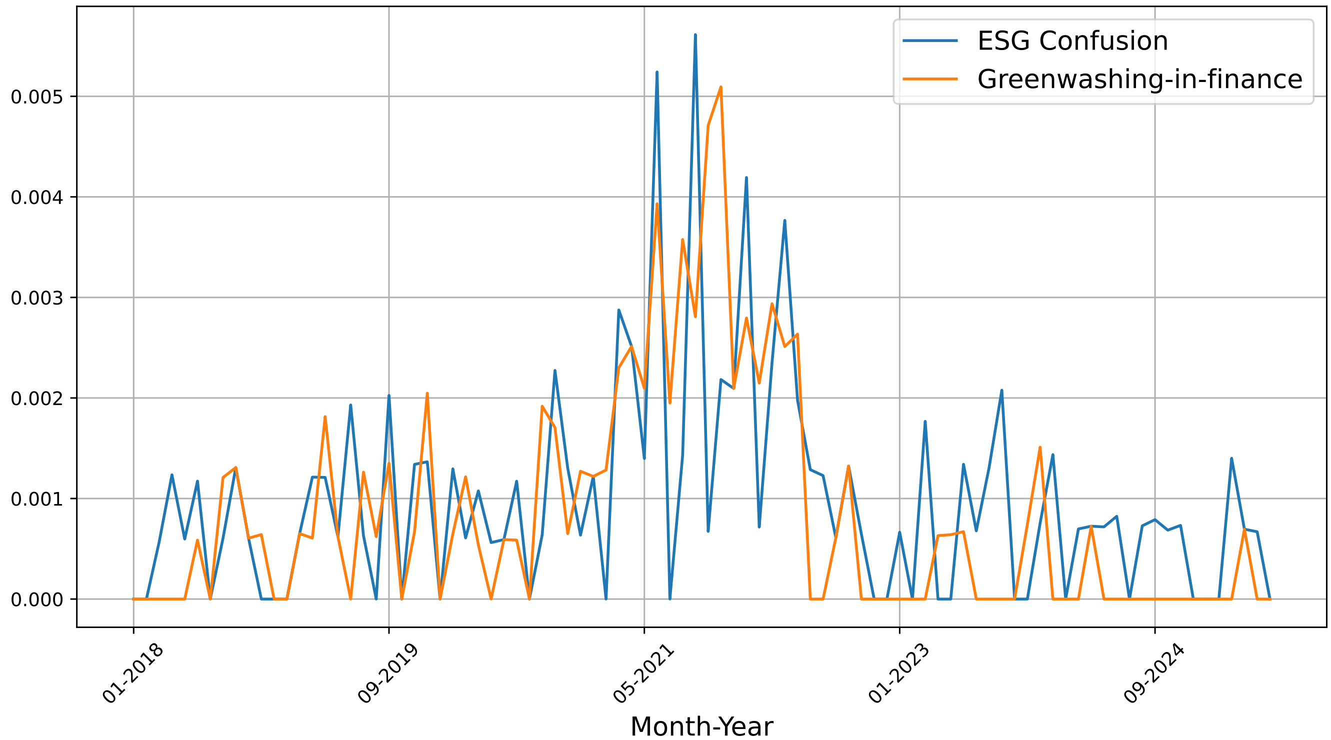

Figure 11: ESG Backlash Index

This figure displays the monthly greenwashing-in-finance index as well as the monthly ESG Backlash index, built following the procedure outlined in Section 5.2.

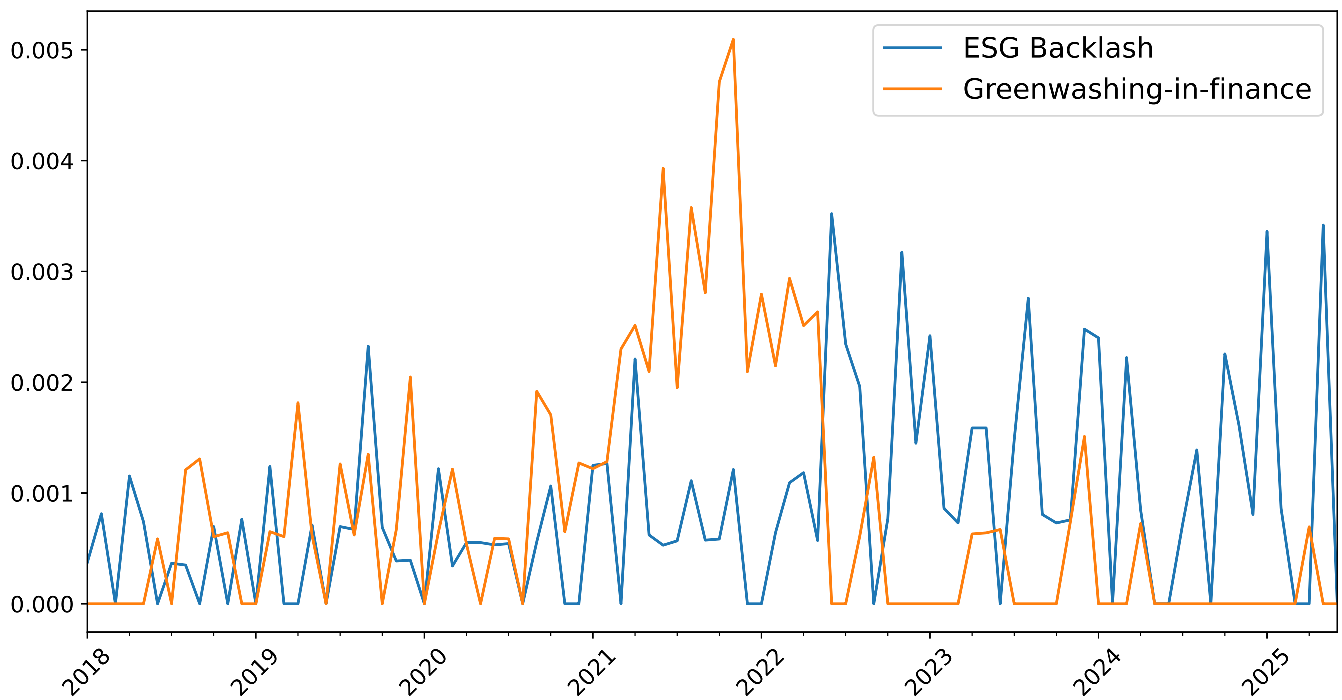

Table 1: Description of Sets Used in the Classification

This table summarizes the training, validation and testing sets used in the two steps of our classification algorithm. The column named '1. Firms' sustainability' contains the volume of articles used in each set in the first step (identification of firms' sustainability risk-related articles), and the column named '2. Greenwashing' contains the numbers corresponding to the second step (identification of greenwashing articles). The classification algorithm is described in Section 3.2.

Step		1. Firms' sustainability		2. Greenwashing	
Training set		1	0	1	0
	# articles	1132	1132	690	690
	Min year	1987	1987	1987	1987
	Max year	2025	2025	2025	2025
Validation set		1	0	1	0
	# articles	300	2100	200	2900
	Min year	1987	1987	1987	1987
	Max year	2025	2025	2023	2025
Testing set		1	0	1	0
	# articles	300	2100	200	2900
	Min year	1987	1987	1987	1987
	Max year	2025	2025	2025	2025

Table 2: Performance Metrics of Classification Algorithms

This table summarizes the performance of the three types of algorithms tested for the classification of firm sustainability-related articles (step 1) and greenwashing articles (step 2). The steps of these algorithms are described in Section 3.2.

Step		Accuracy	Precision	Recall	F-score	AUC
1. Firm sustainability	**This paper, BoW**	**95.7**	**82.7**	**82.7**	**82.7**	**90.1**
	This paper, LLM (OpenAI)	95.1	83.7	75.3	79.3	86.0
	Engle et al. (2020)	88.5	55.5	39.3	46.0	67.4
	Sautner et al. (2023a)	87.3	49.1	52.7	50.8	72.5
2. Greenwashing	**This paper, LLM (OpenAI)**	**95.6**	**78.2**	**52.0**	**62.5**	**79.5**
	This paper, BoW	94.4	56.3	60.0	58.1	78.4
	Engle et al. (2020)	48.4	7.9	66.0	14.1	56.5
	Sautner et al. (2023a)	93.6	75.0	1.6	3.1	50.8

Table 3: Correlation between Firm Sustainability and Climate Risk Indices

This table displays the correlation matrix between our monthly firm sustainability index, and alternative indices of climate risk. 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	**0.70**	1					
EGKLS	**0.11**	-0.17	1				
FMS-policy	**0.57**	0.36	0.29	1			
FMS-summits	**0.32**	0.00	0.53	0.55	1		
FMS-glowarm	**0.44**	0.41	0.37	0.57	0.70	1	
FMS-natdisast	**0.37**	0.28	0.36	0.47	0.61	0.72	1

Table 4: Topic Analysis

This table displays the main 22 topics identified by the keyATM method, as well as the words characterizing them. The keyATM method is described in Section 4.4. The topics are ordered by topic proportion.

Topic	%	Most characteristic words
Disclosure	12%	report, standard, service, data, system, spokesman, information, based, agency, article
Corporation	11%	executive, business, chief, month, share, board, investor, corporate, deal, end
Activism	11%	group, environmental, industry, public, official, national, director, effort, environmentalist, environment
Retail	10%	market, product, price, consumer, sale, cost, customer, service, technology, sell
Social	9%	people, work, problem, long, run, office, safety, old, worker, ago
Lawsuit	7%	case, court, attorney, law, federal, suit, investigation, settlement, judge, state
Automobile	6%	car, vehicle, volkswagen, emission, auto, diesel, fuel, engine, test, ford
Label	3%	product, consumer, label, food, naturel, brand, ingredient, fda, organic, claim
Alternative energies	3%	power, energy, plant, solar, utility, project, coal, electricity, dam, nuclear
Health	3%	water, health, chemical, epa, study, level, test, toxic, lead, cancer
Fossil fuels	3%	oil, exxon, gasoline, spill, shell, gas, ethanol, pipeline, alaska, fuel
Waste	2%	waste, plant, state, official, environmental, site, epa, department, cleanup, water
Marketing	2%	marketing, campaign, advertising, york, people, good, show, child, message, claim
Politics	2%	state, president, air, administration, tax, federal, house, bill, white, california
Emissions	2%	carbon, emission, credit, climate, gas, fuel, global, energy, dioxide, change
Construction	2%	building, light, air, saving, water, house, bulb, device, system, old
Asset management	2%	fund, esg, investor, investment, stock, social, sustainable, asset, investing, firm
Recycling	2%	plastic, paper, recycling, waste, recycled, bag, packaging, material, trash, diaper
Bonds	2%	green, bank, mart, bond, wal, energy, program loan, initiative, certification
Agriculture	2%	food, organic, farmer, crop, farm, plant, corn, genetically, pesticide, agriculture
Forest	1%	tree, land, wood, pacific, palm, acre, area, logging, lumber
Tobacco	1%	cigarette, tobacco, smoke, philip, morris, smoker, smoking, reynolds, filter, product

Table 5: ESG Confusion and Greenwashing

This table reports the results of regressions of unexpected shocks in greenwashing indices on the lagged values of unexpected shocks in ESG confusion. Unexpected shocks are computed as innovations of an AR(1) model. GS_t denotes shocks to the greenwashing index, $GSFin_t$ shocks to the greenwashing-in-finance index, and $ConfS_t$ shocks to the ESG confusion index. The ESG confusion index is built following the procedure outline in Section 5.1. Regressions were performed on time series starting in 2000 in columns (1) and (3) and in 2010 in columns (2) and (4). 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.

	GS_t Since 2000	GS_t Since 2010	$GSFin_t$ Since 2000	$GSFin_t$ Since 2010
$ConfS_{t-1}$	-0.09	-0.13	0.03	-0.13
	(0.15)	(0.16)	(0.08)	(0.16)
$ConfS_{t-2}$	0.34**	0.36***	0.28**	0.36***
	(0.15)	(0.12)	(0.11)	(0.13)
$ConfS_{t-3}$	0.10	0.08	0.07	0.08
	(0.10)	(0.15)	(0.10)	(0.15)
Observations	302	182	302	182
R^2	0.02	0.02	0.08	0.02

Table 6: Summary Statistics

This table presents summary statistics for the funds obtained from Morningstar, from August 2018 to June 2025. 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 7: Flow-Performance Sensitivity and Greenwashing

This table reports the results of panel regressions of equity fund flows (relative to size) on determinants. Regressions are run at the share class level. Columns (1) and (4) correspond to all share classes, columns (2) and (5) to share classes targeting retail investors and columns (3) and (6) for share classes targeting institutional investors. The determinants include fund performance, measured as the fund's past 1-year returns, fund sustainability measured using dummies for 5-globe and 1-globe Morningstar ratings, the greenwashing index and in columns (4) to (6) the firm sustainability indices, as well as the interaction of these variables. All explanatory variables except funds' past performance are lagged by one month. Controls include lagged fund net assets (log), lagged share class size (log), the adjusted net expense ratio and the class age. Category-year fixed effects are included, using the categories of fund defined by Morningstar. Standard errors are double-clustered by fund and time. We present the t-statistics in the parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
1-year return	0.033***	0.029***	0.040***	0.039***	0.037***	0.044***
	(0.004)	(0.003)	(0.005)	(0.010)	(0.009)	(0.012)
$D_{5\text{ globes}}$	0.001	0.001	0.001	0.001	-0.001	0.006*
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{1\text{ globe}}$	0.001	0.001**	0.000	-0.002	0.000	-0.006**
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{5\text{ globes}} \times$ 1-year return	0.009***	0.011***	0.007	0.022**	0.029***	0.004
	(0.003)	(0.003)	(0.004)	(0.010)	(0.011)	(0.015)
$D_{1\text{ globe}} \times$ 1-year return	0.007***	0.008***	0.005	0.031***	0.036***	0.024*
	(0.002)	(0.003)	(0.004)	(0.009)	(0.011)	(0.014)
$GW_{m-1} \times$ 1-year return	-1.708	-0.996	-3.176**	-0.898	-0.113	-2.522
	(1.155)	(1.026)	(1.466)	(1.681)	(1.503)	(2.139)
$GW_{m-1} \times D_{5\text{ globes}}$	0.402*	0.328	0.509*	0.419	0.174	0.863**
	(0.209)	(0.228)	(0.298)	(0.278)	(0.287)	(0.418)
$GW_{m-1} \times D_{1\text{ globe}}$	-0.284	-0.289*	-0.356	-0.476**	-0.356	-0.844**
	(0.184)	(0.170)	(0.306)	(0.226)	(0.228)	(0.374)
$GW_{m-1} \times D_{5\text{ globes}} \times$ 1-year return	1.065	0.242	2.740**	2.631*	2.171	2.972
	(1.022)	(1.178)	(1.377)	(1.479)	(1.562)	(2.261)
$GW_{m-1} \times D_{1\text{ globes}} \times$ 1-year return	-1.077*	-0.891	-1.214	1.743*	2.514**	0.613
	(0.635)	(0.730)	(1.108)	(1.056)	(1.172)	(1.682)
$Firm\ sust_{m-1} \times$ 1-year return				-0.145	-0.165	-0.098
				(0.194)	(0.180)	(0.249)
$Firm\ sust_{m-1} \times D_{5\text{ globes}}$				-0.011	0.038	-0.111*
				(0.038)	(0.039)	(0.064)
$Firm\ sust_{m-1} \times D_{1\text{ globe}}$				0.061*	0.029	0.137**
				(0.031)	(0.035)	(0.058)
$Firm\ sust_{m-1} \times D_{5\text{ globes}} \times$ 1-year return				-0.273	-0.382**	0.056
				(0.183)	(0.195)	(0.295)
$Firm\ sust_{m-1} \times D_{1\text{ globe}} \times$ 1-year return				-0.524***	-0.603***	-0.405
				(0.170)	(0.214)	(0.255)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	602790	408584	194206	602790	408584	194206
N. of groups	8933	5991	2942	8933	5991	2942
R^2	0.041	0.045	0.029	0.042	0.045	0.030

Table 8: Flow-Performance Sensitivity and Greenwashing for Funds Advertised Sustainable

This table reports the results of panel regressions of equity fund flows (relative to size) on determinants. Regressions are run at the share class level. Columns (1) and (4) correspond to all share classes, columns (2) and (5) to share classes targeting retail investors and columns (3) and (6) for share classes targeting institutional investors. The determinants include fund performance, measured as the fund's past 1-year returns, fund sustainability measured using dummies for funds self-labelled sustainable, the greenwashing index and in columns (4) to (6) the firm sustainability indices, as well as the interaction of these variables. All explanatory variables except funds' past performance are lagged by one month. Controls include lagged fund net assets (log), lagged share class size (log), the adjusted net expense ratio and the class age. Category-year fixed effects are included, using the categories of fund defined by Morningstar. Standard errors are double-clustered by fund and time. We present the t-statistics in the parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
1-year return	0.034***	0.031***	0.041***	0.045***	0.043***	0.049***
	(0.004)	(0.004)	(0.005)	(0.010)	(0.009)	(0.012)
Sustainable fund name	0.001	0.000	0.002	-0.001	-0.001	-0.003
	(0.001)	(0.001)	(0.001)	(0.002)	(0.003)	(0.004)
Sust. fund name × Lagged return	-0.002	-0.000	-0.007	-0.010	-0.012	-0.002
	(0.004)	(0.005)	(0.007)	(0.014)	(0.014)	(0.026)
$GW_{m-1}\times$ 1-year return	-1.828	-1.097	-3.312**	-0.509	0.307	-2.213
	(1.135)	(1.025)	(1.407)	(1.630)	(1.472)	(2.042)
$GW_{m-1}\times$ Sust. fund name	0.889***	0.825***	0.954**	0.726**	0.742*	0.589
	(0.275)	(0.307)	(0.468)	(0.360)	(0.420)	(0.633)
$GW_{m-1}\times$ Sust. fund name × 1-year return	1.714	0.565	3.733**	0.563	-1.047	3.831
	(1.109)	(1.302)	(1.754)	(2.050)	(1.911)	(3.287)
$Firm\ sust_{m-1}\times$ 1-year return				-0.238	-0.262	-0.180
				(0.194)	(0.181)	(0.247)
$Firm\ sust_{m-1}\times$ Sust. fund name				0.045	0.023	0.098
				(0.049)	(0.054)	(0.091)
$Firm\ sust_{m-1}\times$ Sust. fund name × 1-year return				0.158	0.256	-0.097
				(0.276)	(0.264)	(0.505)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	602790	408584	194206	602790	408584	194206
N. of groups	8933	5991	2942	8933	5991	2942
R^2	0.040	0.044	0.028	0.041	0.044	0.029

ONLINE APPENDIX

A Manual labelling

A.1 Selection of the articles to review

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.

A.2 Instructions given to the assistants

The assistants were asked to read each article they were given, and to determine whether the article was a sustainability article or not. Sustainability articles are defined as "articles alluding to the integration of environmental and social considerations into corporate behavior and financial decision-making".

For each article classified as a sustainability article, assistants were asked whether the article was a greenwashing article or not. Greenwashing articles are defined as "containing criticisms of companies' claims or marketing related to environmental and social issues for being purposely exaggerated, misleading, false or unsubstantiated." We specified that an article could be a greenwashing article even if it did not contain the word "greenwashing".

A.3 Processing of outputs

A total of 1479 articles were manually labelled. 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.

B Extension of manually labelled set with GPT

We use GPT 4o-mini to extend the set of sustainability articles manually labelled by the assistants, and GPT 4o-Turbo to extend the set of greenwashing articles. We describe in this section the prompts we use in each step of the classification, as well as the extensive prompt engineering process that led to the final prompt.

B.1 Final sustainability prompt

We use the following user message[17] in the GPT prompt that identifies sustainability articles:

> ### Role : You are a **Senior Financial Analyst** at a leading hedge fund. Your expertise lies in identifying **any explicit allusion to the **integration of environmental and social considerations into corporate behavior and financial decision-making** from media articles.
>
> ### Objective : From the provided journal article (‘{input_text}‘), extract **explicitly stated** allusions to the **integration of environmental and social considerations into corporate behavior and financial decision-making** and estimate the percentage of the text that is about the topic.
>
> If the topic is mentioned (score > 0), your explanation must start with **direct quotations** from the article showing that the topic is being discussed, followed by context if needed.
>
> Format your answer as follows:
> - Environmental and social considerations: float between 0 and 1. You must strictly follow this format. Do not skip or reorder any lines.

B.2 Final greenwashing prompt

Due to many failed attempts with GPT 4o-mini, we use the most recent but also more costly GPT 4 Turbo to identify greenwashing articles among sustainability articles.

We use the following user message.

> ### Role : You are a **Senior Journalist** at a leading newspaper. Your expertise lies in **identifying sections of text that explicitly state or quote suspicion, accusations, or investigations of exaggerated, misleading, mislabelled, false

[17] The user message is the component of a prompt that contains the instruction given to the LLM.

or unsubstantiated companies' claims or marketing about their environmental and social actions or green products.**

Objective : From the provided journal article ('{input_text}'), extract sections of text that explicitly state or quote suspicion, accusations, or investigations of exaggerated, misleading, mislabelled, false or unsubstantiated companies' claims or marketing about their environmental and social actions or green products. Do not **infer, invent or speculate**. Only report what is stated in the text.

Exclusion criteria: - Exclude text that accuse an entity or person of exaggerating environmental risks. - Exclude text of companies' actions that align with their own claims. - Exclude text of green actions' drawbacks or unintended consequences if they are not considered intentionally hidden or understated. - Exclude text of lobbying to prevent adoption of a regulation.

Format your answer as follows, filling in each line carefully:
- 1 if you have identified text as described above, 0 otherwise

Do not skip or reorder any lines.

B.3 Validation of the GPT classification

We assess the performance of the GPT classification by comparing the predicted classes to the ones assigned by the assistants in the manually labelled set of articles. Whereas our final prompts achieved as match close to 95% for both steps of the classification, this performance is the result of a long prompt engineering process. We found, indeed, that the classifications were extremely sensitive to the formulation of the prompt. The primary challenge lies in GPT's extrapolative tendencies: when not carefully prompted, the model may infer a company's involvement in greenwashing based on its own interpretation, rather than adhering strictly to the content of the media article.

C 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.

C.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.

To use the method of Engle, Giglio, Kelly, Lee, and Stroebel (2020) in the greenwashing classification, we first need to select a set of documents that we know are related to greenwashing. We use the articles labelled as greenwashing in our training set. A greenwashing vocabulary is extracted from this corpus. Wall Street Journal articles are classified as greenwashing if their cosine similarity with this vocabulary is larger than a chosen threshold. The optimal threshold is optimized using the validation set.

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

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

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

D Entity Linking

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

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

E ESG confusion and negative sentiment indices

E.1 ESG confusion

Below is the prompt that was used to build a training set with articles on ESG confusion.

Role : You are a **Senior Journalist** at a leading newspaper. Your expertise lies in **identifying sections of text that explicitly state or quote the difficulties of measuring the sustainability or the performance of ESG assets. **

Objective : From the provided journal article (‘{input_text}‘), extract sections of text that explicitly state or quote the difficulties of measuring the sustainability or the performance of ESG assets. Do not **infer, invent or speculate**. Only report what is stated in the text.

Distinguish discussions around the following topics:
1 - the difficulties of measuring financial assets’ sustainability: this includes sections of text that mention divergences between ESG ratings, the difficulty of measuring or quantifying sustainability and ESG criteria, the lack of standardization in the reporting of ESG data, the issues about the quality of reported ESG data, and the complexity of ESG ratings and labels.
2 - the difficulties of measuring sustainable assets’ performance: this includes explicit mentions of the disagreements about ESG assets’ performance and risk, and financial managers making misleading, exaggerated or false claims about ESG assets’ performance and risk.

Format your answer as follows, filling in each line carefully:
- Presence: 1 if you have identified text as described above, 0 otherwise Do not skip or reorder any lines.

E.2 ESG backlash

Below is the prompt that was used to build the ESG backlash index from opinion pieces.

Role : You are a **Senior Financial Analyst** at a leading hedge fund reviewing opinion pieces from a newspaper.

Objective : Read the following article (‘{input_text}‘) and determine:
1. Does it discuss ESG (Environmental, Social, and Governance) principles, especially in relation to business strategies, financial decisions, or corporate responsibility?
2. Does it express a negative or critical view of the adoption or effectiveness of ESG principles?

Format your answer as follows:
- Criticism of ESG: 0 or 1; respond with 1 if the article discusses ESG *and* contains a critical or negative tone toward ESG or its use in finance/business. Respond with 0 if the article either doesn’t discuss ESG, or discusses it in a neutral or positive way.
- Quote and explanation: provide quotes from the articles to justify your answer and an explanation You must strictly follow this format. Do not skip or reorder any lines.

F Panel Regressions - Robustness tests

This section presents the results of alternative specifications of the panel regressions described in Section 5.3.

Table A1: Flow-Performance Sensitivity and Greenwashing for all Funds

This table reports the results of panel regressions of all fund flows (relative to size) on determinants. These determinants include fund performance, measured as the fund's past 1-year returns, fund sustainability measured a dummy that is 1 if the fund has a 5-globe (resp. 1-globe) Morningstar rating, the greenwashing index and in columns (4) to (6) the firm sustainability indices, as well as the interaction of these variables. All explanatory variables are lagged by one month except past performance. Controls include lagged fund net assets (log), lagged share class size (log), the adjusted net expense ratio and the class age. Category-year fixed effects are included, using the categories of fund defined by Morningstar. Standard errors are double-clustered by fund and time. We present the t-statistics in the parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
1-year return	0.030***	0.027***	0.037***	0.039***	0.035***	0.048***
	(0.004)	(0.003)	(0.005)	(0.010)	(0.009)	(0.012)
$D_{5\text{ globes}}$	0.001	0.001	0.001	0.001	-0.001	0.005*
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{1\text{ globe}}$	0.000	0.001	-0.001	-0.001	0.000	-0.005**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
$D_{5\text{ globes}}\times$ 1-year return	0.009***	0.011***	0.005	0.025***	0.032***	0.006
	(0.003)	(0.003)	(0.004)	(0.009)	(0.009)	(0.014)
$D_{1\text{ globe}}\times$ 1-year return	0.009***	0.010***	0.006	0.032***	0.036***	0.025*
	(0.002)	(0.003)	(0.004)	(0.009)	(0.011)	(0.014)
$GW_{m-1}\times$ 1-year return	-1.374	-0.742	-2.846**	-0.382	0.142	-1.581
	(1.089)	(0.990)	(1.374)	(1.649)	(1.500)	(2.087)
$GW_{m-1}\times D_{5\text{ globes}}$	0.204	0.129	0.289	0.208	-0.028	0.597*
	(0.173)	(0.210)	(0.242)	(0.246)	(0.274)	(0.362)
$GW_{m-1}\times D_{1\text{ globe}}$	-0.309*	-0.420**	-0.137	-0.438**	-0.469**	-0.477
	(0.158)	(0.166)	(0.260)	(0.207)	(0.233)	(0.312)
$GW_{m-1}\times D_{5\text{ globes}}\times$ 1-year return	0.566	-0.570	3.211**	2.426**	1.736	3.662*
	(0.812)	(1.014)	(1.269)	(1.169)	(1.330)	(2.054)
$GW_{m-1}\times D_{1\text{ globes}}\times$ 1-year return	-1.010*	-0.851	-1.154	1.968*	2.482**	1.140
	(0.613)	(0.726)	(1.019)	(1.011)	(1.147)	(1.643)
$Firm\ sust_{m-1}\times$ 1-year return				-0.192	-0.173	-0.238
				(0.192)	(0.178)	(0.247)
$Firm\ sust_{m-1}\times D_{5\text{ globes}}$				-0.006	0.037	-0.084
				(0.035)	(0.036)	(0.057)
$Firm\ sust_{m-1}\times D_{1\text{ globe}}$				0.036	0.014	0.090**
				(0.027)	(0.033)	(0.045)
$Firm\ sust_{m-1}\times D_{5\text{ globes}}\times$ 1-year return				-0.324**	-0.438***	-0.021
				(0.159)	(0.169)	(0.279)
$Firm\ sust_{m-1}\times D_{1\text{ globe}}\times$ 1-year return				-0.516***	-0.567***	-0.428*
				(0.166)	(0.214)	(0.251)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	911554	636829	274725	911554	636829	274725
N. of groups	15534	10794	4740	15534	10794	4740
R^2	0.044	0.050	0.028	0.045	0.051	0.029

Table A2: Flow-Performance Sensitivity and Greenwashing with Lagged Returns

This table reports the results of panel regressions of equity fund flows (relative to size) on determinants. These determinants include fund performance, measured as the fund's past monthly returns, fund sustainability measured a dummy that is 1 if the fund has a 5-globe (resp., 1-globe) Morningstar rating, the greenwashing index and in columns (4) to (6) the firm sustainability index, as well as the interaction of these variables. All explanatory variables are lagged by one month. Controls include lagged fund net assets (log), lagged share class size (log), the adjusted net expense ratio and the class age. Category-year fixed effects are included, using the categories of fund defined by Morningstar. Standard errors are double-clustered by fund and time. We present the *t*-statistics in the parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Lagged return	0.006	0.011	-0.004	-0.004	0.005	-0.025
	(0.011)	(0.010)	(0.015)	(0.026)	(0.024)	(0.037)
$D_{5 \text{ globes}}$	0.002**	0.001**	0.002	0.002	0.000	0.005
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{1 \text{ globe}}$	0.001**	0.002***	0.001	-0.003	-0.001	-0.007**
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{5 \text{ globes}} \times$ Lagged return	0.026***	0.031***	0.018	0.028	0.036	0.010
	(0.010)	(0.009)	(0.014)	(0.030)	(0.025)	(0.042)
$D_{1 \text{ globe}} \times$ Lagged return	0.006	0.010	0.002	-0.004	-0.003	-0.003
	(0.008)	(0.008)	(0.014)	(0.025)	(0.025)	(0.045)
$GW_{m-1} \times$ Lagged return	6.005	4.742	8.755	6.096	5.030	8.304
	(4.730)	(3.777)	(7.041)	(5.444)	(4.574)	(7.706)
$GW_{m-1} \times D_{5 \text{ globes}}$	0.684***	0.523***	0.978***	0.715**	0.398	1.264***
	(0.187)	(0.200)	(0.284)	(0.291)	(0.309)	(0.422)
$GW_{m-1} \times D_{1 \text{ globe}}$	-0.244	-0.204	-0.374	-0.652***	-0.470**	-1.112***
	(0.176)	(0.170)	(0.295)	(0.221)	(0.234)	(0.390)
$GW_{m-1} \times D_{5 \text{ globes}} \times$ Lagged return	-2.561	-4.675	1.499	-2.159	-4.423	1.862
	(3.781)	(3.585)	(4.998)	(4.940)	(4.587)	(6.510)
$GW_{m-1} \times D_{1 \text{ globes}} \times$ Lagged return	3.071	3.107	2.351	0.974	1.225	-0.116
	(3.146)	(2.729)	(5.504)	(3.071)	(3.146)	(5.487)
$Firm\ sust_{m-1} \times$ Lagged return				0.182	0.090	0.401
				(0.618)	(0.544)	(0.873)
$Firm\ sust_{m-1} \times D_{5 \text{ globes}}$				-0.010	0.028	-0.076
				(0.046)	(0.045)	(0.070)
$Firm\ sust_{m-1} \times D_{1 \text{ globe}}$				0.090***	0.059	0.163***
				(0.035)	(0.038)	(0.062)
$Firm\ sust_{m-1} \times D_{5 \text{ globes}} \times$ Lagged return				-0.045	-0.100	0.148
				(0.632)	(0.551)	(0.894)
$Firm\ sust_{m-1} \times D_{1 \text{ globe}} \times$ 1-year return				0.279	0.302	0.198
				(0.486)	(0.495)	(0.885)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	602790	408584	194206	602790	408584	194206
N. of groups	8933	5991	2942	8933	5991	2942
R^2	0.027	0.032	0.013	0.028	0.032	0.014

Table A3: Flow-Performance Sensitivity and Greenwashing Star Ratings

This table reports the results of panel regressions of equity fund flows (relative to size) on determinants. These determinants include fund performance, measured as the fund's Morningstar star ratings, fund sustainability measured a dummy that is 1 if the fund has a 5-globe (resp, 1-globe) Morningstar rating, the greenwashing index and in columns (4) to (6) the firm sustainability index, as well as the interaction of these variables. All explanatory variables are lagged by one month except funds' past performance. Controls include lagged fund net assets (log), lagged share class size (log), the adjusted net expense ratio and the class age. Category-year fixed effects are included, using the categories of fund defined by Morningstar. Standard errors are double-clustered by fund and time. We present the t-statistics in the parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
$D_{5\ \text{stars}}$	0.018***	0.015***	0.023***	0.026***	0.020***	0.035***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.004)
$D_{5\ \text{globes}}$	0.001	0.001	0.001	0.002	0.000	0.005
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{1\ \text{globe}}$	0.001	0.001*	0.000	-0.004**	-0.002	-0.007**
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{5\ \text{globes}} \times D_{5\ \text{stars}}$	-0.001	0.001	-0.003	-0.006	-0.005	-0.008
	(0.001)	(0.002)	(0.003)	(0.004)	(0.005)	(0.007)
$D_{1\ \text{globe}} \times D_{5\ \text{stars}}$	0.004**	0.008***	-0.001	0.008	0.011	0.003
	(0.002)	(0.002)	(0.003)	(0.006)	(0.008)	(0.009)
$GW_{m-1} \times D_{5\ \text{stars}}$	-0.459*	-0.240	-0.829*	0.436	0.410	0.490
	(0.275)	(0.241)	(0.428)	(0.305)	(0.282)	(0.540)
$GW_{m-1} \times D_{5\ \text{globes}}$	0.644***	0.480**	0.949***	0.719**	0.401	1.291***
	(0.186)	(0.200)	(0.282)	(0.292)	(0.310)	(0.433)
$GW_{m-1} \times D_{1\ \text{globe}}$	-0.178	-0.149	-0.284	-0.584***	-0.432**	-0.980***
	(0.148)	(0.157)	(0.261)	(0.184)	(0.209)	(0.360)
$GW_{m-1} \times D_{5\ \text{globes}} \times D_{5\ \text{stars}}$	0.293	0.379	0.007	-0.242	-0.270	-0.501
	(0.508)	(0.466)	(0.855)	(0.560)	(0.511)	(1.034)
$GW_{m-1} \times D_{1\ \text{globe}} \times D_{5\ \text{stars}}$	-0.980*	-1.299*	-0.590	-0.745	-1.117	-0.342
	(0.561)	(0.757)	(0.696)	(0.848)	(1.079)	(1.033)
$Firm\ sust_{m-1} \times D_{5\ \text{stars}}$				-0.179***	-0.129***	-0.266***
				(0.042)	(0.042)	(0.075)
$Firm\ sust_{m-1} \times D_{5\ \text{globes}}$				-0.019	0.017	-0.088
				(0.042)	(0.042)	(0.069)
$Firm\ sust_{m-1} \times D_{1\ \text{globe}}$				0.092***	0.064*	0.156***
				(0.032)	(0.036)	(0.060)
$Firm\ sust_{m-1} \times D_{5\ \text{globes}} \times D_{5\ \text{stars}}$				0.105	0.130	0.106
				(0.090)	(0.095)	(0.149)
$Firm\ sust_{m-1} \times D_{1\ \text{globe}} \times D_{5\ \text{stars}}$				-0.072	-0.060	-0.075
				(0.124)	(0.164)	(0.174)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	607756	411013	196743	607756	411013	196743
N. of groups	8966	6007	2959	8966	6007	2959
R^2	0.043	0.044	0.036	0.044	0.044	0.037

Table A4: Flow-Performance Sensitivity and Greenwashing in Finance

This table reports the results of panel regressions of equity fund flows (relative to size) on determinants. These determinants include fund performance, measured as the fund's past 1-year returns, fund sustainability measured a dummy that is 1 if the fund has a 5-globe (resp, 1-globe) Morningstar rating, the greenwashing-in-finance index and in columns (4) to (6) the firm sustainability index, as well as the interaction of these variables. All explanatory variables are lagged by one month except past funds' performance. Controls include lagged fund net assets (log), lagged share class size (log), the adjusted net expense ratio and the class age. Category-year fixed effects are included, using the categories of fund defined by Morningstar. Standard errors are double-clustered by fund and time. We present the t-statistics in the parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
1-year return	0.030***	0.028***	0.035***	0.050***	0.045***	0.059***
	(0.004)	(0.003)	(0.005)	(0.011)	(0.010)	(0.014)
$D_{5\ \text{globes}}$	0.001**	0.001*	0.002**	0.002	-0.001	0.006*
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{1\ \text{globe}}$	0.000	0.001*	-0.001	-0.002	-0.000	-0.006**
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{5\ \text{globes}}\times$ 1-year return	0.010***	0.012***	0.008*	0.023**	0.029**	0.007
	(0.003)	(0.003)	(0.004)	(0.011)	(0.012)	(0.015)
$D_{1\ \text{globe}}\times$ 1-year return	0.009***	0.010***	0.008*	0.028***	0.035***	0.017
	(0.002)	(0.003)	(0.004)	(0.010)	(0.013)	(0.017)
$GWFin_{m-1}\times$ 1-year return	-1.060	-0.467	-2.277	2.960	3.138	2.578
	(1.558)	(1.398)	(2.028)	(2.963)	(2.684)	(3.825)
$GWFin_{m-1} \times D_{5\ \text{globes}}$	0.487*	0.614*	0.150	0.560	0.332	0.889
	(0.289)	(0.320)	(0.502)	(0.422)	(0.443)	(0.725)
$GWFin_{m-1} \times D_{1\ \text{globe}}$	-0.083	-0.277	0.209	-0.564*	-0.562	-0.784
	(0.319)	(0.325)	(0.480)	(0.343)	(0.428)	(0.542)
$GWFin_{m-1} \times D_{5\ \text{globes}}\times$ 1-year return	1.066	-0.656	4.868**	3.423	2.422	4.841
	(1.497)	(1.771)	(1.988)	(2.262)	(2.628)	(3.304)
$GWFin_{m-1} \times D_{1\ \text{globes}}\times$ 1-year return	-3.300***	-2.815**	-4.147**	0.695	2.258	-2.222
	(1.114)	(1.305)	(1.732)	(1.975)	(2.460)	(3.246)
$Firm\ sust_{m-1}\times$ 1-year return				-0.421*	-0.380*	-0.501*
				(0.233)	(0.217)	(0.300)
$Firm\ sust_{m-1} \times D_{5\ \text{globes}}$				-0.009	0.035	-0.094
				(0.039)	(0.040)	(0.067)
$Firm\ sust_{m-1} \times D_{1\ \text{globe}}$				0.059*	0.036	0.119**
				(0.031)	(0.036)	(0.059)
$Firm\ sust_{m-1} \times D_{5\ \text{globes}}\times$ 1-year return				-0.260	-0.352	0.027
				(0.197)	(0.220)	(0.295)
$Firm\ sust_{m-1} \times D_{1\ \text{globe}}\times$ 1-year return				-0.414**	-0.527**	-0.202
				(0.199)	(0.258)	(0.317)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	602790	408584	194206	602790	408584	194206
N. of groups	8933	5991	2942	8933	5991	2942
R^2	0.041	0.045	0.028	0.042	0.045	0.030

Table A5: Flow-Performance Sensitivity and Greenwashing with Improved Firm Sustainability Index

This table reports the results of panel regressions of equity fund flows (relative to size) on determinants. These determinants include fund performance, measured as the fund's past 1-year returns, fund sustainability measured as a dummy that is 1 if the fund has a 5-globe (resp., 1-globe) Morningstar rating, the greenwashing index and in columns (4) to (6) the improved firm sustainability indix, as well as the interaction of these variables. Here, the firm sustainability index is "cleaned" from all greenwashing articles. All explanatory variables are lagged by one month except funds' past performance. Controls include lagged fund net assets (log), lagged share class size (log), the adjusted net expense ratio and the class age. Category-year fixed effects are included, using the categories of fund defined by Morningstar. Standard errors are double-clustered by fund and time. We present the t-statistics in the parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
1-year return	0.033***	0.029***	0.040***	0.039***	0.037***	0.044***
	(0.004)	(0.003)	(0.005)	(0.010)	(0.009)	(0.012)
$D_{5\text{ globes}}$	0.001	0.001	0.001	0.001	-0.001	0.006*
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{1\text{ globe}}$	0.001	0.001**	0.000	-0.002	0.000	-0.006**
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
$D_{5\text{ globes}} \times$ 1-year return	0.009***	0.011***	0.007	0.022**	0.029***	0.004
	(0.003)	(0.003)	(0.004)	(0.010)	(0.011)	(0.015)
$D_{1\text{ globe}} \times$ 1-year return	0.007***	0.008***	0.005	0.031***	0.036***	0.024*
	(0.002)	(0.003)	(0.004)	(0.009)	(0.011)	(0.014)
$GW_{m-1} \times$ 1-year return	-1.708	-0.996	-3.176**	-1.042	-0.279	-2.620
	(1.155)	(1.026)	(1.466)	(1.532)	(1.366)	(1.947)
$GW_{m-1} \times D_{5\text{ globes}}$	0.402*	0.328	0.509*	0.408	0.212	0.752**
	(0.209)	(0.228)	(0.298)	(0.251)	(0.264)	(0.369)
$GW_{m-1} \times D_{1\text{ globe}}$	-0.284	-0.289*	-0.356	-0.415*	-0.327	-0.708**
	(0.184)	(0.170)	(0.306)	(0.213)	(0.209)	(0.344)
$GW_{m-1} \times D_{5\text{ globes}} \times$ 1-year return	1.065	0.242	2.740**	2.358*	1.789	3.029
	(1.022)	(1.178)	(1.377)	(1.358)	(1.450)	(2.028)
$GW_{m-1} \times D_{1\text{ globes}} \times$ 1-year return	-1.077*	-0.891	-1.214	1.219	1.910*	0.207
	(0.635)	(0.730)	(1.108)	(0.927)	(1.003)	(1.501)
$Improved\ firm\ sust_{m-1} \times$ 1-year return				-0.145	-0.165	-0.098
				(0.194)	(0.180)	(0.249)
$Imp.firm\ sust_{m-1} \times D_{5\text{ globes}}$				-0.011	0.038	-0.111*
				(0.038)	(0.039)	(0.064)
$Imp.firm\ sust_{m-1} \times D_{1\text{ globe}}$				0.061*	0.029	0.137**
				(0.031)	(0.035)	(0.058)
$Imp.firm\ sust_{m-1} \times D_{5\text{ globes}} \times$ 1-year return				-0.273	-0.382**	0.056
				(0.183)	(0.195)	(0.295)
$Imp.firm\ sust_{m-1} \times D_{1\text{ globe}} \times$ 1-year return				-0.524***	-0.603***	-0.405
				(0.170)	(0.214)	(0.255)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	602790	408584	194206	602790	408584	194206
N. of groups	8933	5991	2942	8933	5991	2942
R^2	0.041	0.045	0.029	0.042	0.045	0.030