Decoding Sustainable Investment Strategies: Bridging Intentions and Outcomes

Keer Yang[a] Ayako Yasuda[b]

December 2025

Abstract

We study sustainability intent in U.S. mutual funds and whether stated objectives translate into differences in behavior. Using machine learning applied to fund prospectuses, we classify sustainable strategies as financial, moral, or impact oriented. Among 1,523 funds managing $1.7 trillion in 2023, 88% of assets are financially motivated, 10% morally motivated, and just 2% impact oriented. Financial funds tilt toward firms with high ESG ratings; moral funds rely on exclusionary screens and exhibit low flow–performance sensitivity. Only impact funds are associated with reductions in portfolio firms' carbon intensity. Most U.S. sustainable capital is therefore not deployed for real-economy impact.

Keywords: Impact Investing; Climate Finance; Machine Learning; ESG; Socially Responsible Investing; Ethical exclusion; Materiality; Sustainable Finance; Shareholder Voting

---

We are grateful for the comments of Vikas Agarwal, Aleksandar Andonov, Brad Barber, Ioannis Branikas, Hoa Briscoe-Tran, Murillo Campello, Will Cong, Yaniv Grinstein, Deeksha Gupta, Galina Hale, Xing Huang, Wei Jiang, Laura Lindsey, David Ng, Nadya Malenko, Marcus Opp, Kelly Posenau, Aneesh Raghunandan, Yuan Ren, Alberto Rossi, Mirco Rubin, Tao Shu, Kelly Shue, Paul Smeets, Denis Sosyura, Laura Starks, Artem Streltsov, Sunil Wahal, Xiang Yin, Adam Zhang, and conference and seminar participants at GSU CEAR-Finance Conference, CAFIN Conference, GRASFI 2024 Conference, Minnesota Family and Friends Conference, SGFIN Annual Research Conference on Sustainability, FMA Annual Meetings, Arizona State University, Audencia Business School, Cornell University, EDHEC, Tsinghua University, UC Davis, University of Amsterdam, University of Tokyo, York University, and Zhejiang University. Sayar Banerjee and Yan Jia provided valuable research assistance. We each acknowledge the generous funding from the UC Davis Academic Senate Committee on Research and the Center for Investor Welfare. All errors are our own.

[a] University of California, Davis, Graduate School of Management; E-mail: kkeyang@ucdavis.edu

[b] University of California, Davis, Graduate School of Management; E-mail: asyasuda@ucdavis.edu

1. Introduction

The 21st century presents unprecedented global challenges—climate change, biodiversity loss, rising inequality—that are increasingly salient to investors and policymakers. Sustainable investment has gained traction as one potential approach to addressing these issues while also delivering financial returns. However, sustainable investing encompasses fundamentally distinct objectives. Investors may allocate capital to sustainable funds for at least three reasons: to enhance financial performance through ESG integration, to align investments with moral or ethical values, or to induce real-world change in portfolio companies.[1]

These distinctions have material implications for the functioning of sustainable finance. From a policy perspective, achieving environmental and social goals through private capital markets requires that funds be genuinely impact-oriented, that households and institutions have preferences aligned with those goals[2], and that capital be effectively matched to the appropriate funds. Understanding the prevalence and behavior of these different fund types is therefore central to evaluating both the promise and the limitations of sustainable investment.

Despite the rapid growth in sustainable fund assets under management (AUM), relatively little is known about the motivations that drive these investment strategies or whether funds behave in ways consistent with their stated objectives. Existing classification systems—widely used by practitioners, regulators, and researchers—often fail to distinguish between fundamentally different sustainability goals.[3] In particular, financial, moral, and impact motivations are frequently conflated, and U.S. regulatory frameworks lack standardized definitions. As a result, empirical studies and policy debates often rest on untested assumptions about what sustainable funds aim to achieve.

This ambiguity has created space for competing narratives about ESG investing. Critics have argued that ESG funds pursue ideological goals at the expense of shareholder value, implicitly assuming that most such funds aim to induce real-world change and therefore underperform.[4] Supporters, meanwhile, have expressed frustration at the lack of observed

---

[1] Fama and French (2007), Benabou and Tirole (2010), Morgan and Tumlinson (2019), Chowdhry et al. (2019), Pedersen, Fitzgibbons, and Pomorski (2021), Pastor, Stambaugh and Taylor (2021), Broccardo et al. (2022), Oehmke and Opp (2025), Berk and van Binsbergen (2025), Landier and Lovo (2025), and Green and Roth (2025).
[2] See Riedl and Smeets (2017), Barber, Morse, and Yasuda (2021), Bauer, Ruof, and Smeets (2021), Heeb et al. (2023), Giglio et al. (2025), Bonnefon et al. (2025).
[3] Berg, et al. (2022), Christensen et al. (2022), Avramov et al. (2022), Raghunandan et al (2022), Diaz-Rainey et al. (2023), Allcott et al. (forthcoming).
[4] Rajgopal et al. (2025).

impact, assuming that funds promise real-world outcomes but fail to deliver. Both critiques hinge on a shared but largely unexamined assumption: that many ESG-labeled funds aim to reduce negative externalities. Whether that assumption holds in practice remains an open empirical question.

As Starks (2023) noted in her AFA presidential address, confusion over the meaning of sustainable investing often arises "due to differences in whether motivation arises from value or values"—that is, whether ESG characteristics are relevant because they are financially material or because they reflect nonpecuniary preferences (pp. 1837–1838). She calls for research that explicitly distinguishes between these motives and investigates their implications.

We take up this challenge by developing a scalable, data-driven method for classifying mutual funds by their sustainability objectives and analyzing whether fund actions align with those goals and lead to differential outcomes at portfolio companies.

We begin by delineating three core motivations for sustainable investing: financial, moral, and impact. Financially motivated funds seek to improve portfolio performance by integrating ESG-related risks and opportunities. These funds treat ESG inputs as instrumental tools for stock selection, risk reduction or identifying firms with superior intangible capital.[5] Morally motivated funds aim to align with investor values, often using negative screening to exclude sectors such as tobacco, fossil fuels, or weapons. Impact-oriented funds seek to generate measurable improvements in environmental or social outcomes. Rather than avoiding "brown" firms, impact funds may deliberately invest in them with the goal of driving change through active engagement.

This distinction between moral and impact investing is particularly salient. Both reflect nonpecuniary preferences, but they operate through different mechanisms. Moral investors avoid firms that conflict with their values, thereby reducing complicity in harm. Impact investors seek to reduce harm by improving firms' practices through ownership and engagement. For example, a moral fund may exclude a high-emissions utility; an impact fund may overweight the same firm with the intention of reducing its emissions over time through stewardship.[6] These strategies imply different portfolio compositions, stewardship behaviors, and expected outcomes.

[5] Financial motivation corresponds to "ESG-aware" investors in Pedersen, Fitzgibbons, and Pomorski (2021).
[6] Moral and impact motivation correspond to "narrow" and broader "impact" mandates, respectively, in Oehmke and Opp (2025).

In this paper, we investigate the prevalence, behavior, and effectiveness of three types of sustainable mutual funds—financial, moral, and impact. We address three core questions:

1. In what proportions do U.S. mutual funds pursue the three sustainable investing strategies?
2. Do funds act consistently with their stated sustainability goals?
3. Which types of fund investments, if any, are associated with real-world social or environmental improvements?

To classify fund intentions, we develop a fine-tuned BERT model that identifies how each fund frames sustainability—financially, morally, or in terms of real-world impact. This yields a fund-level measure of sustainability intent, which we validate on a labeled sample. We then use this measure as an explanatory variable in a series of economic tests. Specifically, we examine whether sustainability intent predicts differences in portfolio composition, fund performance, investor flows, stewardship behavior, and subsequent firm-level emissions.

We identify over 1,500 sustainable mutual funds managing $1.7 trillion in assets as of 2023—a fifteen-fold increase from $130 billion in 2014, when just under 150 such funds existed. Despite this growth, the vast majority of sustainable AUM remains financially motivated: 88% of assets are managed by financial funds, 10% by moral funds, and only 2% by impact funds. In dollar terms, just $29 billion is allocated to impact funds, representing only 0.08% of the total U.S. mutual fund market. This stark asymmetry is a key empirical finding.

The observed distribution of fund types reflects the joint outcome of three forces. First, the latent preferences of households and institutions shape demand for different sustainability goals.[7] Second, fund managers choose investment strategies based on their perceptions of that demand, as well as their own incentives and constraints.[8] Third, frictions in the search and matching process—driven by limited transparency, investor misunderstanding, or intermediary influence—affect the allocation of capital across fund types.[9] While we do not attempt to disentangle these drivers, the observed skew in capital allocation is itself informative: in the absence of regulatory or structural changes, most capital labeled as "sustainable" is unlikely to be deployed toward generating real-world impact.

[7] Krueger, Sautner, and Starks (2020), Heeb et al. (2023), Bonnefon et al. (2025).
[8] Gibson et al. (2022), Liang et al. (2022), Kim and Yoon (2023), Parise and Rubin (forthcoming).
[9] Lee et al. (2020), Heeb et al. (2025).

We show that sustainability intent predicts systematic differences in portfolio composition and fund characteristics. Financial funds tend to be larger and hold firms with high ESG ratings and low carbon intensity, consistent with ESG integration for risk management. Moral funds systematically exclude sin sectors and charge lower fees. Impact funds are smaller, charge higher expense ratios, and overweight high-emissions sectors, consistent with an engagement-oriented strategy.

We next examine whether sustainability intent is associated with differences in fund performance, investor behavior, and stewardship. While we find no meaningful differences in risk-adjusted performance across fund types, sustainability intent predicts systematic differences in investor flows and voting behavior. Moral funds exhibit lower flow–performance sensitivity, consistent with a less performance-driven investor base. Impact-oriented funds are more likely to support outcome-oriented environmental and social (ES) shareholder proposals both relative to financial and moral funds, and relative to other (e.g., disclosure-oriented) types of ES proposals.

To assess whether sustainability intent is associated with real-economy outcomes, we use a difference-in-differences design to examine whether firms newly added to the portfolios of different fund types exhibit changes in carbon intensity relative to matched firms. We find that firms added by impact-oriented funds subsequently reduce carbon intensity relative to matched controls, while no comparable pattern is observed for firms added by financial or moral funds.

Yet because impact-oriented funds manage a very small share of assets, their aggregate influence remains limited. Financial and moral funds largely act in ways that are consistent with their stated objectives, but those objectives do not include generating measurable environmental or social improvements at the firm level. Absent changes in regulation, disclosure, or investor understanding, most capital labeled as sustainable is therefore unlikely to generate measurable real-economy impact.

Contribution and Related Literature. The paper makes several contributions to existing literature. First, it provides a new classification of mutual funds based on sustainability intent—financial, moral, or impact—derived from the fund's own language and validated at the sentence level. There is an extensive literature that compares sustainable funds (and investors) to traditional funds and examine how their motivations, behavior and outcomes differ. Several studies show that prosocial investors are more likely to choose sustainable funds (Riedl and Smeets, 2017), some accept lower returns for sustainability (Bauer, Ruof, and Smeets, 2021), and there is heterogeneity in willingness to pay for impact across legal environments, geographies, and time (Barber, Morse, and Yasuda (2021), Baker, Egan, and Sarkar (2022)). Others find that investors' primary motivation for investing in sustainable funds is higher returns (Giglio et al. (2023)). Our work contributes to this literature by classifying the umbrella category of sustainable funds into 3 categories (financial, moral, and impact) based on distinct investment goals. In a related study, Lowry, Wang, and Wei (2025) classify ESG funds into high vs. low levels of commitment using Lewellen and Lewellen (2022)'s "incentives to engage" measure. Their commitment measure is designed to capture presence of value-driven (i.e., financial) sustainability intent. Our approach captures a value-driven intent using a different, text-based method, and it further identifies and distinguishes moral and impact – the two values-driven (i.e., nonpecuniary) sustainability goals.

Our classification also relates to the growing literature that employs machine learning and AI as methodological tools in financial research generally, and more specifically in sustainable finance.[10] Existing literature has documented a lack of coherence among different ESG rating agencies and a challenge this poses (Cohen et al. (2020), Berg et al. (2022), Christensen et al. (2022)). This insight, coupled with the advance of AI and natural language processing (NLP) models, have motivated researchers to develop new sustainability measures based on textual data.[11] In mutual funds research, Abis (2022) and Abis and Lines (2024) apply unsupervised topic models to classify mutual funds' overall investment strategies. We contribute to this literature by applying a BERT model to mutual fund prospectuses and extracting their sustainability intent. We build and perform both BERT vs. generative AI models against a labeled sample ("the ground truth") and find that the BERT model performs better and thus preferrable for our application to the U.S. fund prospectuses.

By applying this classification, we find that the majority of funds using ESG language articulate primarily financial objectives, with relatively few expressing a distinct impact-oriented mandate. Our results complement those in two contemporaneous working papers that examine related questions. Abis, Buffa, and Sadashivam (2024) manually classify mutual funds'

[10] Gentzkow, Kelly, and Taddy (2019) provide a comprehensive survey of textual analysis in economics and finance research. Goldstein, Spatt, and Ye (2021) offer a broader overview of "big data" applications in finance. Cao et al. (2024) examine the use of AI and machine learning for analyzing alternative data, including textual data. Jiang, Wang, and Yang (2025) focus on the emerging field of ESG research based on AI methodologies.
[11] Baker et al. (2024), Engel et al. (2020), Li et al. (2022), Sautner et al. (2022), Briscoe-Tran (2022), Bingler et al. (2024), Rajan et al. (2022), Michaely et al. (2023), and Duchin et al. (2025).

sustainability strategies based on fund prospectuses into opportunistic, exclusionary, and impact, and find that only a small fraction of funds are impact-oriented. Beyond mutual funds, Edmans, Gosling, and Jenter (2024) survey global institutional investors to elicit their ESG objectives and show that financial considerations dominate, with impact constrained by costs and frictions. Our approach differs from these studies in that we use a manually built labeled sample ("the ground truth") to train and test the BERT model and then apply the validated model at scale to the universe of mutual funds. We also examine the relationship between funds' sustainability intent and their stewardship and outcomes at the portfolio company level.

Second, the paper connects to the literature that examines the implication of investor sustainability intent on fund behavior and stewardship. Several studies find that sustainable mutual funds exhibit less flow-performance sensitivity than conventional mutual funds (e.g., Bollen (2007), Benson and Humphrey (2008), Renneboog, ter Horst, and Zhang (2011), Bialkowski and Starks (2018), and El Ghoul and Karoui (2017)). Hartzmark and Sussman (2019) document that the introduction of Morningstar Globe ratings attracts investor flows, consistent with nonpecuniary preferences. However, Giglio et al. (2023), Barber, Morse, and Yasuda (2021), and Gantchev, Giannetti and Li (2024) provide evidence that financial motivation is the dominant driver for some of the investors in sustainable funds. Our work complements this literature by showing that funds with moral sustainability intent exhibit lower flow-performance sensitivity, whereas sustainable funds with financial or impact goals do not. While the moral and financial funds results suggest that on average values investors sort into moral funds and value investors sort into financial funds, the impact fund result suggests information frictions, mixed clientele, or some other mechanism that warrants further exploration.[12]

Another strand of literature examines engagement and shareholder voting on environmental and social issues.[13] Some of the prior works focus on value-driven investors (e.g., hedge funds), while few studies explicitly examine voting behavior of investors or funds with nonpecuniary preferences (Shanker (2025)). Our work contributes to this literature by studying the differences across financial, moral, and impact intent on shareholder proposal votes. We

---

[12] Bonnefon et al. (2025) and Heeb et al. (2023) suggest that retail investors' preferences are more aligned with moral funds than impact funds. Heeb et al. (2025) suggests that retail investors overestimate the impact of generic sustainable products.

[13] Dimson, Karakas, and Li (2015, 2025), Flammer et al. (2021), Gantchev et al. (2022), Heath et al. (2023), Hoepner et al. (2024), Naaraayanan et al. (2021), Akey and Appel (2020), Chen et al. (2020), He, Kahraman, and Lowry (2023), Michaely et al. (2024), Lowry, Wang, and Wei (2025), Couvert (2025), and Shanker (2025).

show that while all three types of sustainable funds exhibit higher propensity to vote in support of environmental and social shareholder proposals than conventional funds, impact funds exhibit the highest support, and this is driven by their higher support of proposals related to outcomes at portfolio firms rather than greater disclosures.

Third, the paper relates to the literature that connect the sustainability objective to real-economy outcomes. Lowry, Wang, and Wei (2025) focus on cross-sectional variation in strength of financial motivation of sustainable funds and show that firms bought by more highly motivated ("committed") funds and intensively engaged by them following severe ESG incidents subsequently experience decreases in both their ESG risk index and their emissions. These strongly financially motivated funds view engagements and positive changes in real outcomes as value drivers, since severe ESG incidents elevate the potential negative impact of not achieving such outcomes on firm valuations, and act accordingly.

Our paper adds to this strand of literature by comparing the real-economy outcomes of investments by impact funds vs. financial and moral funds. Theoretical models of investor nonpecuniary preferences in Oehmke and Opp (2025) show that a sustainable fund has real-world impact (i.e., changes the firm's technology choice) if and only if the fund sufficiently prioritizes the *reduction* in negative externalities from the fund's investment, and is willing to invest in brown firms if that results in negative social cost reduction. This preference corresponds to impact orientation in our paper. In contrast, if a sustainable fund only cares about the *level* of negative externalities generated by firms in its portfolio, such a fund only finds it optimal to invest in already green firms and is therefore does not generate impact.[14] This preference corresponds to moral motivation. We empirically study these predictions in several ways. First, we show that impact funds systematically hold larger portfolio weights in high-emission sectors such as utilities, which also have lower average E scores. Second, we show that impact funds have higher propensity to support E and S shareholder proposals than financial and moral funds, particularly outcome-oriented proposals. Finally, we show that firms held by impact-oriented funds reduce emission intensity relative to matched firms, while no such pattern

[14] In a different framework, Hartzmark and Shue (2023) also find that sustainable investing strategies of tilting away from brown firm and towards (already) green firms—employed via exclusion by moral funds or ESG integration by financial funds—generates a counterproductive negative impact. Also see Gupta et al. (2026), Dangl et al. (forthcoming), and Dasgupta et al. (2025).

is observed for financial or moral funds.[15] Our findings provide strong support for the effectiveness of impact funds that prioritize reduction in negative externalities as results of investments, as opposed to minimizing levels of carbon footprint of fund portfolios.

The remainder of the paper proceeds as follows. Section 2 describes the data and our classification of fund sustainability intent. Section 3 examines how sustainability intent maps into portfolio composition and ESG-related firm characteristics. Section 4 analyzes differences in fund performance, investor flows, and voting behavior. Section 5 studies whether sustainability intent is associated with firm-level emissions outcomes. Section 6 concludes.

2. Empirical Method and Data

This section describes our empirical methodology and data. In Section 2.1, we present our classification of investors' sustainability goals. In Section 2.2, we introduce an objective empirical method that classifies investors' sustainability goals based on the textual descriptions of their investment strategies and apply it to all U.S. actively managed mutual funds and ETFs. In Section 2.3, we describe the additional data sources used in our study.

2.1 Classification of Sustainability Goals

We classify investors' sustainability goals along two conceptual dimensions that build on existing frameworks in the literature. The first distinguishes whether sustainability characteristics enter the investor's objective function directly or only through their implications for financial performance. Following Starks (2023), we refer to these as *values*-driven versus *value*-driven investing. Value-driven investors incorporate ESG information insofar as it affects expected returns or risk, while values-driven investors derive utility directly from environmental or social characteristics, independent of financial materiality. This distinction closely parallels Pedersen, Fitzgibbons, and Pomorski's (2021) categorization of ESG-aware versus ESG-motivated investors.

We label value-driven sustainability objectives as financial value. Investors with this objective maximize standard risk–return preferences while using ESG data to identify and manage financially material risks and opportunities, such as regulatory exposure, litigation risk,

[15] In related work, Atta-Darkua, Glossner, Krueger and Matos (2023) find that institutional investors joining climate pledges reduce their public equity portfolio emissions mostly by re-balancing towards low-emission stocks rather than by engagement.

reputational capital, or transition risk. ESG characteristics matter instrumentally, not intrinsically.

Our second dimension differentiates values-driven investors based on how they assess the desirability of sustainability outcomes. Drawing on philosophical terminology, we distinguish between impact (consequentialist) and categorical morality (deontological) objectives. Impact-oriented investors evaluate investments based on the external outcomes they generate—for example, reductions in emissions or improvements in workplace safety—and therefore tend to emphasize engagement and stewardship. Categorical moral investors, by contrast, evaluate investments based on adherence to ethical principles, such as avoiding particular industries or activities, regardless of downstream consequences.

This distinction has direct implications for portfolio construction. An impact-oriented investor may intentionally invest in firms with poor initial environmental performance in order to influence future behavior, whereas a morally motivated investor categorically excludes such firms to avoid moral compromise. Although both investors possess nonpecuniary preferences, their optimal portfolios, engagement strategies, and expected outcomes differ sharply.

Failing to distinguish between these objectives can lead to misleading empirical interpretations. Many studies implicitly assume that all values-driven investors prefer high-ESG firms, yet this assumption is inconsistent with impact investing, which often targets firms or sectors with substantial scope for improvement. Moreover, commonly used ESG ratings—such as MSCI scores—are primarily designed to capture financially relevant risk exposures rather than moral alignment or real-world impact. As a result, identical ESG scores may be interpreted very differently across investor types.

Taken together, these two dimensions yield three distinct sustainability goals: financial value, categorical morality, and impact. This framework provides a disciplined basis for classifying sustainable funds by stated intent and for generating testable predictions about portfolio composition, investor behavior, stewardship, and real-economy outcomes. In Section 2.2, we operationalize this framework using a text-based classification of mutual fund prospectus disclosures.

2.2 Sustainable Fund Classification

Currently there are no formal regulations in the U.S. that identify sustainable funds or distinguish between the three sustainable fund types. Thus, a household that wants to invest in a sustainable fund needs to either conduct independent search or rely on third-party service providers (e.g., a financial adviser, an investment newsletter, a website with annual top fund manager lists) to identify a fund that meets their sustainability goal. We propose an empirical method of classifying an investor's sustainability goals based on the text description of their investment strategy.

2.2.1 A Text-Based Empirical Method

We examine the "Principal Investment Strategy" section of the prospectus because it is a required disclosure and is intended to "tell you how the fund intends to achieve its investment objective" (U.S. SEC, 2016). Funds typically discuss their sustainability goals in this section. We extract text from the "Principal Investment Strategy" section of the fund prospectus for our classification. Since the fund also discusses how it intends to achieve its non-sustainability objectives (i.e., financial goals) in the same section, we first use an ESG keywords list to isolate sentences that discuss ESG-related topics and discard the rest.[16] We download the prospectuses (summary prospectus (Form 497K) and/or statutory prospectus (Form 485)) of the funds on the list from the SEC EDGAR system and extract the Principal Investment Strategy section from 2014 to 2023, resulting in a total of 75,886 fund-quarter observations.[17]

We build a supervised machine-learning model-based method that uses a manually created sample of classified texts to train the model and then leverages the model's ability to classify a large body of texts objectively and consistently. Since the three investor goals are distinguished from each other in how non-financial traits/outcomes are valued by investors, not what topics or data points the investors track, it is critical that the method can extract the investor's intent, not just keywords they use. More generally, we need a method that interprets a whole sentence to extract its meaning, not merely word pairs or topic nouns.

For this purpose, we use a Bidirectional Encoder Representations from Transformers (BERT) model. BERT is a natural language processing (NLP) method released by Google in 2018. The innovative feature of this model is its ability to process words in relation to one another within a given text. BERT comes pre-trained on a large source of text provided by

---

[16] Our ESG keyword list is sufficiently broad that any sentence from the Principal Investment Strategy section without ESG-related keywords is interpreted as discussing investment strategies unrelated to sustainability. The list of ESG keywords we used for this screening step is provided in the Appendix A1.

[17] Each fund–quarter observation is assigned the most recent available Principal Investment Strategy.

Google and is ready to be used for natural language process tasks. The pre-trained model can then be fine-tuned with a smaller training data sample that we provide and then used to conduct specific NLP tasks such as classification. Our empirical method uses custom-trained BERT models to classify sentences from fund prospectuses as “financial value,” “categorical morality,” or “impact investing” at the sentence level and then constructs fund-level measures based on these sentence classifications. To train and test the BERT model, we need to build a training and testing sample of example sentences that sufficiently and distinctly express each of the three sustainable investing goals. We sample 362 of the 75,886 fund-quarter observations and identified 3,575 ESG-related sentences for manual classification. To construct our sample, we oversample true positives (those classified as Financial, Moral, or Impact) from imbalanced samples to improve classification balance and performance, which is a common strategy in machine learning and textual analysis (He and Garcia, 2009). Among the 3,575 ESG-related used for training and testing, 2,438 are drawn from the prospectuses a random subset of U.S. sustainable mutual funds compiled and published by Morningstar. We oversample this group because we expect these funds to be more likely to describe their sustainability goals in the investment strategy section of their fund prospectuses than other funds. We supplement the training and testing sample using funds not included in the Morningstar fund list.[18] These 3,575 sentences are then manually classified by the authors as “financial”, “moral”, “impact”, or “unclassified”.

To label a sentence as “financial value”, we require that the sentence states that the investor uses ESG (or non-financial) information for the purpose of improving financial performance. To label a sentence as “categorical morality”, we require that the sentence states that the investor excludes certain categories of investments (e.g., industries) from the portfolio for (implicitly) ethical reasons. When a sentence explicitly states that the purpose of exclusion is not related to ethics or SRI, then such a sentence fails to meet the “categorical morality” criteria. Finally, to label a sentence as “impact investing”, we require that the sentence states that the investor uses ESG information specifically for externality considerations. When a sentence also

[18] There is a priori no fixed sample size that is appropriate to be used as a training and testing sample. In selecting the fund-quarters to be included in the training and testing sample, we opted to use a given fund only once. This is because a fund prospectus text tends not change much from year to year, and our aim is to build a training sample with as much variation in sentence structures and phrase patterns as possible. More details on the training and testing sample construction are provided in Internet Appendix Section 2.

mentions financial performance as the motivation for monitoring externality, we label it as both "financial value" and "categorical morality"/ "impact investing". We labeled a total of 729 sentences as "financial", 288 sentences as "moral", and 305 sentences as "impact".[19] Among the 3,575 manually classified observations, we randomly split the data into an 80% training sample and a 20% testing sample using stratified sampling to ensure balanced class representation.

Once we have built a training and testing sample, the BERT model is trained on the training sample to learn the types of sentences that qualify as "financial value", "categorical morality", and "impact investing". These three categories are trained separately. We then ask the model to classify sentences in the testing sample into "financial", "moral", and "impact". As we also have our manually labeled "true" classification, we can compare the predictions made by the BERT model with our manual labels by calculating four different model performance metrics: accuracy, precision, recall, and f1.

[Insert Table 1: BERT Model Performance]

Accuracy is the ratio of the sum of true positives and true negatives divided by the total number of observations. Precision is the ratio of true positives divided by the sum of true positives and false positives. Recall is the ratio of true positives divided by the sum of true positives and false negatives. f1 is defined as $[\frac{2}{\frac{1}{Precision}+\frac{1}{Recall}}]$. We computed the model performance metrics separately for each label.

We find that the BERT model performs well in classifying the sustainability goal expressed in a sentence. Comparing the three goals, accuracy ranges from 91% (Financial) to 98% (Moral), implying that when a sentence is classified as either a given type (positive) or not (negative), it is correct in most cases. Precision ranges from 83% (Impact) to 85% (Moral), meaning when a sentence is classified as a given type (positive cases), most of them are true positives. Recall ranges from 66% (Impact) to 89% (Moral), meaning most of the true positive cases are classified as such. F1 ranges from 73% (Impact) to 87% (Moral), meaning most false negatives and false positives are correctly identified by the model. These performance levels are

[19] Note that the majority (2,278) of the sentences were unclassified either because it does not state any sustainability goal (though screened in because it contained some keywords), or because its sustainability statement was too ambiguous to fall clearly into one category. More details on the classification criteria, including example sentences for each of the 4 classification categories that we manually coded, can be found in Internet Appendix Section 1.

in line with other studies using BERT models in the literature (e.g., see Bingler et al. (2024) and Rajan et al. (2022)).

After establishing the reliability of our BERT model, we apply it to all sentences from the 75,886 fund–quarter observations and construct fund–quarter–level measures based on the sentence-level classifications. A fund is assigned the category with the largest percentage share of all sustainable sentences. This definition reflects the primary sustainable goal of each fund, as a fund's sustainability type is determined by the category with the largest share of sustainable sentences. Funds may also have secondary motives, represented by a smaller number of sustainable sentences from other categories.

When there is a two-way (three-way) tie for the largest percentage share, the fund is assigned to each of the tied categories. In such cases, the fund pursues more than one sustainable goal as a primary objective. This results in 1,067 fund–quarter observations being assigned to multiple categories, representing less than 2% of the total sample.

2.2.2 Comparison to GenAI Models

One alternative classification framework is to use generative AI (GenAI) models, such as ChatGPT, Google's Gemini series, and Claude, which can be prompted to classify sustainable investment goals. Two different approaches can be used to apply GenAI models to classify sustainable goals. First, researchers can directly ask a GenAI model to perform the classification without providing it with training data. Under this approach, the model's pretraining data implicitly serves as the ground truth rather than researcher-defined labels. This methodology is problematic because retail investors and the media have been documented to often misunderstand sustainable goals (see Heeb et al. (2025), Hartzmark and Sussman (2019), and Gantchev et al. (2024)), and such misunderstandings reflected in texts appearing in news articles and online discussions may be transmitted to the model through its pretraining data. Internet Appendix Table IA2 shows that the ChatGPT-5 model consistently exhibits lower performance than the BERT model.

The second approach to using GenAI is to provide the model with the training sample through prompting or retrieval-augmented generation (RAG), which can partially align GenAI outputs with researcher-defined classifications. However, this approach is subject to output variability and limited reproducibility due to the probabilistic nature of text generation in GenAI models. When researcher-defined ground truth is available, a supervised BERT model offers greater stability, full reproducibility, and avoids hallucinations.[20] Internet Appendix Section 3 provides a more detailed discussion.

2.3 Other Data

Our empirical analysis draws on several additional data sources. We construct fund-quarter-level ESG characteristics variables by merging the CRSP mutual fund database and the MSCI ESG ratings database. From the CRSP mutual fund database, we obtain the average market value of each stock held by the fund for the quarter, which we use as the portfolio weight of the security. From the MSCI ESG ratings database, we obtain MSCI's ESG ratings for stocks held by our sample mutual funds. Combining the two, we calculate the fund-quarter -level ESG ratings that are the weighted-average ESG ratings of all MSCI-rated stocks held by the fund. We use both the industry-adjusted and unadjusted ratings to address different hypotheses.

We construct fund–year–month–level fund flows and fund alphas from the CRSP Mutual Fund dataset. At the fund–year–month level, a fund's sustainability type is determined based on the most recently available quarterly fund type. We construct firm-month–level and fund–month–level Scope 1&2 and Scope 1&2&3 carbon emission intensity measures, defined as metric tons of emissions per million dollars of revenue, using the Trucost Environmental dataset. Trucost Environmental reports firms' annual emission intensity. We convert these annual measures into firm-month emission intensity using a weighted average of the two adjacent fiscal years, with weights equal to 12 minus the distance (in months) from each fiscal year. We use both industry-adjusted and unadjusted emission intensity measures. Fund–month–level Scope 1&2 and Scope 1&2&3 carbon emission intensity measures are constructed as the value-weighted average of firm-level emissions across all stocks with non-missing emission data held by the fund.

Finally, from the ISS database, we obtain the fund's voting records on shareholder proposals at each of the stocks it held. We match the ISS database FundID with the FundID in our sample using the methodology proposed by Sulaeman and Ye (2023). Specifically, we use SEC Form N-PX and the reference identifier (NPXFileID) provided by the ISS database to link each ISS fund to its corresponding SEC Form N-PX. We then use name matching to identify the fund identifier (series_id) from Form N-PX and link it to our database. We calculate the fund-

[20] Huang et al. (2023) also discuss and demonstrate a similar disadvantage of GenAI models.

level ESG activism variable defined as the likelihood of supporting shareholder proposals on environmental and social issues.

3. Sustainable Fund Characteristics

3.1 Overview of the U.S. Sustainable Funds Universe

Using our empirical method, we classify the universe of CRSP mutual funds into (i) Keyword-only (has 1 or more sustainability keywords in the investment strategy but no sentence containing a sustainable strategy), (ii) Sustainable (at least 1 sentence classified as containing a sustainable strategy), and (iii) Non-sustainable (no sustainable keywords in the investment strategy. Funds are classified each year using their annual prospectus disclosures. Figure 1 presents the annual sustainable fund count (Panel A) and assets under management (Panel B) during the sample period based on our BERT classification model. We find that 1,523 sustainable funds managed $1.7 trillion in 2023, compared to 145 sustainable funds managing $0.1 trillion in 2014. As % of the fund universe, the sustainable funds grew from 0.5% to 4.7% of the AUM, and from 1.4% to 12.9% of the fund count. Keyword-only funds grew significantly to 3,849 funds managing $5T as of 2023 and comprising 33% of the fund count and 14% of AUM. Together, sustainable and keyword-only funds grew faster than non-sustainable funds, pushing down non-sustainable funds' market share slightly from 86.3% (in 2014) to 81.3% of AUM in 2023. In terms of AUM per fund, sustainable funds are the smallest, managing about $1.12 billion per fund in 2023, vs. $1.30 billion for keyword-only funds and $4.53 billion for non-sustainable funds. Overall, sustainable funds are fast-growing but remain a minor segment of the mutual fund industry in the U.S.

[Insert Figure 1: Growth of Sustainable Mutual Funds, 2014–2023]

[Insert Figure 2: Distribution of Sustainable Fund Classifications]

In Figure 2, we present the breakdown the sustainable funds into financial, moral, and impact funds. The first observation is the increasing dominance of financial funds among the sustainable funds over time, and declining shares of impact and moral funds. In 2014, 71 out of 14521 funds (49%) were financial funds managing 82% of total AUM. By 2023, 1,211 out of

[21] The fund count sums up to 146 in 2014 in Figure 2 due to rounding, as 2 funds are dual tie-funds and are assigned 0.5 fund count each in 2 categories.

1,523 funds (80%) were financial funds managing 88% of total AUM. The AUM managed by financial funds grew 14 times from $108B to 1.470 trillion, while impact funds' AUM grew only 5 times from $6B to $29B during the sample period. Moral funds' growth is in the middle, growing 10 times from $18B to $176B. To the best of our knowledge, this overwhelming prevalence of financial value-driven sustainable funds has not been well documented in the literature.22

The findings suggest that while the impact funds' influence has increased in an absolute sense in the last decade, it has also diminished significantly relative to that of the financially driven sustainable funds. Any analysis that is conducted that treats sustainable funds as a homogenous group with a uniform sustainability goal is thus complicated by this time-varying heterogeneity of value vs. values (and moral vs. impact) investment objectives.

Table 2 presents the summary statistics of the sustainable funds in comparison to benchmark funds. Panel A presents the general fund characteristics for benchmark, financial, moral, and impact funds. We define the benchmark funds as CRSP equity non-index funds and ETFs whose prospectuses contain no sustainability keywords and are never on the Morningstar sustainable fund list during the sample period. A unit of observations is a fund-quarter from 2014 Q2 to 2023 Q4.

[Insert Table 2: Summary Statistics]

As shown in Figure 1 and 2, the number of sustainable funds grew significantly during the sample period. Averaged across the 39 quarters, there are approximately 314 financial funds, 90 moral funds, and 46 impact funds, versus 1,524 benchmark funds. All fund types invest approximately 90% of portfolio assets in equity. Financial and impact funds hold relatively more concentrated portfolios with fewer than 100 stocks on average, while moral funds hold nearly 200 stocks in the portfolio – more than benchmark funds. Benchmark funds are larger than sustainable funds on average, and among sustainable funds financial funds are the largest ($1.226B) while impact funds are the smallest with the average size of $327M. We also compare the expense ratios for funds with retail class offerings. Consistent with more resource intensive investment strategies of impact funds, their expense ratios at 1.47% are higher than other fund types, whereas moral funds, with primary focus on ex ante screening, offer the lowest expense

[22] Also see Abis et al. (2024) for contemporaneous and complementary evidence.

ratios of 1.27%. Financial funds offer expense ratios that are indistinguishable from benchmark funds (1.34% vs. 1.35%).

[Insert Figure 3: Sector Weights, ESG Ratings, and Emission Intensity]

In Panel B, we compare the sector weights of the three sustainable fund types with benchmark funds as well as the broad U.S. markets. We compute the sector weights by averaging across the annual year-end weights during the sample period. We also report the average E and S ratings and the emission intensity of each sector during the sample period. The 19 sector definitions are based on NAICS. We also present the same data in Figure 3, Panel A-C for ease of visualization. We find that financial funds hold portfolios with sector weights that closely track benchmark funds. In contrast, impact funds significantly overweight utilities and manufacturing, and underweight information and finance. Utilities and manufacturing contain operators with significant carbon footprints, pollution and other negative environmental externalities, partially reflected in relatively high emission intensity and low E ratings. In contrast, information and finance have relatively low emission intensity and high E ratings. As argued by Hartzmark and Shue (2023) and others, investing in already green firms is inconsistent with impact investing objectives. Instead, investing in brown sectors that have initially poor environmental records and induce improvements has a greater potential for impact creation. The significant sector weight deviations of impact funds are broadly consistent with pursuit of this impact objective. We further examine if impact funds use proxy voting to support ES shareholder proposals, and if the environmental performance of firms held by impact funds improves during the holding periods in later sections.

We further examine sectors commonly excluded by moral funds that express exclusionary investment goals in their prospectuses, including “Tobacco”, “Oil & Gas Extraction”, “Natural Gas Distribution”, “Coal Mining”, “Casino & Gambling”, and “Aerospace” (including weapons manufacturers). The results are presented in the Internet Appendix Table IA3. Overall, we find that moral funds with exclusionary investment goals in their prospectuses indeed underweight the sectors they claim to avoid.

3.2 ESG Characteristics of Fund Holdings

To assess whether sustainable funds differ in the ESG characteristics of their holdings, we examine the portfolio-level environmental and social (E and S) ratings across fund types[23]. We expect distinct patterns reflecting differences in fund objectives. Impact funds, by design, seek to generate real-world improvements and are therefore more likely to invest in firms with relatively low E and S ratings, with the intention of engaging and supporting their progress over time. In contrast, financial funds prioritize enhancing risk-adjusted returns and thus have incentives to allocate capital to firms with high E and S ratings, which may signal better management of ESG-related risks or greater exposure to financially material sustainability opportunities.[24] Benchmark funds, which do not explicitly integrate ESG considerations, are expected to lie between these two. Based on this reasoning, we formulate two empirical predictions: (i) financial funds hold portfolios with higher ESG ratings than benchmark funds, and (ii) impact funds hold portfolios with lower ESG ratings than financial funds.

Table 3 presents summary statistics and regression analyses testing these predictions.[25] Panel A reports the value-weighted ESG, E, and S ratings of fund holdings for benchmark, sustainable, financial, and impact funds. When constructing the fund-level quarterly analysis, we take the average fund holdings at the calendar-quarter level, which alleviate the concern that fund strategically time their trades in ESG stocks as documented in Parise and Rubin (2025). Variable definitions are provided in Appendix A2.

[Insert Table 3: ESG Characteristics of Fund Holdings]

We find that financial funds hold stocks with higher ESG ratings than benchmark funds. This result holds for both the E and S components. In contrast, impact funds hold stocks with substantially lower ESG ratings than financial funds, with the difference driven primarily by lower E ratings. However, these unadjusted comparisons may conflate time trends in ESG ratings with fund type effects: MSCI ESG scores have increased over time, and financial funds

---

[23] Pástor et al. (2023) find that institutional investors tilt their portfolios based on ESG characteristics.

[24] See, for example, Giglio et al. (2021), Pástor et al. (2022, 2024), Bolton and Kacperczyk (2021, 2023), and Aswani et al. (2024).

[25] In fewer than 2% of observations, there is a two-(or three-)way tie for the largest percentage share, in which case the fund is assigned to each of the tied categories. For example, in our regression analyses, a fund tied as both financial and moral is included as a financial fund for regressions comparing financial and benchmark funds and as a moral fund for regressions comparing moral and benchmark funds. This fund is designated as both a financial fund and a moral fund (i.e., both the financial and moral fund dummy is equal to one) for regressions comparing financial and moral funds.

have become more prevalent in later years. To address this, we estimate regressions with year–quarter fixed effects to compare ESG characteristics conditional on time.

Panel B reports the results of these regressions. In Panel B-1, we find that financial funds hold stocks with significantly higher ESG, E, and S ratings relative to benchmark funds in the same quarter. These results are robust to clustering standard errors at the fund level.

Panel B-2 compares impact funds to financial funds. We find that impact funds hold stocks with significantly lower E ratings, though the difference becomes statistically insignificant when standard errors are clustered. For S ratings, by contrast, impact funds hold stocks with significantly higher scores. Several explanations are plausible. As shown in Table 2 and Figure 3, impact funds allocate more heavily to sectors with lower average E ratings—such as utilities and manufacturing—while financial funds invest more in sectors with higher average E ratings, such as information technology and finance. Notably, utilities tend to have relatively high S ratings, whereas information and finance score lower on that dimension. Thus, differences in sector allocations may explain some of the observed variation in fund-level E and S scores. Alternatively, financial funds may select firms that are ESG "leaders" within sectors, while impact funds may deliberately invest in sector "laggards" to facilitate improvement. Finally, ESG ratings may imperfectly capture the types of real-world outcomes—such as emissions reduction or worker safety—that impact funds seek to promote.

To disentangle whether these differences are driven by sector allocation or within-sector stock selection, we conduct two decomposition exercises reported in Panel C. In columns (1)-(3), we assign each stock its *sector-average* score and compute the fund's sector-weighted ESG rating, thereby isolating the effects of sector tilts on fund-level ESG characteristics. In columns (4)-(6), we instead construct sector-adjusted ESG, E, and S scores by subtracting each stock's sector average (at the 6-digit NAICS level) and value-weighting within the portfolio. This isolates a fund's selection of leaders versus laggards *within* each sector.

In Panel C-1, column (1) shows that financial funds overweight sectors with higher average ESG ratings. Column (4) shows that financial funds also invest in firms with significantly higher ESG scores relative to their sector peers, compared to benchmark funds. Together, these findings confirm that financial funds both tilt toward already "green" sectors and select ESG leaders within those sectors.

In contrast, Panel C-2 column (2) shows that impact funds overweight sectors with lower E ratings, and column (3) shows they tend to allocate more heavily to sectors with higher S ratings, such as utilities. At the same time, column (6) shows that impact funds invest in stocks with significantly lower S ratings than their sector peers. These results suggest that impact funds differ from financial funds both in their sector-level allocations and within-sector selection. Specifically, impact funds appear to target firms and sectors with weaker ESG performance, consistent with an engagement-based strategy focused on driving improvements over time.

Overall, these findings underscore the heterogeneity within sustainable funds. Financial funds, which dominate the sustainable fund universe, systematically invest in firms and sectors with strong ESG profiles—consistent with the objective of enhancing financial returns through ESG integration. In contrast, the smaller subset of impact funds allocates capital to firms and sectors with weaker ESG performance, consistent with a strategy aimed at generating real-world change.

We also decompose the environmental rating, with results reported in the Internet Appendix Table IA4. We find that impact funds hold stocks with significantly higher Pollution & Waste and Environmental Opportunities scores—consistent with investing in companies whose core businesses (e.g., renewable energy, green buildings, and clean tech) contribute directly to environmental solutions. These represent a class of investments where the business model itself serves as a generator of positive environmental impact. However, we also find that impact funds hold companies with significantly lower Natural Capital scores relative to financial funds. These firms are more likely to be associated with operations that generate negative externalities—such as deforestation, water depletion, or unsustainable sourcing—and thus face regulatory, physical, or reputational risks. While MSCI ratings may only imperfectly capture the full extent of these externalities, this finding is consistent with the interpretation that impact funds allocate capital to relatively poor environmental performers, potentially as part of an engagement-based strategy aimed at improvement.

3.3 Carbon Emission Intensity of Fund Portfolios

The findings from the E rating decomposition highlight two distinct modes of sustainable investing: financial funds concentrate in firms with already strong environmental performance, while impact funds pursue a dual strategy of investing in environmental solution providers and firms with room for improvement. However, MSCI ratings capture a broad, relative assessment across industries and may not fully reflect the actual environmental footprint of fund holdings. To complement this analysis and evaluate funds' exposure to environmental externalities more directly, we next examine the carbon intensity of portfolio companies.

Because financial funds prioritize hedging climate-related risks, we hypothesize that financial funds hold portfolios with lower carbon intensity than benchmark funds. In contrast, impact funds aim to generate positive externalities by investing in firms with higher emissions intensity and seeking to reduce their emissions through engagement. Thus, we conjecture that impact funds hold portfolios with higher carbon intensity than financial funds.

[Insert Table 4: Carbon Emission Intensity of Fund Portfolios]

Carbon intensity is measured as Scope 1&2 and Scope 1&2&3 emissions per unit of revenue (metric tons per million dollars of revenue). We also construct within-sector-ranked carbon intensity measures, defined as the decile rank of a firm's emission intensity within its 6-digit NAICS industry. Finally, to isolate sector allocation effects, we calculate sector-average carbon intensity for each fund by assigning the sector's average emissions value to each stock and computing the fund's value-weighted portfolio score.

Panel A of Table 4 reports mean values of these metrics for benchmark, sustainable, financial, and impact funds. Consistent with our hypotheses, we find that financial funds hold portfolios with substantially lower carbon intensity than benchmark funds under both raw and sector-average metrics. This indicates that financial funds select firms with low emissions and tilt toward sectors with lower overall carbon intensity, such as information technology and finance (as shown in Table 2 and Figure 3). Since financial firms constitute the majority of sustainable funds, our results align with Jin et al. (2024), who document that ESG funds disproportionately hold low-emission firms relative to an optimal portfolio benchmark. In contrast, impact funds hold portfolios with higher carbon intensity than both financial and benchmark funds. This is attributable to sector-level allocation. Specifically, impact funds overweight high-intensity sectors such as utilities.

Panel B presents regression results with year–quarter fixed effects. Columns (2) and (4) cluster standard errors at the fund level. Panel B-1 confirms that financial funds hold portfolios with significantly lower carbon intensity than benchmark funds during the same periods. Panel

B-2 and B-3 shows that impact funds hold portfolios with significantly higher carbon intensity than both benchmark funds and financial funds.

Panel C decomposes fund-level carbon intensity into cross-sector and within-sector components, mirroring the decomposition in Table 3 for ESG characteristics. Panel C-1 columns (1) and (2) show that financial funds overweight sectors with lower average emissions, while Panel C-2 columns (1) and (2) show that impact funds overweight higher-emission sectors. Panel C-1 columns (3) and (4) show that financial funds do not hold firms with significantly different emissions intensity than benchmark funds *within* the same sectors. In contrast, Panel C-2 columns (3) and (4) show that impact funds hold firms with significantly lower within-sector emissions intensity than financial funds. This result suggests that impact funds target lower emitters within high-emission industries, focusing on "good students" that show a willingness to reduce emissions in sectors where reductions in emission intensity are consequential and economically meaningful. Taken together, we find that financial funds seek to reduce portfolio exposure to climate-related risks by investing in low-emission firms and sectors, consistent with their financial risk–return objectives. Impact funds, by contrast, invest in high-emission firms and sectors—presumably with the intention of reducing externalities and promoting environmental improvement over time. These portfolio differences reflect the distinct goals and tactics of the two fund types.

4. Sustainable Funds Performance and Actions

4.1 Performance and Flow–Performance Sensitivity by Sustainable Fund Type

This section examines whether mutual funds' stated sustainable investment strategies are associated with differences in financial performance and flow–performance sensitivity. Prior literature finds that socially responsible funds tend to exhibit lower flow–performance sensitivity than conventional funds, consistent with the idea that investors with nonpecuniary preferences are less likely to withdraw capital after underperformance (Bialkowski and Starks (2015); Renneboog, Ter Horst, and Zhang (2011); Bollen (2007)). Motivated by this, we formulate two conjectures. First, moral and impact funds may earn lower alphas than benchmark funds, while financial funds are expected to deliver performance comparable to benchmarks. Second, financial funds should exhibit similar flow–performance sensitivity as benchmark funds, whereas moral and impact funds should exhibit lower sensitivity.

Importantly, both conjectures are joint hypotheses: they require not only that fund managers pursue their stated sustainable investing strategies, but also that investors sort into funds aligned with their preferences. If investors fail to distinguish between fund types—due to search frictions or a lack of transparency—then performance and flow patterns may not vary across categories.

Table 5 presents performance results. Panel A reports time-series regressions of value-weighted net-of-fee monthly excess returns using the CAPM and Fama–French three-factor models. Panel B compares the returns of sustainable fund categories relative to benchmark funds. We find that financial and impact funds exhibit significantly lower loadings on HML, indicating greater tilt toward growth stocks. For financial funds, this tilt may reflect their heavier allocations to the technology sector (see Table 2 and Figure 2). For impact funds, the tilt is likely attributable to investments in firms explicitly targeting environmental or social opportunities. In contrast, moral funds exhibit factor exposures statistically indistinguishable from benchmark funds.

[Insert Table 5: Performance and Factor Tilts]

None of the sustainable fund categories earn alphas significantly different from those of benchmark funds.[26] However, the point estimates are economically meaningful: both moral and impact funds display negative and relatively large alphas, while financial funds' alphas are close to zero. These patterns are broadly consistent with our conjecture that moral and impact strategies may involve costs not borne by financial or benchmark strategies.[27]

Table 6 investigates how flows respond to performance. Following Barber, Huang, and Odean (2016), we construct alphas using an 18-month exponential decay function applied to prior monthly returns, with the CAPM alpha serving as the performance metric given its stronger predictive power for flows. Column (1) compares moral and benchmark funds. We find that moral funds earn significantly higher unconditional flows (roughly 0.3% per month). More notably, moral funds exhibit significantly attenuated flow–performance sensitivity. A one-percentage-point decline in alpha reduces flows by 0.823% for benchmark funds, but by only

[26] Lindsey et al. (2024) and O'Hara and Streltsov (2024) find no significant cost associated with ESG and faith-based values investing, respectively.

[27] Edelen (1999) attributes the commonly observed underperformance of open-end mutual funds to the costs arising from liquidity-driven trading. In our analysis, we assume that the extent of liquidity-motivated trading is uniform across different types of sustainable funds.

0.499% (= 0.823 − 0.324) for moral funds. These findings are consistent with moral funds attracting investors who derive utility from ethical exclusionary screens and thus are less responsive to performance.

[Insert Table 6: Flow Performance Sensitivity]

Column (2) compares impact and benchmark funds. Unlike moral funds, impact funds do not exhibit statistically different flow–performance sensitivity relative to benchmarks. The insignificant difference result is robust when using an alternative three-month decay function and when excluding fixed effects; the results are presented in the Internet Appendix Table IA5.

Overall, we find no evidence that impact funds are systematically matched with less performance-sensitive investors. Several interpretations are plausible. Investors who value the positive externalities generated by impact funds may fail to identify them due to the absence of a transparent labeling or disclosure system in the U.S., which may make investor–fund matching more difficult. This challenge may be exacerbated by the nature of impact strategies themselves: whereas moral funds typically rely on exclusionary screens that are relatively simple and easily identified in prospectus language, impact strategies often involve more complex, multidimensional goals—such as engagement or thematic investing—that are harder to detect and evaluate from standard disclosures. Consequently, nonpecuniary investors with impact-oriented preferences may struggle to identify funds aligned with their goals. Alternatively, nonpecuniary impact investors may be less willing to sacrifice financial returns than morally motivated investors. Finally, impact funds may attract pecuniary investors drawn to the perceived alpha-generating potential of E- or S-themed opportunities.

Column (3) compares financial and benchmark funds. As expected, we find no statistically significant difference in flow–performance sensitivity between the two groups. This is consistent with financial funds competing for flows from traditional pecuniary investors. All three results comparing each of the sustainable funds and benchmark funds are robust to the use of an alternative three-month decay function and to the exclusion of fixed effects. Our findings underscore the importance of distinguishing among sustainable funds by stated objectives. Existing literature typically relies on third-party ESG fund lists or fund name screening, which do not differentiate among financial, moral, and impact goals. Such classifications may mask meaningful heterogeneity. Our methodology reveals that while moral funds exhibit flow dynamics consistent with matching between nonpecuniary investors and strategies aligned with their ethical values, financial and impact funds do not. These results suggest that investor behavior, fund disclosure, and fund strategy all jointly determine sustainable fund outcomes, and that effective investor–fund matching may vary substantially between moral funds and impact funds.

4.2 Proxy Voting on Environmental and Social Shareholder Proposals

We now examine whether funds' voting behavior on shareholder proposals aligns with their stated investment objectives. Proxy voting is a key mechanism through which institutional investors can influence corporate behavior, particularly with respect to environmental and social (ES) issues (Iliev and Lowry (2015), Brav et al. (2024), Di Giuli et al. (2025)).[28] If impact funds are genuinely motivated to generate positive externalities, we would expect them to exhibit higher support for ES-related shareholder proposals, even when such proposals may impose costs on firms. In contrast, financial funds are expected to support ES proposals when they align with long-term value creation or mitigate transition, regulatory, or reputational risks. To the extent that financial funds systematically integrate ESG considerations into valuation, they should support ES proposals more frequently than benchmark funds, but less frequently than impact funds. The voting behavior of moral funds is more ambiguous ex ante, given the diverse ethical concerns and norms that may motivate this group.

Following He, Kahraman, and Lowry (2023), we identify ES-related proposals using a dummy variable and calculate, for each fund-quarter, the percentage of such proposals that a given fund supported. We also examine the incidence of opposing, abstaining from, or not casting a vote on ES proposals, enabling a more granular analysis of voting behavior across fund types.[29]

[Insert Table 7: Fund Votes on Environmental and Social Shareholder Proposals]

Table 7 reveals marked differences in voting behavior across fund categories. Panel A shows that each type of sustainable fund—financial, moral, and impact—supports ES proposals

[28] The literature examines the relative effectiveness of exit versus engagement in achieving social impact. For example, Edmans et al. (2022) study conditional and unconditional exclusion of brown firms. Broccardo et al. (2024) show that exit is ineffective unless a large share of investors act in a socially responsible manner. In this section, we focus on the engagement decision.

[29] We do not exclude sustainable index funds from our sample, as Appel et al. (2016) suggest that passive mutual funds also influence firms' governance. While Heath et al. (2022) find that passive index funds monitor less, passive index sustainable funds are a minority in our sample and inclusion would only bias against finding positive results.

at substantially higher rates than benchmark funds (40.7%–58.5% vs. 25.1%). This pattern suggests that, regardless of specific motivation, sustainable funds exhibit greater engagement on ES issues. Sustainable funds also have significantly lower rates of abstention and non-voting, indicating a more active stance in proxy governance.

Importantly, substantial heterogeneity exists within the sustainable category. Impact funds show the highest support for ES proposals (58.5%), followed by moral funds (47.0%), and financial funds (40.7%). This gradient aligns with our conceptual framework: impact funds have the strongest incentives to use proxy voting to drive improvements in corporate ESG performance, while financial funds are more selective, backing proposals they deem financially material. Moral funds fall in between, reflecting an orientation toward ethical values rather than impact generation or financial performance per se.

While all moral funds by construction have the largest percentage of sustainable sentences classified as moral, some of them also have secondary sustainability intent that may affect their voting behavior. We thus define pure moral funds as those with only moral sentences (thus no financial or impact sentences) and compare them to other moral funds. We find that pure moral funds' support for ES proposals is nearly identical to that of financial funds. In contrast, the higher support for environmental and social (ES) proposals among moral funds is mainly driven by moral funds with also some impact sentences. The results are reported in the Internet Appendix Table IA6.

Panel B presents regression results controlling for year–quarter fixed effects. In Panel B-1, we find that financial funds are 11.8 percentage points more likely to vote in favor of ES proposals than benchmark funds, 9.6 percentage points less likely to vote against, and 2.6 percentage points less likely to abstain. Panel B-2 shows that, relative to other sustainable funds, impact funds are an additional 19.8 percentage points more likely to vote in favor of ES proposals and 20.8 percentage points less likely to vote against. These effects are economically large and statistically significant.

Next, we focus on outcome-focused E&S proposals and examine whether impact funds vote differently on these proposals. If impact funds are truly impact-driven, they should be especially supportive of outcome-focused proposals relative to financial or moral funds. In contrast, for other E&S proposals (e.g., those that advocate for financially material disclosures), the behavioral differences among the three sustainability fund types are a priori less clear.[30] Panel C reports regression results estimated at the fund–proposal level, where the dependent variable is an indicator equal to one if the fund voted Yes on the proposal. We define a proposal as Outcome Proposal if it falls into an ISS category whose proposal descriptions (*AgendaGeneralDesc* and *ItemDesc*) are not related to ESG reporting.[31]

Column (1) includes all sustainable funds, while column (2) compares impact funds with benchmark funds. In both specifications, the interaction between the Impact Fund and Outcome Proposal indicators is positive and statistically significant, indicating that impact funds are more likely to support outcome-focused E&S proposals than disclosure-focused ones. Among all sustainable funds, impact funds are 18.1 percentage points more likely to vote in favor of outcome-focused proposals, increasing to 21.2 percentage points relative to benchmark funds. Outcome-focused proposals receive lower average support, consistent with higher implementation costs. Overall, the results suggest that impact funds disproportionately support proposals aimed at achieving real outcomes despite their lower average support. Taken together, the voting patterns documented in Table 7 provide strong evidence that funds act in ways that reflect their stated objectives. Impact funds' elevated support for ES proposals is consistent with the use of proxy voting as a deliberate tool for engagement and influence. Financial funds, while more supportive than conventional peers, adopt a more measured approach—likely weighing financial materiality. Moral funds exhibit intermediate behavior, consistent with a values-based orientation that varies across ethical dimensions.

These findings complement our earlier results on portfolio composition. Impact funds not only invest in firms with the potential for ESG improvement but also use governance mechanisms to encourage such improvements. The alignment between stated objectives and observed actions—across both asset allocation and voting behavior—offers new insight into how different sustainable investment motivations are operationalized in practice.

[30] We thank Nadya Malenko for suggesting this analysis.

[31] This definition consists of proposals with the following ISSAgendaItem IDs: S0224, S0411, S0416, S0703, S0732, and S0745. These proposals seek to eliminate corporate involvement in harmful products and fossil fuels and to implement policies that promote employment welfare and human rights standards. The detailed descriptions are provided in Internet Appendix Section 5.

5. Emission Outcomes

5.1 E and S Performance Improvements During Holding Periods

5.1.1 ESG Ratings Analysis

A critical test of impact funds' effectiveness is whether they actually succeed in improving the environmental and social performance of their portfolio companies. If impact funds genuinely select firms with poor E and S performance but improvement potential—as suggested by our findings in Sections 3.2 and 3.3—and actively engage through proxy voting (as shown in Section 4.2), we should observe greater improvements in ESG outcomes for companies held by impact funds compared to those held by other fund types.

To evaluate this, we examine changes in two sets of outcome metrics: ESG ratings and carbon emission intensity. For each new stock holding reported by a fund in month t, we track the change in ESG, E, and S ratings (as well as emission intensity) at t+12, t+18, and t+24. Stocks are included only if they remain in the fund portfolio at each respective time point, [32] which introduces survivor bias—funds may sell firms whose E and S performance deteriorates, selectively retaining firms that improve. However, as long as this culling behavior is similar across fund types, observed differences in improvement rates among the remaining holdings can still be informative about fund managers' influence.

[Insert Table 8: Changes in ESG Ratings and Carbon Emission Intensity]

Panel A of Table 8 presents average changes in ESG, E, and S ratings across benchmark, financial, moral, and impact funds. We find that, in general, ESG ratings improve over time among stocks that remain in fund portfolios. Comparing benchmark and financial funds, we do not observe consistent differences in the extent of ESG rating improvements: for example, financial fund holdings show smaller improvements in E ratings but somewhat larger improvements in S ratings relative to benchmark funds. Results for moral funds are similarly mixed.

[32] In cases where a fund's classification changes during a stock's holding period—due to updates in the fund's prospectus language—we treat the investment as exited in the month when the fund is no longer classified under the type it was assigned at the time of the stock's purchase. As a robustness check, we allow the holding period to continue beyond the classification "switch" as long as the fund retains at least one sentence associated with the original category. For example, if a fund initially classified as an impact fund later becomes classified as a moral fund but still contains at least one sentence classified as impact (i.e., becomes a hybrid moral–impact fund), we continue to treat the investment as ongoing until the stock is sold. Our main results are robust to this alternative specification.

By contrast, impact funds exhibit consistently stronger improvements in both E and S ratings relative to benchmark funds. Although ESG ratings may not perfectly capture real-world improvements, this evidence suggests that impact fund investments are associated with positive ESG performance trajectories at the portfolio firm level.

To examine more direct environmental outcomes, we turn next to changes in portfolio companies' carbon emission intensity.

5.1.2 Carbon Emission Improvements During Holding Periods

In this section we examine changes in emission intensity at the investment holdings level. We define investment holding periods analogously to the previous section. The sample sizes are different because MSCI and Trucost coverage of firms differ (Trucost covers more firms than MSCI). We measure both changes in Scope 1 & 2 and Scope 1, 2 & 3 emission intensity from the purchase month t to t+12, t+18 and t+24 for benchmark, financial, moral and impact funds.

Panel B of Table 8 reports the average change in carbon emission intensity during investment holding periods across fund types. On average, we observe a decline in emission intensity for portfolio companies held by all fund categories. This pattern likely reflects both the survivor bias discussed in the prior ESG analysis—where firms that worsen on emissions may be dropped from portfolios—and a broader secular trend of declining corporate emission intensity during the sample period, as firms adopted more carbon-efficient technologies.

Comparing sustainable and benchmark funds, we find that sustainable fund holdings generally exhibit greater reductions in emission intensity. However, the magnitude of emission intensity decline is not significantly different between financial and impact funds. In contrast, moral fund holdings show more modest reductions. These findings suggest that while sustainable funds as a group are associated with greater carbon intensity improvements than benchmark funds, there is heterogeneity in performance across sustainable fund types.

5.2 Emission Intensity Changes: Difference-in-Differences Analysis

A limitation of the previous analysis is that it does not account for heterogeneity in the types of portfolio companies held by different funds. Reducing emission intensity for an airline company, for example, is fundamentally different from doing so for a cement producer. Additionally, overlap in fund ownership may confound attribution: if a firm purchased by a financial fund was also held—currently or previously—by an impact fund, any observed reduction in emissions could be mistakenly attributed to the financial fund. Additionally, we observe that firms held by all sustainable type funds exhibit reductions in emission intensity. This decline in emissions may arise either because sustainable funds select firms that are already on a trajectory of declining emission intensity (a selection effect), or because fund ownership influences firm management and leads to reductions in emissions (a treatment effect). To disentangle these two effects, we (i) implement a stacked difference-in-differences (DiD) framework and (ii) estimate the dynamic effects of sustainable fund ownership. We adopt this specific DiD implementation for two reasons.

First, the recent DiD literature highlights problems that arise when treated firms are used as controls under staggered treatment timing (Callaway and Sant'Anna, 2021; Baker et al., 2022). In our setting, where firms are held by different impact funds at different times, this concern naturally applies. Therefore, for each treated firm, we identify a matched control firm based on *ROA* and *log(total assets)* within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between t–24 and t+24.We stack the treatment firm and matched control firms as our sample of analysis. This way, the treated firm is never used as control firms and alleviate the staggered DID concern. Moreover, as our control firms and treatment firms are from the same sector, this allows us to account for heterogeneity in the types of portfolio companies held by different funds and be able to compare emission within sector.

Second, treated firms may be held by sustainable funds multiple times during the sample period, generating multiple treatment events. To accommodate settings with multiple treatments in close succession, we follow Sandler and Sandler (2014) and allow multiple event-time dummies to be activated simultaneously. This approach allows us to compare treated and control firms prior to treatment and to trace the dynamic evolution of emission reductions, helping to distinguish between selection and treatment effects. Specifically, we estimate a dynamic DiD model as follows:

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

where $Holding_{i,t}^{k}$ is a dummy variable equal to one for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$.[33] We define a fund's holding session as the uninterrupted period during which a stock appears in the fund's portfolio Holding sessions lasting fewer than six months are excluded from the analysis. The outcome variable is Scope 1&2 and Scope 1&2&3 carbon emission intensity. All regressions include firm and year-month fixed effects, and standard errors are clustered at the firm level. We drop the constant term and retain the variable at $Holding_{i,t}^{k}$ at k = 0, so that we can directly compare treated and control firms at the time of treatment.

[Insert Figure 4: DiD Analysis of Firm Emission Intensity during Holding Period]

Figure 4 plots the estimated dynamic coefficients for Scope 1&2 emission intensity. For impact funds (Panel A of Figure 4), we find no evidence of pre-trend violations: treated firms and matched controls follow parallel trajectories prior to investment. Starting around month +10, treated firms exhibit significantly lower emission intensity, with effects lasting approximately six months and remaining negative thereafter. These results are consistent with the hypothesis that impact fund engagement contributes to meaningful reductions in emissions, with a lag that reflects the time required for operational change.

The timing of emission reductions among firms held by impact funds helps disentangle whether these reductions reflect impact funds selecting firms that are already on a trajectory of declining emission intensity (a selection effect) or whether fund ownership influences firm

[33] Suppose firm i is held by two impact funds during the sample period. Fund 1 starts holding firm i in January 2018, and Fund 2 starts holding firm i in January 2019. For all $t \leq$June 2016, all $Holding_{j=i,t}^{k} = 0$ for $k \in [-18,18]$. At $t =$July 2016, $Holding_{i,t}^{-18} = 1$ because Fund 1 begins holding firm i in January 2018, which is 18 months after 2016m7; all other event-time dummies equal zero. As time advances, the corresponding event-time dummy associated with Fund 1 shifts forward. For example, at $t =$June 2017, $Holding_{i,t}^{-7} = 1$ and all other dummies equal zero. At $t =$July 2017, two event-time dummies are activated simultaneously. Specifically, $Holding_{i,t}^{-6} = 1$ because Fund 1 begins holding firm i in 2018m1 (6 months ahead), and $Holding_{i,t}^{-18} = 1$because Fund 2 begins holding firm i in 2019m1 (18 months ahead). Similarly, at $t =$August 2017, $Holding_{i,t}^{-5} = 1$ and $Holding_{i,t}^{-17} = 1$, with all other dummies equal to zero. At $t =$January 2018, $Holding_{i,t}^{0} = 1$ for the start of Fund 1's holding of firm i, while $Holding_{i,t}^{-12} = 1$ reflects start of Fund 2's holding one year later. At $t =$July 2019, $Holding_{i,t}^{+6} = 1$ reflecting post-treatment periods following the start of Fund 2's holding of firm i, and $Holding_{i,t}^{+18} = 1$, reflecting post-treatment periods following the start of Fund 1's holding of firm i. At $t =$August 2019, only $Holding_{i,t}^{+7} = 1$ remains active, since this date lies more than 18 months after Fund 1's start of holding. For all $t \geq$August 2020, all event-time dummies equal zero. If two funds of the same type begin holding firm i at the same time, we still set the dummy equal to one. In the Internet Appendix Figure IA1, we provide robustness results in which, in such cases, we replace the indicator with the total number of simultaneous holding events that share the same entry date.

management and leads to reductions in emissions (a treatment effect). If impact funds endogenously select firms that are already on a path of declining emission intensity, treated firms would exhibit lower emissions than control firms even prior to being held by impact funds, along with a downward pre-trend. Instead, we find that the significant decline in emissions occurs only after firms enter impact fund portfolios, with no evidence of differential pre-trends. These results suggest that being held by impact funds has a treatment effect, in that impact fund ownership reduces portfolio firms' emission intensity.

While we cannot fully rule out the possibility that impact funds happen to invest in firms at the exact moment their emissions begin to decline, this scenario is not likely. It would require funds to accurately anticipate the precise timing of emission reductions, an event typically driven by firm-specific and largely unobservable operational, technological, or regulatory factors. Therefore, it is highly unlikely that the observed reduction is driven by impact funds selecting firms that are just on the verge of declining emission intensity.

Our definition of impact funds does not distinguish between environmentally focused and socially focused funds. To strengthen the analysis, we restrict the sample to firms held by environmentally focused impact funds, defined as those in the top quintile of support for environmental-related shareholder proposals. The results are presented in Panel B of Figure 4. Consistent with Panel A, we find no evidence of pre-trend violations. Beginning around month +7, treated firms exhibit significantly lower emission intensity relative to control firms, with effects persisting for the remaining twelve months during the analysis period. These findings indicate that the reduction in emission intensity is economically significant. The results for Scope 1&2&3 are reported in the Internet Appendix Table IA7. We find robust results, with no evidence of pre-trends, and treated firms exhibiting significantly lower emission intensity after treatment. The stronger result for environmentally focused impact funds provides corroborating evidence that the fund's sustainability intent is associated with real-world outcomes.

For financial funds (Panel C of Figure 4), we observe anomalous pre-trend patterns: treated firms show significantly lower emission intensity than controls beginning as early as –12 and persisting through approximately +8. The coefficients gradually revert to zero after that point. These pre-treatment patterns suggest that observed emission reductions may stem from prior or concurrent ownership by impact funds. Given that impact fund effects appear to begin around +10, lagged influence from earlier impact fund holdings may be confounding the financial fund results.

To test this, we impose a stricter restriction: we exclude any financial fund investment in firms that were held by impact funds between –18 and +18. Results from this cleaner specification (Panel D of Figure 4) show that pre-trend anomalies largely disappear, except for modest significance between –5 and –2. Crucially, we find no treatment effect after financial fund investment. This suggests that financial fund ownership alone is not associated with emission reductions at portfolio firms.

For moral funds (Panel E and Figure 4), we find weak evidence of pre-trends but no treatment effects. This pattern suggests that moral funds may select firms already on an improving trajectory with respect to emission intensity, but their ownership does not appear to induce further improvement.

Taken together, these results corroborate and extend our earlier findings. Impact funds select firms with relatively poor initial ESG performance (Sections 3.2 and 3.3), actively engage through proxy voting (Section 4.2), and achieve measurable improvements in environmental performance during their holding periods (Section 5.1). The consistency of evidence across portfolio composition, engagement, and outcome dimensions strengthens confidence in our classification methodology and demonstrates that impact funds' actions align with their stated objectives.

The ESG improvement analysis also provides important insights into the real-world consequences of different sustainable investing approaches. Financial funds tend to invest in already "green" companies, and moral funds avoid controversial sectors through exclusionary screens. By contrast, impact funds are uniquely positioned to contribute directly to environmental improvements through their ownership and engagement strategies. Despite managing only 2% of sustainable fund assets—and just 0.08% of total mutual fund assets—impact funds may generate disproportionate positive externalities compared to other sustainable investment strategies. This highlights a key tension in sustainable finance: the fastest-growing categories of ESG investment are not necessarily those with the greatest potential to deliver meaningful environmental and social change. Policymakers, fund managers, and investors must grapple with this disconnect in order to better align capital flows with sustainability outcomes.

6. Conclusion

This paper studies how mutual funds articulate sustainability goals and whether those stated intentions are informative about investment behavior and real-economy outcomes. Using a text-based classification of fund disclosures, we distinguish between financial, moral, and impact-oriented sustainability intent. The classification is validated at the sentence level and serves as an organizing variable for the empirical analysis.

We document that most funds using ESG language articulate sustainability primarily in financial terms, while relatively few express a distinct impact-oriented mandate. This descriptive pattern highlights the importance of distinguishing among different sustainability objectives that are often grouped together under a single ESG label. We show that sustainability intent predicts systematic differences in portfolio composition, investor behavior, and stewardship. In particular, moral funds exhibit lower flow-performance sensitivity, while impact-oriented funds are more likely to support outcome-oriented environmental and social shareholder proposals.

Turning to real-economy outcomes, we find that firms added to the portfolios of impact-oriented funds subsequently reduce carbon intensity relative to matched firms, while no comparable pattern is observed for firms added by financial or moral funds. These results suggest that distinctions in stated sustainability intent are informative about downstream firm behavior. At the same time, the absence of post-investment emissions changes for financial and moral funds underscores the need to separate sustainability intent from broader ESG labeling when evaluating claims about real-world impact.

Our findings also help clarify how sustainability-oriented funds are interpreted in public and academic discussions. Using fund disclosures, we find that the majority of assets in sustainability-labeled mutual funds—including those managed by large asset managers—are allocated to funds with financially motivated sustainability objectives rather than impact-oriented mandates. These funds emphasize the use of ESG information for risk management and valuation, rather than engagement aimed at generating measurable negative externalities reductions. Consistent with this distinction, we find limited evidence that financially motivated funds pursue outcome-oriented strategies at scale. Taken together, these patterns underscore the importance of transparent frameworks for distinguishing among different sustainability objectives.

Several directions for future research emerge from our findings. An open question is why assets are so heavily concentrated in financially motivated sustainability funds, and whether this reflects investor preferences, regulatory constraints, or career incentives faced by fund managers. Future work could also examine how disclosure regimes, such as the EU's Sustainable Finance Disclosure Regulation, affect fund strategies and capital allocation, and whether similar patterns arise in other institutional settings, including pensions, credit markets, and private equity.

References

Abis, Simona, 2022, “Man vs. Machine: Quantitative and Discretionary Equity Management”, unpublished working paper.

Abis, Simona, Andrea M. Buffa, and Meha Sadasivam, 2024. “Different Shades of ESG Funds”, unpublished working paper.

Abis, Simona, and Anton Lines, 2024, “Broken Promises, Competition, and Capital Allocation in the Mutual Fund Industry”, Journal of Financial Economics, 162, 103948.

Akey, P. and Appel, I., 2020. Environmental Externalities of Activism. Unpublished working paper.

Allcott, H., Montanari, G., Ozaltun, B. and Tan, B., 2023. “An economic view of corporate social impact”, forthcoming in the Journal of Finance.

Appel, Ian R., Todd A. Gormley, and Donald B. Keim, 2016. Passive investors, not passive owners. *Journal of Financial Economics*, *121*(1), pp.111-141.

Aswani, Jitendra, Aneesh Raghunandan, and Shiva Rajgopal, 2024. Are carbon emissions associated with stock returns?. *Review of Finance*, *28*(1), pp.75-106.

Atta-Darkua, V., Glossner, S., Krueger, P. and Matos, P., “Decarbonizing Institutional Investor Portfolios: Helping to Green the Planet or Just Greening Your Portfolio?”, unpublished working paper.

Avramov, D., Cheng, S., Lioui, A. and Tarelli, A., 2022. Sustainable investing with ESG rating uncertainty. *Journal of financial economics*, *145*(2), pp.642-664.

Baker, A.C., Larcker, D.F., McClure, C.G., Saraph, D. and Watts, E.M., 2024. Diversity washing. *Journal of Accounting Research*, *62*(5), pp.1661-1709.

Baker, A.C., Larcker, D.F. and Wang, C.C., 2022. How much should we trust staggered difference-in-differences estimates?. *Journal of Financial Economics*, *144*(2), 370-395.

Baker, Malcolm, Mark L. Egan, and Suproteem K. Sarkar (2022) *How Do Investors Value ESG?*. No. w30708. National Bureau of Economic Research.

Barber, B. M., X. Huang, and T. Odean, 2016. Which factors matter to investors? Evidence from mutual fund flows. *The Review of Financial Studies*, *29*(10), pp.2600-2642.

Barber, Brad, Adair Morse, and Ayako Yasuda (2021) “Impact Investing” Journal of Financial Economics 139, 162-185

Bauer, R., Ruof, T., Smeets, P., 2021. Get Real! Individuals Prefer More Sustainable Investments. *The Review of Financial Studies*, 34, Issue 8, p. 3976–4043.

Benson, K.L. and Humphrey, J.E., 2008. Socially responsible investment funds: Investor reaction to current and past returns. *Journal of Banking & Finance*, *32*(9), pp.1850-1859.

Berg, F., Kölbel, J.F. and Rigobon, R., 2022. Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, *26*(6), pp.1315-1344.

Berk, Jonathan, and Jules H. van Binsbergen. (2025) "The impact of impact investing." Journal of Financial Economics 164, 103972.

Białkowski, Jedrzej, and Laura T. Starks, 2015. SRI funds: Investor demand, exogenous shocks and ESG profiles. Working Paper.

Bingler, Julia and Kraus, Mathias and Leippold, Markus and Webersinke, Nicolas (2024) "How Cheap Talk in Climate Disclosures relates to Climate Initiatives, Corporate Emissions, and Reputation Risk", *Journal of Banking & Finance*, *164*, 107191.

Bollen, Nicolas PB., 2007. Mutual fund attributes and investor behavior. *Journal of financial and quantitative analysis*, *42*(3), pp.683-708.

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

Bolton, Patrick, and Marcin Kacperczyk, 2023. Global pricing of carbon-transition risk. *The Journal of Finance*, *78*(6), pp.3677-3754.

Bonnefon, J-F, A. Landier, P. Sastry, and D. Thesmar, (2025) "The moral preferences of investors: Experimental evidence", *Journal of Financial Economics* 163 103955.

Brav, A., W. Jiang, T. Li, and J. Pinnington, (2024). Shareholder monitoring through voting: New evidence from proxy contests. *The Review of Financial Studies*, *37*(2), 591-638.

Briscoe-Tran, H., 2025. Do employees have useful information about firms' Esg practices?. *Fisher College of Business Working Paper*, (2021-03), p.21.

Broccardo, Eleonora, Oliver Hart, and Luigi Zingales, (2022). Exit versus voice. *Journal of Political Economy*, *130*(12), pp.3101-3145.

Callaway, B. and Sant'Anna, P.H., 2021. Difference-in-differences with multiple time periods. *Journal of econometrics*, *225*(2), pp.200-230.

Cao, Sean Shun, Wei Jiang, and Lijun Gillian Lei, 2024. Applied AI for finance and accounting: Alternative data and opportunities. *Pacific-Basin Finance Journal*, *84*, p.102307.

Chen, T., Dong, H. and Lin, C., 2020. Institutional shareholders and corporate social responsibility. *Journal of Financial Economics*, *135*(2), pp.483-504.

Chowdhry, Bhagwan, Shaun William Davies, and Brian Waters (2019) "Investing for Impact" The Review of Financial Studies 32: 864–904.

Christensen, D.M., Serafeim, G. and Sikochi, A., 2022. Why is corporate virtue in the eye of the beholder? The case of ESG ratings. *The Accounting Review*, *97*(1), pp.147-175.

Cohen, L., Gurun, U.G. and Nguyen, Q.H., 2020. *The ESG-innovation disconnect: Evidence from green patenting* (No. w27990). National Bureau of Economic Research.

Couvert, Maxime, 2025, "What is the Impact of Mutual Funds' ESG Preferences on Portfolio Firms?", unpublished working paper.

Dangl, Thomas, Michael Halling, Jin Yu, and Josef Zechner, "Social preferences and corporate investment", forthcoming in Journal or Financial Studies.

Dasgupta, Amil, Jenter, Dirk, Mathews, Richmond D. and Voss, Paul, 2025, "Are Socially Responsible Funds Viable?", unpublished working paper.

Diaz-Rainey, I., Griffin, P.A., Lont, D.H. Matero-Marquez, A., Zamora-Ramirez, C. (2023) Shareholder Activism on Climate Change: Evolution, Determinants, and Consequences. *J Bus Ethics*, 1-30.

Di Giuli, Alberta, Alexandre Garel, Roni Michaely, and Arthur Romec, 2025. Climate change and mutual fund voting on climate proposals. *Management Science*.

Dimson, E., Karakaş, O. and Li, X., 2015. Active ownership. *The Review of Financial Studies*, *28*(12), pp.3225-3268.

Dimson, E., Karakaş, O. and Li, X., 2025. Coordinated engagements. *European Corporate Governance Institute–Finance Working Paper*, *721*(6).

Duchin, R., J. Gao, and Q. Xu, (2025) "Sustainability or Greenwashing: Evidence from the Asset Market for Industrial Pollution", *The Journal of Finance* 80, no. 2 (2025): 699-754.

Edelen, Roger M, 1999. Investor flows and the assessed performance of open-end mutual funds. *Journal of Financial Economics*, *53*(3), pp.439-466.

Edmans, Alex, Tom Gosling, and Dirk Jenter, 2024. Sustainable investing: Evidence from the field. *FEB-RN Research Paper*, (18).

Edmans, Alex, Doron Levit, and Jan Schneemeier, 2022. Socially responsible divestment. *Centre for Economic Policy Research Working Paper.*

El Ghoul, S. and Karoui, A., 2017. Does corporate social responsibility affect mutual fund performance and flows?. *Journal of Banking & Finance*, *77*, pp.53-63.

Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel, 2020. Hedging climate change news. *The Review of Financial Studies*, *33*(3), pp.1184-1216.

Flammer, C., Toffel, M. and Viswanathan, K., 2021. Shareholder activism and firms' voluntary disclosure of climate change risks. *Strategic Management Journal*, *42*, 1850-1879.

Gantchev, N., Giannetti, M. and Li, R., 2022. Does money talk? Divestitures and corporate environmental and social policies. *Review of Finance*, *26*(6), pp.1469-1508.

Gantchev, N., Giannetti, M. and Li, R., 2024. Sustainability or performance? Ratings and fund managers' incentives. *Journal of Financial Economics*, 155, p.103831.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. (2019) "Text as data." Journal of Economic Literature 57, 535-574.

Gibson Brandon, R., Glossner, S., Krueger, P., Matos, P. and Steffen, T., 2022. Do responsible investors invest responsibly?. *Review of Finance*, *26*(6), pp.1389-1432.

Giglio, Stefano, Bryan Kelly, and Johannes Stroebel, 2021. Climate finance. *Annual review of financial economics*, *13*(1), pp.15-36.

Giglio, S., Matteo M., Johannes Stroebel, Zhenhao Tan, Stephen Utkus, and Xiao Xu, 2025, Four facts about ESG beliefs and investor portfolios, *Journal of Financial Economics* 164 (2025): 103984.

Goldstein, Itay, Chester Spatt, Mao Ye. (2021) "Big Data in Finance" The Review of Financial Studies 34, 3213–3225.

Goldstein, Itay, Alexandr Kopytov, Lin Shen and Haotian Xiang (2022) "On ESG Investing: Heterogeneous Preferences, Information, and Asset Prices" Unpublished working paper

Green, D., and B. Roth. (2025) "The allocation of socially responsible capital." *The Journal of Finance* 80, no. 2 (2025): 755-781.

Gupta, D., Kopytov, A., & Starmans, J. (2026). The pace of change: Socially responsible investing in private markets. *The Review of Financial Studies*, *39*(1), 30-78.

Hartzmark, S. M. and Sussman, A. B. (2019): Do investors value sustainability? A natural experiment examining ranking and fund flows, Journal of Finance 74, 2789–2837.

Hartzmark, Samuel M., and Kelly Shue. (2023) Counterproductive sustainable investing: The impact elasticity of brown and green firms. Unpublished working Paper.

He, Haibo, and Edwardo A. Garcia. (2009) "Learning from imbalanced data." *IEEE Transactions on knowledge and data engineering* 21, no. 9: 1263-1284.

He, Yazhou Ellen, Bige Kahraman, and Michelle Lowry, (2023) "ES Risks and Shareholder Voice", *The Review of Financial Studies* 36, 4824-4863.

Heath, D., Macciocchi, D., Michaely, R. and C. Ringgenberg, M., 2023. Does socially responsible investing change firm behavior?. *Review of Finance*, *27*(6), pp.2057-2083.

Heeb, Florian, Julian F. Kölbel, Falko Paetzold and Stefan Zeisberger (2023) "Do Investors Care about Impact?" Review of Financial Studies 36, 1737-1787.

Heeb, Florian, Julian Kölbel, and Camilla Weder, (2025) "Beliefs about the Climate Impact of Green Investing", unpublished working paper.

Hoepner, A.G., Oikonomou, I., Sautner, Z., Starks, L.T. and Zhou, X.Y., 2024. ESG shareholder engagement and downside risk. *Review of Finance*, *28*(2), pp.483-510.

Huang, A.H., Wang, H. and Yang, Y., 2023. FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, *40*(2), pp.806-841.

Iliev, Peter, and Michelle Lowry, (2015). Are mutual funds active voters?. *The Review of Financial Studies*, *28*(2), 446-485.

Jiang, Wei, Meng Wang, and Baozhong Yang, 2024. Measuring Sustainability with AI. In *Artificial Intelligence, Finance, and Sustainability: Economic, Ecological, and Ethical Implications* (pp. 33-57). Cham: Springer Nature Switzerland.

Jin, Dunhong, Roni Michaely, and Menghan Wang, 2024. How green is green? Anatomy of ESG funds' selection. Working Paper.

Kim, S. and Yoon, A., 2023. Analyzing active fund managers' commitment to ESG: Evidence from the United Nations Principles for Responsible Investment. *Management science*, *69*(2), pp.741-758.

Krueger, Philipp, Zacharias Sautner, and Laura T. Starks, 2020, The importance of climate risk for institutional investors, *Review of Financial Studies* 33, 1067–1111.

Landier, Augustin and Lovo, Stefano (2025) "Socially Responsible Finance: How to Optimize Impact" *The Review of Financial Studies* 38, no. 4 (2025): 1211-1258.

Lee, Matthew, Arzi Adbi, and Jasjit Singh (2020) "Categorical cognition and outcome efficiency in impact investing decisions" Strategic Management Journal 41:86–107

Lewellen, J. and Lewellen, K., 2022. Institutional investors and corporate governance: The incentive to be engaged. *The Journal of Finance*, *77*(1), pp.213-264.

Li, J. and Wu, D., 2020. Do corporate social responsibility engagements lead to real environmental, social, and governance impact?. *Management Science*, *66*(6), 2564-2588.

Liang, H., Sun, L. and Teo, M., 2022. Responsible hedge funds. *Review of Finance*, *26*(6), pp.1585-1633.

Lindsey, L.A., Pruitt, S. and Schiller, C., 2024. The cost of ESG investing. *Unpublished working paper.*

Lowry, M., Wang, P. and Wei, K.D., 2025. Are all ESG funds created equal? Only some funds are committed, forthcoming in the Review of Financial Studes.

Michaely, R., Ordonez-Calafi, G. and Rubio, S., 2024. Mutual funds' strategic voting on environmental and social issues. *Review of Finance*, *28*(5), pp.1575-1610.

Michaely, R., S. Rubio, and I. Yi (2023) "Voting Rationales", Unpublished working paper

Morgan, John and Justin Tumlinson (2019) "Corporate Provision of Public Goods" Management Science 65, 4489-4504.

Naaraayanan, S.L., Sachdeva, K. and Sharma, V., 2021. The real effects of environmental activist investing. *European Corporate Governance Institute–Finance Working Paper*.

Oehmke, Martin and Opp, Marcus M. (2025) "A Theory of Socially Responsible Investment" Review of Economic Studies 92, 1193-1225.

O'Hara, M. and Streltsov, A., 2024. Keeping the Faith (and the Returns): An AI Approach to Values-based Investing. *Unpublished Working Paper*.

Parise, Gianpaolo, and Mirco Rubin, 2025. Green window dressing. forthcoming in the *Journal of Finance.*

Pástor, Lubos, Robert Stambaugh, and Lucian Taylor (2021) "Sustainable Investing in Equilibrium" Journal of Financial Economics, 142(2), 550-571

Pástor, Lubos, Robert Stambaugh, and Lucian Taylor, 2022. Dissecting green returns. *Journal of financial economics*, *146*(2), pp.403-424.

Pástor, Lubos, Robert Stambaugh, and Lucian Taylor, 2023. *Green tilts* (No. w31320). National Bureau of Economic Research.

Pástor, Lubos, Robert Stambaugh, and Lucian Taylor, 2024. *Carbon burden* (No. w33110). National Bureau of Economic Research.

Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski (2021) "Responsible Investing: The ESG-efficient Frontier" Journal of Financial Economics 142(2) 572–597

Raghunandan, A. and Rajgopal, S., 2022. Do ESG funds make stakeholder-friendly investments?. *Review of Accounting Studies*, *27*(3), pp.822-863.

Rajan, R. and Ramella, Pietro and Zingales, Luigi (2022) "What Purpose Do Corporations Purport? Evidence from Letters to Shareholders", Unpublished working paper.

Rajgopal, Shiva, Anup Srivastava, and Rong Zhao, 2025. Economic substance behind Texas political anti-ESG sanctions. *Management Science*.

Renneboog, Luc, Jenke Ter Horst, and Chendi Zhang., 2011. Is ethical money financially smart? Nonfinancial attributes and money flows of socially responsible investment funds. *Journal of Financial Intermediation*, *20*(4), pp.562-588.

Riedl, A. and Smeets, P. (2017), Why Do Investors Hold Socially Responsible Mutual Funds?. The Journal of Finance, 72: 2505-2550.

Sandler, Danielle H., and Ryan Sandler, 2014. Multiple event studies in public finance and labor economics: A simulation study with applications. *Journal of Economic and Social Measurement*, *39*(1-2), pp.31-57.

Sautner, Z., Van Lent, L., Vilkov, G. and Zhang, R., 2023. Firm-level climate change exposure. *The Journal of Finance*, *78*(3), pp.1449-1498.

Shanker, Harshini (2025) "Ethical Capital, Coordination, and the Correction of Externalities", unpublished working paper.

Starks, Laura (2023) "Presidential Address: Sustainable Finance and ESG Issues—*Value* versus *Values"* Journal of Finance 78, 1837-1872.

Sulaeman, Johan, and Qiaozhi Ye. (2023) "Who Do You Vote For? Same-Race Voting Preferences in Director Elections." Unpublished Working Paper.

U.S. Securities and Exchange Commission. (2016) "How to Read a Mutual Fund Prospectus" from the SEC website.

Appendix

A1: Keywords used to identify ESG-Sentences.
'esg','environment','social','governance','sustainable','sustainability','abortion','lgbt','gay','lesbian','tobacco','gambling','alcohol','pornography','gun','energy','fossil','fuel','green','impact','responsible','clean','minority','minorities','poverty','girl','girls','male','female','fair','maternity','paternity','equal','equality','discrmination','non-discrimination', 'sexual', 'harassment', 'safety', 'diversity', 'civic', 'trafficking','ethics','gender','race','ethnicity', 'climate', 'renewable', 'energy', 'vote', 'voting','proxy', 'transform', 'transformation', 'dialogue', 'engage', 'engagement', 'transition'

A2: Variable Definitions

Variable Name	Definition
Financial$_{it}$	equal to 1 if the most recent available prospectus for fund i as of quarter t contains the highest percentage of Financial sentences among all BERT-classified sustainable sentences (Financial, Moral, and Impact).
Moral$_{it}$	equal to 1 if the most recent available prospectus for fund i as of quarter t contains the highest percentage of Moral sentences among all BERT-classified sustainable sentences (Financial, Moral, and Impact).
Impact$_{it}$	equal to 1 if the most recent available prospectus for fund i as of quarter t contains the highest percentage of Impact sentences among all BERT-classified sustainable sentences (Financial, Moral, and Impact).
Pure Financial$_{it}$	equal to 1 if the most recent available prospectus for fund *i* contains at least one sentence classified as Financial, and no sentences classified as Moral or Impact.
Pure Moral$_{it}$	equal to 1 if the most recent available prospectus for fund *i* contains at least one sentence classified as Moral, and no sentences classified as Financial or Impact
Pure Impact$_{it}$	equal to 1 if the most recent available prospectus for fund *i* contains at least one sentence classified as Impact, and no sentences classified as Financial or Moral
Environmental Impact$_{it}$	equal to 1 if the fund is an impact fund and is in the top quintile based on its support for environmental-related shareholder proposals.
Holding-quarters$_{it}$	the average number of quarters for which the stocks in the portfolio have been held by fund *i* as of quarter *t*.
Retail Expense Ratio$_{it}$	the average expense ratio for retail funds classes for fund *i* as of quarter *t*.
ESG Score$_{jt}$	the average of the evaluated MSCI Environmental and Social key issues scores that firm *j* received from MSCI in quarter *t*
E Score$_{jt}$	the average of the evaluated MSCI Environmental key issues scores that firm *j* received from MSCI in quarter *t*
S Score$_{jt}$	the average of the evaluated MSCI Social key issues scores that firm *j* received from MSCI in quarter *t*
Sector-Average-ESG Score$_{it}$	assign the quarter-sector average *ESG Score*$_{jt}$ at the 6-digit NAICS level to each stock. $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(ESG\ Score_{kt} * MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$
Sector-Average-E Score$_{jt}$	assign the quarter-sector average *E Score*$_{jt}$ at the 6-digit NAICS level to each stock. $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(E\ Score_{kt} * MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$

Sector-Average-S Score$_{jt}$	assign the quarter-sector average *S Score$_{jt}$* at the 6-digit NAICS level to each stock. $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(S\ Score_{kt}*MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$
Sector-Adjusted-ESG Score$_{jt}$	*ESG Score$_{jt}$* minus *Sector-Average-ESG Score$_{it}$*.
Sector-Adjusted-E Score$_{jt}$	E Score$_{\text{jt}}$ minus *Sector-Average-E Score$_{jt}$*.
Sector-Adjusted-S Score$_{jt}$	*S Score$_{jt}$* minus *Sector-Average-S Score$_{jt}$*.
Scope 1&2(&3) Carbon Emission Intensity$_{jt}$	the Scope 1&2(&3) emissions, measured in metric tons per million dollars of revenue, that firm *j* reports at the fiscal year-end immediately following quarter t
Sector-Average Scope 1&2(&3)$_{jt}$	assigning the quarterly sector-level Scope 1&2(&3) emission intensity at the 6-digit NAICS level to each stock $\frac{\sum_{for\ all\ k\ in\ the\ industry\ of\ firm\ j}(Emission*MarketValue_{kt})}{\sum_{\text{for all k in the industry of firm j}}(\text{MarketValue}_{\text{kt}})}$
Within-Sector-ranked Scope 1&2(&3)$_{jt}$	the decile rank of firm *j*'s Scope 1&2(&3) emission intensity within the 6-digit NAICS code.
12 (18, or 24) -Month ESG (E, or S) Score Change	Change in MSCI ESG (E, or S) ratings of firm *j* held by fund *i* between the time of initial investment and 12 (18, or 24) months post-investment
12 (18, or 24) -Month Scope 1&2 Intensity Change	Change in Scope 1&2(&3) Intensity of firm *j* held by fund *i* between the time of initial investment and 12 (18, or 24) months post-investment. A firm's carbon emission at year-month *t* is calculated as the weighted average of the last available and the next available (annually measured) carbon intensity. The weight on the last available carbon intensity equals 12 minus the difference between the current month and the reporting month divided by 12, with the remaining weight applied to the next available carbon intensity.
ES "For" Vote$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* voted for in quarter *t*
ES "Against" Vote$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* voted against in quarter *t*
ES "Abstain" Vote$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* abstained from voting in quarter *t*
ES "Do Not Vote"$_{it}$	the percentage of Environmental and Social shareholder proposals received by firms held by fund *i* that fund *i* did not cast a vote in quarter *t*
Outcome Proposal	equal to 1 if the proposal is in an ISS category whose proposal descriptions by ISS (AgendaGeneralDesc and ItemDesc) are not related to ESG reporting, consisting of ISSAgendaItemIDs S0224, S0411, S0416, S0703, S0732, and S0745.

Table 1: BERT Model Performance

This table reports the performance of the BERT model trained to classify ESG-related sentences. 2,950 sentences in the training sample and 625 in the testing sample are manually labeled into four different categories: Financial, Moral, Impact, or None. The model is trained on the training sample for classification. The trained BERT model is applied to the testing sample to obtain the ESG classification given by the BERT model. Four different model performance measures are calculated to measure the accuracy of BERT classification. Accuracy is the ratio of (true positives + true negatives) divided by the total number of observations (fraction of correct classifications). Precision is the ratio of true positives divided by the sum of true positives and false positives. Recall is the ratio of true positives divided by the sum of true positives and false negatives. *f1* is defined as $[\frac{2}{\frac{1}{Precision}+\frac{1}{Recall}}]$.

	Accuracy	Precision	Recall	f1
Financial	0.9088	0.8467	0.7888	0.8167
Moral	0.9840	0.8500	0.8947	0.8718
Impact	0.9440	0.8276	0.6575	0.7328

Table 2: Summary Statistics

This table reports summary statistics for three categories of sustainable funds from 2014 Q2 to 2023 Q4. Panel A presents general fund characteristics. The number of funds equals the number of fund-quarter observations divided by 39 (the number of quarters in the sample period). % of equity in the portoflio equals the value of equity holdings divided by the fund's total NAV as of the quarter end. The number of stocks equals the number of common stocks held in the fund portfolio. The fund size is an average among all fund-quarters. The number of funds with retail class is the number of fund-quarter observations for funds with retail class offering(s) divided by 39. The retail expense ratio equals the expense ratio for retail funds. Panel B presents industry weights for the 20 sectors included in the NAICS, defined using quarter-end holding values of sector stocks divided by all equity holdings. Benchmark funds are CRSP equity non-index funds or ETFs whose prospectuses contain no sentences with sustainability keywords and are never on the Morningstar sustainable fund list during the sample period. Sector weights among the listed stocks are calculated using year-end market capitalizations. The industry average MSCI E and S ratings, as well as Scope 1&2 and Scope 1&2&3 emission intensities are also reported. A fund is assigned the category with the largest % share of all sustainable sentences.

Panel A: General Characteristics

	Benchmark	Financial	Moral	Impact
# of funds (per quarter)	1524	314	90	46
% of equity in the portfolio	86.0%	88.6%	87.7%	89.8%
# of stocks in the portfolio	138.1	99.3	183.9	90.1
Fund size ($M)	$2,985	$1,226	$769	$327
# of funds with retail class	992	193	51	25
Retail Expense Ratio	1.35%	1.34%	1.27%	1.47%

Panel B: Sector Weights

	Listed Stocks	Benchmark	Financial	Moral	Impact	E ratings	S Ratings	Scope 1&2	Scope 1&2&3
Agriculture, Forestry, Fishing and Hunting	0.1%	0.2%	0.1%	0.2%	0.1%	4.85	5.18	615	973
Mining, Quarrying, and Oil and Gas Extraction	2.4%	5.6%	4.6%	3.8%	3.7%	3.91	5.43	611	793
Utilities	3.3%	2.5%	2.6%	2.4%	8.1%	6.04	6.72	1879	2082
Construction	0.5%	1.0%	0.7%	0.9%	1.5%	5.29	4.56	57	224
Manufacturing	40.3%	34.2%	33.0%	38.1%	44.7%	6.87	5.56	145	352
Wholesale Trade	2.3%	2.5%	2.1%	2.9%	3.1%	6.68	5.53	116	280
Retail Trade	8.4%	5.8%	5.4%	5.7%	3.4%	6.86	4.51	49	129
Transportation and Warehousing	2.3%	3.4%	3.0%	2.8%	2.1%	6.33	4.82	823	945
Information	14.8%	13.1%	14.0%	12.1%	9.2%	8.04	5.19	39	90
Finance and Insurance	15.4%	15.4%	15.0%	14.4%	10.4%	7.75	5.01	35	74
Real Estate and Rental and Leasing	0.6%	4.2%	5.0%	4.8%	2.6%	6.45	5.59	92	138
Professional, Scientific, and Technical Services	3.6%	5.0%	5.2%	5.1%	4.6%	7.12	5.08	31	91
Management of Companies and Enterprises	0.3%	0.7%	1.6%	0.9%	1.3%	6.82	4.91	116	172
Administrative and Support and Waste Management and Remediation Services	2.9%	3.0%	4.1%	3.2%	3.2%	7.44	5.44	98	177
Educational Services	0.0%	0.2%	0.4%	0.2%	0.2%	5.71	3.38	62	139
Health Care and Social Assistance	0.6%	0.9%	0.8%	0.7%	0.7%	6.21	4.80	39	107
Arts, Entertainment, and Recreation	0.4%	0.0%	0.0%	0.0%	0.0%	6.09	4.26	52	118
Accommodation and Food Services	1.7%	0.4%	0.3%	0.3%	0.1%	6.69	4.77	129	240
Other Services (except Public Administration)	0.1%	1.8%	2.0%	1.4%	0.9%	5.77	4.51	152	221

Table 3: ESG Characteristics of Fund Holdings

This table reports ESG ratings of portfolio holdings by fund type. ESG Score is the average of the evaluated MSCI Environmental and Social Key Issues scores. E Score is the average MSCI Environmental Key Issues scores. S Score is the average MSCI Social Key Issues scores. For each fund-quarter, we calculate the value-weighted ESG ratings for all MSCI-rated stock holdings. In Panel A, the mean ESG, E and S scores are provided for the benchmark funds, Sustainable funds, Financial funds, and Impact funds. In Panel B, the regression results are reported for regressing the fund-quarter average ESG, E and S ratings on the financial fund dummy and impact dummy. Panel B-1's sample includes Financial Funds and Benchmark Funds. Panel B-2's sample includes Financial Funds and Impact Funds. *t*-stats are reported in parentheses. In Panel C, columns (1), (2), and (3), the fund's sector-average ESG, E, and S ratings are regressed on the fund indicator variables. To construct the fund's sector-average ratings, we assign the quarter-sector average rating at the 6-digit NAICS level to each stock and calculate the value-weighted fund portfolio rating. In columns (4), (5), and (6), the sector-adjusted ESG, E, and S ratings are regressed on the fund indicator variables. To construct the sector-adjusted ratings, we subtract the quarter-sector average ratings at the 6-digit NAICS level from the raw ratings. Panel C-1's sample includes Financial Funds and Benchmark Funds. Panel C-2's sample includes Impact Funds and Financial Funds.

Panel A: Average Ratings

	(1)	(2)	(3)	(4)
	Benchmark	Sustainable	Financial	Impact
ESG Score	5.716	6.056	6.099	5.988
E Score	6.370	6.729	6.784	6.583
S Score	5.059	5.381	5.412	5.392
Observations	59423	16463	12248	1790

Panel B: Regression Analysis

Panel B-1: Financial Funds vs Benchmark Funds

	(1)	(2)	(3)	(4)	(5)	(6)
	ESG Score	ESG Score	E Score	E Score	S Score	S Score
Financial Fund	0.160***	0.160***	0.173***	0.173***	0.146***	0.146***
	(20.702)	(5.139)	(15.900)	(3.978)	(22.706)	(5.735)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes	No	Yes
Observations	71658	71658	71658	71658	71671	71671
Adjusted R2	0.159	0.159	0.118	0.118	0.196	0.196

Panel B-2: Impact Funds vs Financial Funds

	(1)	(2)	(3)	(4)	(5)	(6)
	ESG Score	ESG Score	E Score	E Score	S Score	S Score
Impact Fund	-0.011	-0.011	-0.104***	-0.104	0.085***	0.085
	(-0.596)	(-0.154)	(-4.069)	(-1.016)	(4.971)	(1.284)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes	No	Yes
Observations	13808	13808	13808	13808	13812	13812
Adjusted R2	0.134	0.134	0.073	0.073	0.182	0.182

Panel C: Within-Sector and Cross-Sector Analysis

Panel C-1: Within-Sector and Cross-Sector Financial Funds vs Benchmark Funds

	(1) Sector-Average ESG Score	(2) Sector-Average E Score	(3) Sector-Average S Score	(4) Sector-Adjusted ESG Score	(5) Sector-Adjusted E Score	(6) Sector-Adjusted S Score
Financial Fund	0.058***	0.087***	0.028**	0.102***	0.086***	0.118***
	(3.659)	(3.172)	(2.157)	(4.969)	(3.331)	(6.500)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes	Yes	Yes
Observations	71671	71671	71671	71658	71658	71671
Adjusted R2	0.330	0.200	0.364	0.017	0.010	0.027

Panel C-2: Within-Sector and Cross-Sector Impact Funds vs Financial Funds

	(1) Sector-Average ESG Score	(2) Sector-Average E Score	(3) Sector-Average S Score	(4) Sector-Adjusted ESG Score	(5) Sector-Adjusted E Score	(6) Sector-Adjusted S Score
Impact Fund	0.081**	-0.115*	0.277***	-0.093*	0.009	-0.193***
	(2.143)	(-1.675)	(7.489)	(-1.859)	(0.130)	(-3.183)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13812	13812	13812	13808	13808	13812
Adjusted R2	0.324	0.154	0.358	0.008	0.003	0.028

Table 4: Carbon Emission Intensity of Fund Portfolios

This table reports the carbon emission intensity of the funds' portfolio companies. Carbon Intensity is defined as Scope 1&2, and Scope 1&2&3 Carbon Emission Intensity, measured as metric tons of emissions per million dollars of revenue. A fund's sector-average Carbon Emission Intensity is measured by assigning the quarterly sector-level emission intensity at the 6-digit NAICS level to each stock and then calculating the value-weighted carbon emission intensity. Within-sector-ranked Carbon Emission Intensity is the decile rank of Carbon Emission Intensity within the 6-digit NAICS code. Panel A reports the mean Scope 1&2 Carbon Emission Intensity, Scope 1&2&3 Carbon Emission Intensity, Within-sector-ranked Carbon Emission Intensity, and sector-average Carbon Emission Intensity for Benchmark Funds, Sustainable Funds, Financial Funds, and Impact Funds. In Panel B, regression results are reported for regressing Scope 1&2, and Scope 1&2&3 Carbon Emission Intensity on the Financial Fund dummy and Impact Fund dummy. Panel B-1's sample includes Financial Funds and Benchmark Funds, Panel B-2's sample includes Impact Funds and Benchmark Funds, while Panel B-3's sample includes Financial Funds and Impact Funds. *t*-statistics are reported in parentheses. Panel C presents regressions of sector-average (columns (1) and (2)) and within-sector-ranked (columns (3) and (4)) Scope 1&2 and Scope 1&2&3 carbon emission intensity on fund indicator variables. Panel C-1 includes Financial Funds and Benchmark Funds, while Panel C-2 includes Financial Funds and Impact Funds. Regressions include year-quarter fixed effects, with standard errors clustered at the fund level where indicated.

Panel A: Average Carbon Intensity

	(1) Benchmark	(2) Sustainable	(3) Financial	(4) Impact
Carbon Intensity Scope 1&2	161.866	132.658	126.414	182.124
Carbon Intensity Scope 1&2&3	285.434	247.472	235.650	321.398
Sector-average Scope 1&2	168.094	146.766	135.754	248.990
Sector-average Scope 1&2&3	290.921	261.185	244.044	392.188
Within-sector-ranked Scope 1&2	4.637	4.669	4.702	4.453
Within-sector-ranked Scope 1&2&3	4.937	4.975	5.023	4.713
Observations	60191	16757	12513	1809

Panel B: Regression Analysis

Panel B-1: Financial Funds vs Benchmark Funds

	(1) Scope 1&2	(2) Scope 1&2	(3) Scope 1&2&3	(4) Scope 1&2&3
Financial Fund	-20.344***	-20.344***	-22.539***	-22.539***
	(-11.826)	(-3.859)	(-11.477)	(-3.365)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes
Observations	72696	72696	72696	72696
Adjusted R^2	0.022	0.022	0.041	0.041

Panel B-2: Impact Funds vs Benchmark Funds

	(1) Scope 1&2	(2) Scope 1&2	(3) Scope 1&2&3	(4) Scope 1&2&3
Impact Fund	28.806***	28.806*	51.127***	51.127***
	(6.923)	(1.710)	(10.844)	(2.635)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes
Observations	61995	61995	61995	61995
Adjusted R^2	0.016	0.016	0.032	0.032

Panel B-3: Impact Funds vs Financial Funds

	(1) Scope 1&2	(2) Scope 1&2	(3) Scope 1&2&3	(4) Scope 1&2&3
Impact Fund	50.584***	50.584***	75.560***	75.560***
	(13.654)	(2.894)	(17.181)	(3.707)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes
Observations	14087	14087	14087	14087
Adjusted R^2	0.028	0.028	0.058	0.058

Panel C: Within-Sector and Cross-Sector Analysis

Panel C-1: Within-Sector and Cross-Sector Financial Funds vs Benchmark Funds

	(1) Sector-Average Scope 1&2	(2) Sector-Average Scope 1&2&3	(3) Within-sector-ranked Scope 1&2	(4) Within-sector-ranked Scope 1&2&3
Financial Fund	-15.534***	-17.702**	0.043	0.039
	(-2.784)	(-2.556)	(1.039)	(0.990)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes
Observations	72704	72704	72696	72696
Adjusted R^2	0.024	0.047	0.018	0.030

Panel C-2: Within-Sector and Cross-Sector Impact Funds vs Financial Funds

	(1) Sector-Average Scope 1&2	(2) Sector-Average Scope 1&2&3	(3) Within-Sector-ranked Scope 1&2	(4) Within-sector-ranked Scope 1&2&3
Impact Fund	106.620***	136.566***	-0.257**	-0.320***
	(3.146)	(3.764)	(-2.344)	(-2.947)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes
Observations	14090	14090	14087	14087
Adjusted R^2	0.058	0.091	0.010	0.017

Table 5: Performance and Factor Tilts

This table reports sustainable funds' return performance and factor exposures relative to benchmark funds. The dependent variable is the monthly average net-of-fee excess return for each fund category, including three sustainable categories and the benchmark, computed from March 2014 to December 2023 using equal weights across funds. Panel A presents time-series regressions of the monthly average excess return on the Fama-French three factors, including and reporting the constant. Panel B reports panel regressions at the category–month level, comparing sustainable fund categories with the benchmark fund category. Interactions between fund category indicators and the factors capture differential loadings. Columns (1)-(2) include a panel of two time series, the Financial and Benchmark Fund categories; (3)-(4) the Moral and Benchmark Fund categories; and (5)-(6) the Impact and Benchmark Fund categories. Fund Category equals one for Financial, Moral, and Impact Funds in columns (1)-(2), (3)-(4), and (5)-(6), respectively.

Panel A: Time Series Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Benchmark Excess Return		Financial Excess Return		Moral Excess Return		Impact Excess Return	
MktRf	0.897***	0.877***	0.936***	0.916***	0.927***	0.900***	0.969***	0.939***
	(51.030)	(55.594)	(47.803)	(46.566)	(52.368)	(60.017)	(42.004)	(41.646)
SMB		0.097***		0.101***		0.138***		0.165***
		(3.700)		(3.065)		(5.515)		(4.392)
HML		0.097***		0.034		0.094***		-0.017
		(5.286)		(1.484)		(5.404)		(-0.644)
Constant	-0.251***	-0.205***	-0.247***	-0.212**	-0.303***	-0.246***	-0.284***	-0.241**
	(-3.071)	(-2.919)	(-2.714)	(-2.421)	(-3.677)	(-3.683)	(-2.646)	(-2.396)
Observations	117	117	117	117	117	117	117	117
Adjusted R^2	0.957	0.969	0.952	0.956	0.959	0.974	0.938	0.946

Panel B: Sustainable Funds Compared to Benchmark Funds

	(1)	(2)	(3)	(4)	(5)	(6)
	Financial Excess Return		Moral Excess Return		Impact Excess Return	
MKT	0.897***	0.877***	0.897***	0.877***	0.897***	0.877***
	(48.229)	(49.177)	(50.857)	(56.999)	(43.750)	(45.067)
MKT*Fund Category	0.038	0.039	0.030	0.022	0.072**	0.062**
	(1.458)	(1.560)	(1.194)	(1.034)	(2.470)	(2.258)
Fund Category	0.004	-0.007	-0.052	-0.041	-0.033	-0.036
	(0.033)	(-0.063)	(-0.445)	(-0.421)	(-0.243)	(-0.290)
SMB		0.097***		0.097***		0.097***
		(3.273)		(3.794)		(2.999)
HML		0.097***		0.097***		0.097***
		(4.675)		(5.419)		(4.285)
SMB* Fund Category		0.003		0.040		0.068
		(0.077)		(1.116)		(1.477)
HML* Fund Category		-0.063**		-0.003		-0.114***
		(-2.148)		(-0.110)		(-3.557)
Constant	-0.251***	-0.205**	-0.251***	-0.205***	-0.251***	-0.205**
	(-2.902)	(-2.582)	(-3.060)	(-2.993)	(-2.633)	(-2.366)
Observations	234	234	234	234	234	234
Adjusted R2	0.954	0.962	0.958	0.971	0.947	0.956

Table 6: Flow Performance Sensitivity

This table analyzes the sensitivity of fund flows to performance. For each year-month, fund alpha is calculated as the difference between realized excess return and predicted excess return based on the CAPM model, where the market factor loading is estimated using the previous 60 months of observations. Monthly alpha for each period is then computed accordingly. The independent variable, Alpha at time t, is calculated as a weighted average of returns from t – 1 to t – 18 using an exponential decay function with a decay parameter λ = 0.2. $\text{Alpha}_{jt} = \frac{\sum_{s=1}^{18} e^{-\lambda(s-1)}\tilde{\alpha}_{t-s}}{\sum_{s=1}^{18} e^{-\lambda(s-1)}}$. Column (1) includes a sample of Moral Funds and Benchmark Funds; columns (2) includes a sample of Impact Funds and Benchmark Funds; and column (3) includes a sample of Financial Funds and Benchmark Funds. Fund Type is an indicator variable identifying the fund type. In column (1), Fund Type equals one if the fund is a Financial Fund; in column (2), if the fund is a Moral Fund; and in column (3), if the fund is an Impact Fund. Control variables include lagged log total AUM, log age, lagged expense ratio, and a load dummy. Standard errors are double clustered by fund and year-month.

	(1) Moral Flow Rate	(2) Impact Flow Rate	(3) Financial Flow Rate
Alpha*Fund Type	-0.324***	-0.003	-0.049
	(-2.751)	(-0.014)	(-0.313)
Alpha	0.823***	0.823***	0.831***
	(12.303)	(12.315)	(12.391)
Fund Type	0.267**	0.394	0.150
	(2.059)	(1.521)	(1.243)
Flow Rate at t-19	0.051***	0.051***	0.052***
	(5.272)	(5.161)	(5.821)
Year Month FE	Yes	Yes	Yes
Double Clustering at Fund Year Month	Yes	Yes	Yes
Observations	100845	97446	114384
Adjusted R2	0.007	0.007	0.006

Table 7: Fund Votes on Environmental and Social Shareholder Proposals

This table reports the analysis of sustainable funds' votes on environmental and social shareholder proposals. For each fund-year, we calculate the percentages of Environmental and Social (ES) proposals received by the companies held by the fund for which the fund voted "For," "Against," "Abstain," and "Do Not Vote". In Panel A, summary statistics are provided for benchmark, financial, moral, and impact funds. In Panel B, the regression results are reported for regressing the voting percentages on the financial fund dummy and impact dummy variables. *t*-stats are reported in parentheses. All regressions in Panel B include year–quarter fixed effects, with standard errors clustered at the fund level. Panel C reports the results from regressions estimated at the fund–proposal level. The dependent variable is an indicator equal to one if the fund voted Yes on the proposal. The key independent variables are Outcome Proposal, a dummy variable indicating whether the proposal is an outcome proposal; Impact Fund, a dummy variable identifying impact funds; and the interaction between Impact Fund and Outcome Proposal. A proposal is defined as an Outcome Proposal if it falls into an ISS category whose proposal descriptions (AgendaGeneralDesc and ItemDesc) are not related to ESG reporting. This definition includes proposals with the following ISSAgendaItemIDs: S0224, S0411, S0416, S0703, S0732, and S0745. Column (1) includes all sustainable funds, and column (2) includes impact funds and benchmark funds. Firm fixed effects are included, and standard errors are clustered at the proposal level.

Panel A: Summary Statistics

	(1) Benchmark	(2) Financial	(3) Moral	(4) Impact
Vote "For" ES	25.1%	40.7%	47.0%	58.5%
Vote "Against" ES	63.7%	53.1%	47.2%	34.1%
Abstain ES Vote	7.0%	2.3%	2.4%	2.8%
Do not vote	2.8%	2.1%	2.0%	1.1%
Observations	20440	4317	1164	564

Panel B: Regression Results

Panel B-1: Financial vs Benchmark

	(1) Vote "For"	(2) Vote "Against"	(3) Abstain	(4) Do not Vote
Financial Fund	11.809***	-9.587***	-2.552***	-0.341
	(8.055)	(-6.378)	(-4.540)	(-0.812)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes
Observations	24757	24757	24757	24757
Adjusted R^2	0.135	0.090	0.036	0.052

Panel B-2: Impact Fund Analysis

	(1) Vote "For"	(2) Vote "Against"	(3) Abstain	(4) Do not Vote
Impact Fund	19.809***	-20.837***	-0.147	-1.102**
	(5.265)	(-5.341)	(-0.164)	(-2.293)
Year-quarter FE	Yes	Yes	Yes	Yes
Clustering at Fund	Yes	Yes	Yes	Yes
Observations	5591	5591	5591	5591
Adjusted R^2	0.115	0.129	0.018	0.042

Panel C: Outcome Proposal Analysis

	(1) Vote "For"	(2) Vote "For"
Impact Fund*Outcome Proposal	0.181***	0.212***
	(3.326)	(3.317)
Outcome Proposal	-0.061**	-0.070***
	(-2.242)	(-3.186)
Impact Fund	0.201***	0.349***
	(19.218)	(33.393)
Sample	All Sustainable Funds	Impact and Benchmark Funds
Firm FE	Yes	Yes
Clustering at Proposal	Yes	Yes
Observations	71651	251183
Adjusted R^2	0.113	0.136

Table 8: Changes in ESG Ratings and Carbon Emission Intensity During Holding Periods

This table reports the changes in portfolio company ESG ratings and Carbon Emission Intensity during periods when firms are held by different fund types. The starting time of a stock's holding session in an impact fund is defined as the first year-month when the stock appears in the fund's holdings and the fund is classified as an impact fund at the time of entry. The holding session is defined as the consecutive months during which the stock is observed in the fund's holdings without interruption, and the fund remains classified as an impact fund throughout. The end time of the holding session is defined as the last year-month when the stock appears in the holdings. A similar definition applies to benchmark, financial, and moral funds. For funds with at least 12 (18, or 24) months of holdings, the change in ESG ratings and Carbon Emission Intensity is computed as the difference between the values 12 (18, or 24) months after the entry time and those at the entry time. Carbon intensity is interpolated between reporting years using weighted averages of last and next available data. Panel A reports the average change in ESG ratings across different fund types, while Panel B reports the average change in carbon emission intensity across different fund types.

Panel A: Changes in ESG Scores During Holding Periods

	(1) Benchmark	(2) Financial	(3) Moral	(4) Impact
12-Month ESG Score Change	0.126	0.151	0.136	0.152
12-Month E Score Change	0.165	0.143	0.111	0.153
12-Month S Score Change	0.088	0.158	0.162	0.148
Observations	1897871	258266	210784	32950

	(5) Benchmark	(6) Financial	(7) Moral	(8) Impact
18-Month ESG Score Change	0.194	0.230	0.206	0.221
18-Month E Score Change	0.246	0.222	0.187	0.260
18-Month S Score Change	0.141	0.235	0.226	0.179
Observations	1668374	234078	195567	27005

	(9) Benchmark	(10) Financial	(11) Moral	(12) Impact
24-Month ESG Score Change	0.260	0.283	0.247	0.326
24-Month E Score Change	0.320	0.278	0.200	0.343
4-Month S Score Change	0.197	0.286	0.295	0.302
Observations	1377209	201844	181546	21827

Panel B: Changes in Emission Intensity During Holding Periods

	(1) Benchmark	(2) Financial	(3) Moral	(4) Impact
12-Month Scope 1&2 Intensity Change	-4.276	-7.286	-4.148	-9.421
12-Month Scope 1&2&3 Intensity Change	-5.465	-9.251	-4.681	-13.448
Observations	2040420	254761	210948	33796

	(5) Benchmark	(6) Financial	(7) Moral	(8) Impact
18-Month Scope 1&2 Intensity Change	-6.994	-12.370	-9.347	-12.758
18-Month Scope 1&2&3 Intensity Change	-9.401	-17.221	-11.888	-16.666
Observations	1786126	227448	193925	28022

	(9) Benchmark	(10) Financial	(11) Moral	(12) Impact
24-Month Scope 1&2 Intensity Change	-10.312	-17.243	-14.765	-16.075
24-Month Scope 1&2&3 Intensity Change	-14.519	-27.100	-21.820	-20.241
Observations	1447923	193831	178498	22452

Figure 1: Growth of Sustainable Mutual Funds, 2014–2023

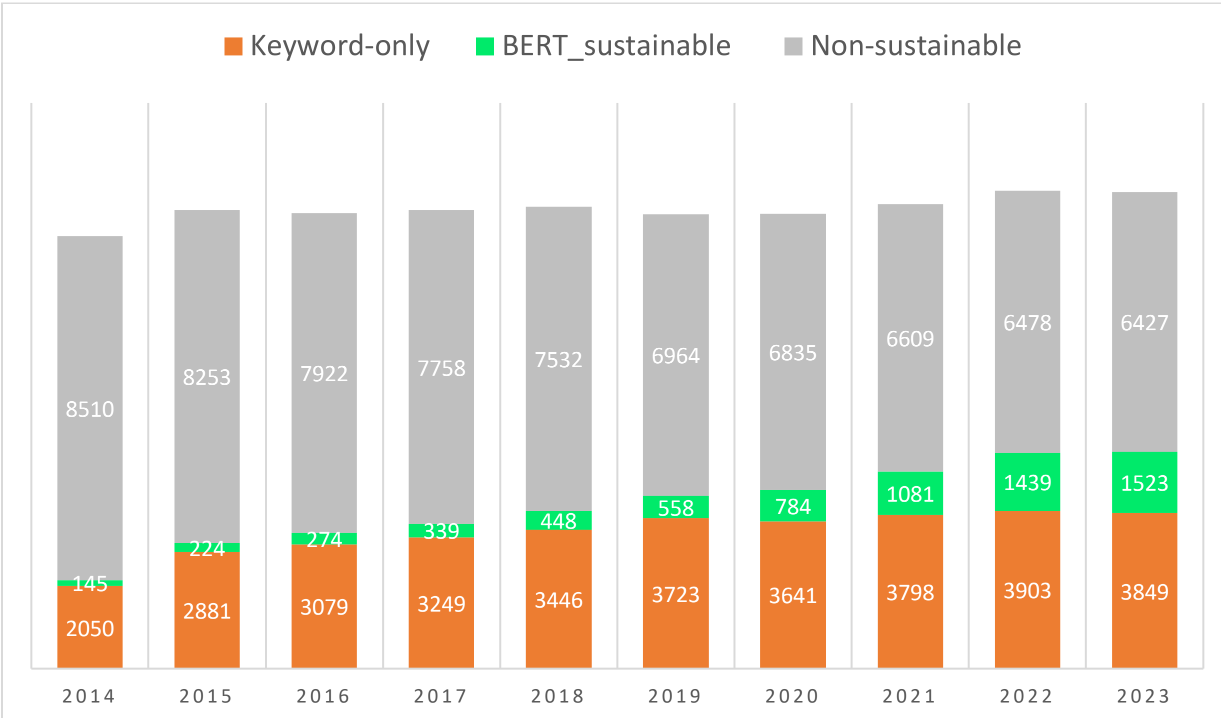

Panel A: Fund Count

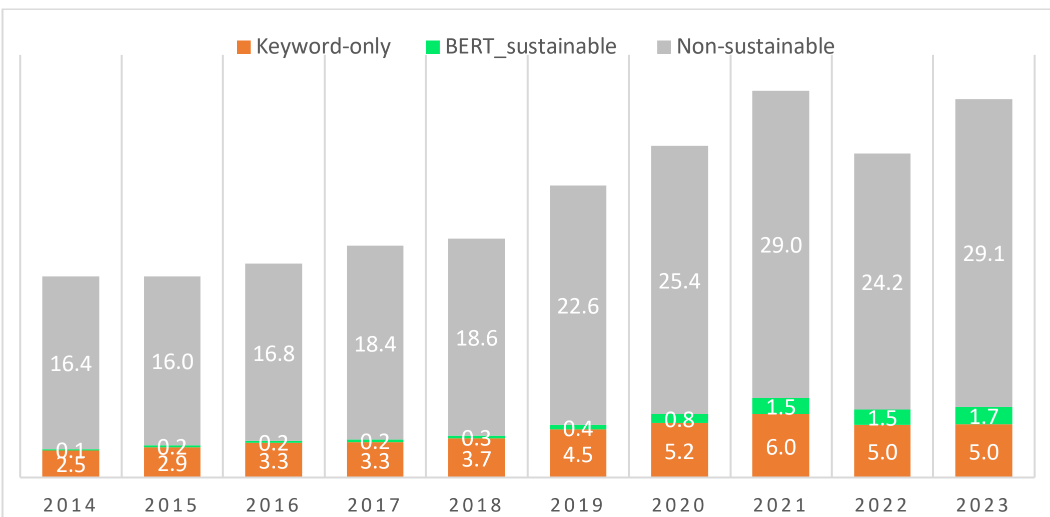

Panel B: Fund AUM ($T)

This figure plots the number of funds and total assets under management (AUM) from 2014Q2 to 2023Q4 for three groups: sustainable funds, keyword-only funds, and non-sustainable funds. Sustainable Funds are funds that the BERT model identifies as financial, moral, or impact. "Keywords-only" refers to funds that mention at least one keywords related to sustainability in the "Principal Investment Strategy" section of its prospectus, but that the BERT model identifies zero sentences as meeting the sustainability criteria. "Non-sustainable" refers to funds that do not mention any sustainability-related keywords.

Figure 2: Distribution of Sustainable Fund Classifications

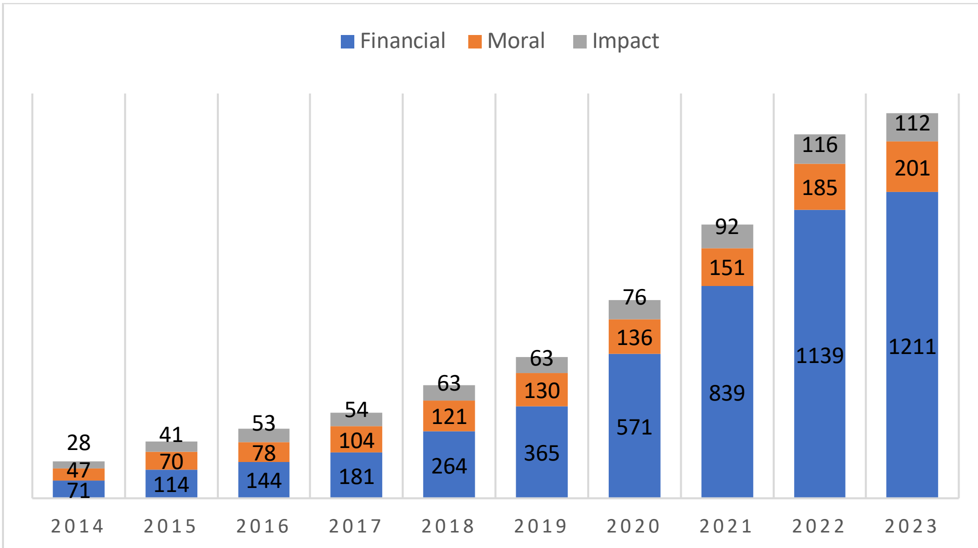

Panel A: Fund AUM ($B)

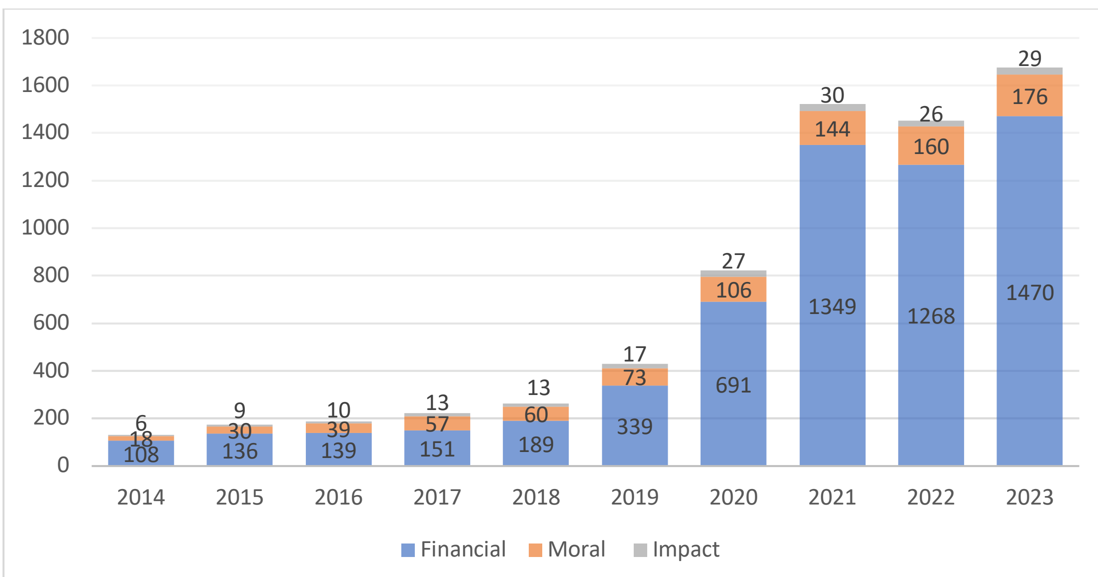

Panel B: % of AUM

This figure plots the cross-sectional distribution of sustainable fund classifications by year. Classifications are based on the % of sustainable sentences for each fund that are classified as financial, moral, and impact. A fund is assigned the category with the largest % share of all sustainable sentences. When there is a two-way (three-way) tie for the largest % share, each category receives 50% (33%) of the share count/AUM of the fund.

Figure 3: Sector Weights, ESG Ratings, and Emission Intensity

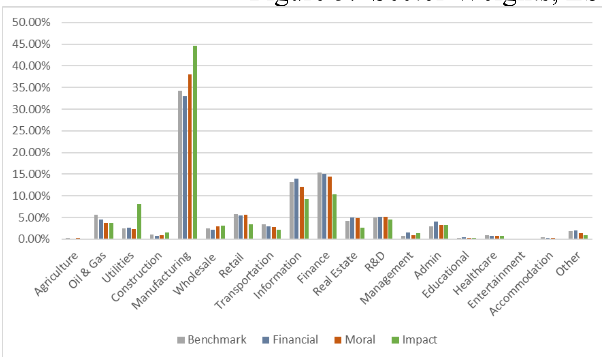

Panel A: Sector Weights

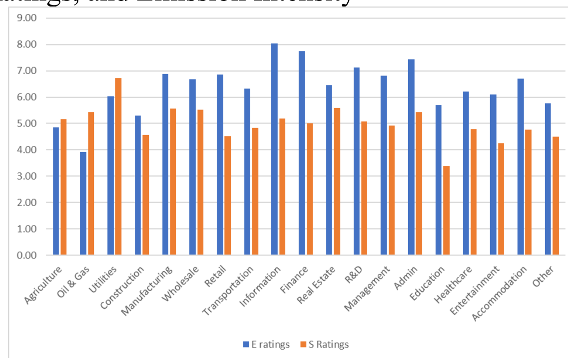

Panel B: Sector Average E and S Ratings

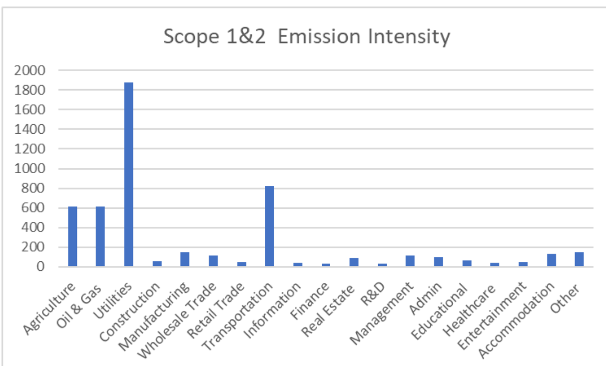

Panel C: Scope 1&2 Emission Intensity

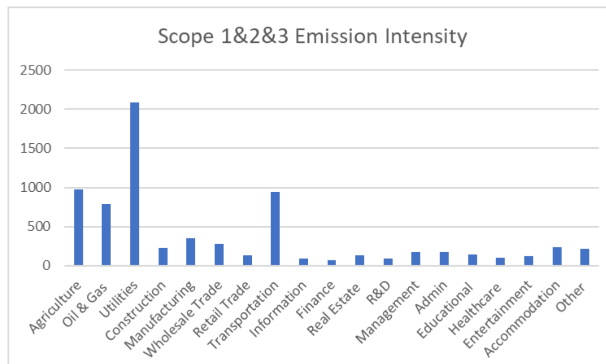

Panel D: Scope 1&2&3 Emission Intensity

This figure compares portfolio sector weights of sustainable fund holdings by type and sector-average E and S ratings and emission intensity of covered firms by MSCI and Trucost, respectively. Panel A shows the sector portfolio weights for benchmark, financial, moral, and impact funds. Panel B shows the average MSCI covered firms' E and S ratings by sector. Panel C and D show the average Trucost covered firms' scope 1&2 and 1&2&3 emission intensity by sector, respectively.

Figure 4: DiD Analysis of Firm Emission Intensity during Holding Periods

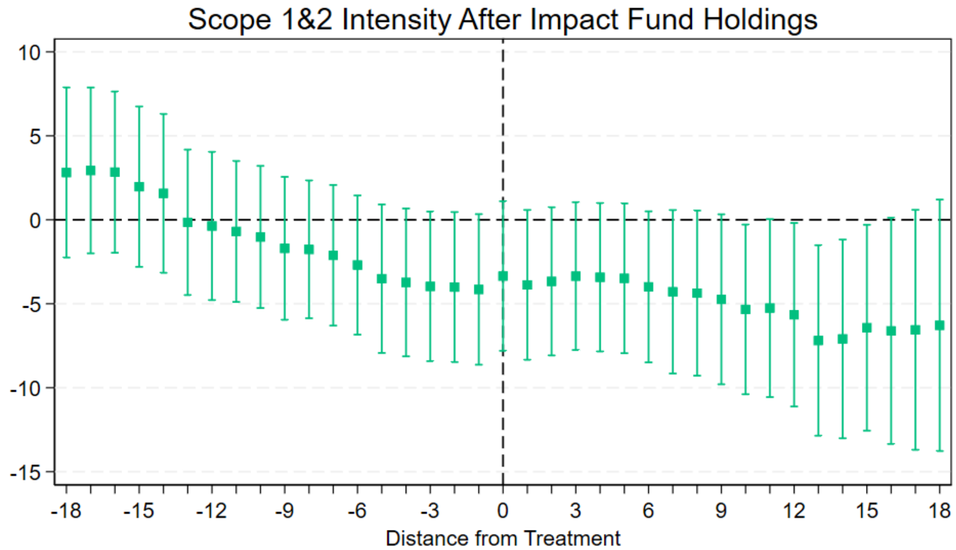

Panel A: Impact Fund Holdings

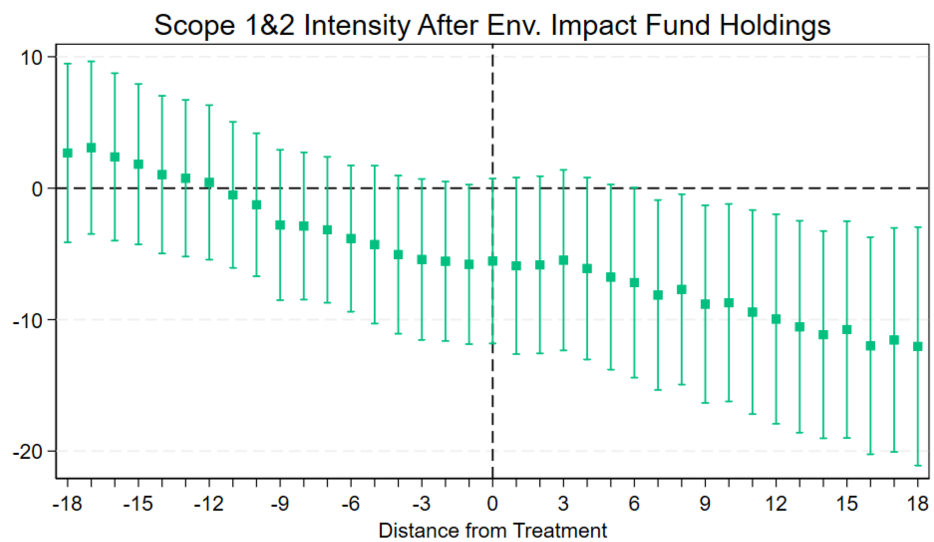

Panel B: Environmental Impact Fund Holdings

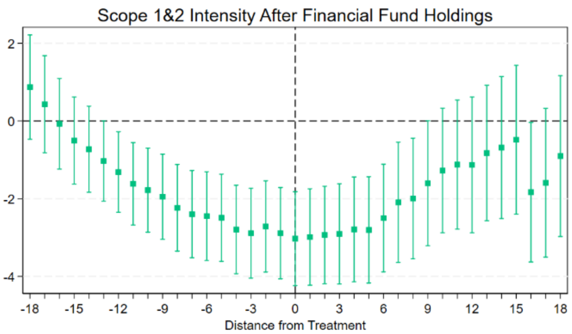

Panel C: Financial Fund Holdings

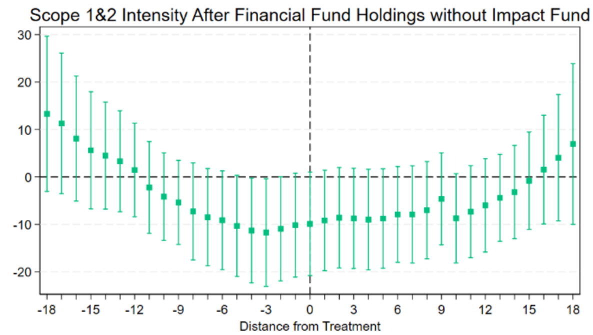

Panel D: Financial Fund Holdings with No Overlapping Impact Fund Holdings

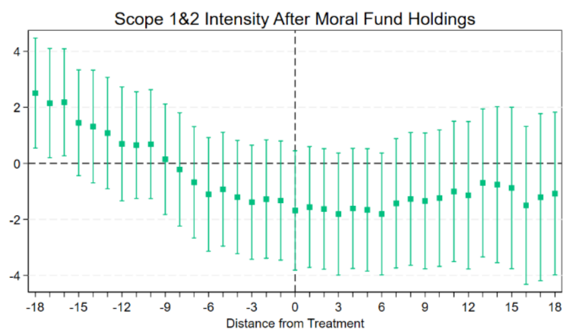

Panel E: Moral Fund Holdings

This figure shows the dynamic treatment effect on firms' Scope 1 & 2 carbon intensity after acquisition by sustainable funds. We estimate a dynamic treatment effect model.

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

Where $Holding_{i,t}^{k}$ is a dummy variable equal to one for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$. For each treated firm—defined as a stock newly purchased by a sustainable fund—we identify a matched control firm based on ROA and log(total assets) within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between t–24 and t+24. The outcome variable is Scope 1&2 carbon intensity. Firm and year-month fixed effects are included, and standard errors are clustered at the firm level. Confidence intervals are reported at the 90% level. Panel A reports results for impact fund holdings, Panel B reports environmental impact fund holdings, focusing on funds in the top quintile based on their support for environmental related shareholder proposals, Panel C reports results for financial fund holdings, Panel D reports results for financial fund holdings without firms that are also held by impact funds between *t*-18 to *t*+18, and Panel E reports results for moral fund holdings.

Internet Appendix for *Decoding Sustainable Investment Strategies: Bridging Intentions and Outcomes*

1. Classification of ESG-Sentences

Financial Sentences

The first category is Financial. A sentence is classified as Financial if it indicates that the fund employs ESG information to enhance financial performance. In general, sentences in this category share the following characteristics.

(1) direct references to using using ESG for financial performance

"...but the firm believes that responsible corporate behavior with respect to ESG factors can contribute to positive and sustainable long-term financial performance."

(2) direct references to using using ESG as risk factor

"...evaluating ESG-related risks as part of its research recommendations."

(3) direct references to using using ESG as material issues

"The fund will not invest in companies that Newton deems to have material environmental, social or governance issues.."

(4) key terms indicating ESG use for financial performance, not externalities

- ESG as attractive investment attributes
- ESG as Sustainable business practices
- ESG controversies

"Newton may also monitor certain companies in whose securities the fund has invested for emerging environmental, social or governance controversies and issues and may update a company's ESG quality review rating on the basis of such monitoring."

- ESG ratings

"As part of this research process, Invesco may use third-party ESG ratings, company reporting, and engagement with management."

- ESG issues

"The Fund's Adviser may consider information about environmental, social and governance issues (also referred to as ESG) in its bottom-up stock selection process when making investment decisions."

(5) Index indicating Financial motives

- ESG MSCI
- ESG Sustainalytics US Inc.

(6) use ESG information as an input when making investment selections

"As such, the Adviser evaluates ESG factors as part of the investment analysis process and this forms an integral component of the Adviser's quality rating for all companies."

Investment analysis process suggests that ESG information is used for financial performance.

(7) references to companies benefiting from ESG transitions

"Stocks are included in the Underlying Index based on the Index Provider's evaluation that such companies will substantially benefit from a societal transition toward the use of cleaner energy, zero-CO 2 renewables, and conservation."

Moral Sentences

The second category is Moral. A sentence is classified as Moral if it indicates that the fund excludes certain types of companies. We assume that such exclusions are based on ethical rather than financial considerations. A sentence is not labeled as Moral if it explicitly states that exclusions are made for non-Moral reasons. In our manual review, we treat exclusion criteria as Moral-related unless they are specifically described as non-Moral.

"The Fund is fossil-fuel free, as it does not invest in companies that derive significant revenues from the extraction, exploration, production or refining of fossil fuels; the Fund may invest in companies that use fossil fuel-based energy to power their operations or for other purposes."

Impact Sentences

The third category is Impact. A sentence is classified as Impact if it indicates that the fund uses ESG information specifically for externality considerations, without reference to financial performance. Specifically referencing externality considerations without mentioning financial performance, and indicating that the fund values impact in its own right.

(1) specifically referencing externality considerations without mentioning financial performance, and indicating that the fund values impact in its own right.

"Rockefeller's sustainability and impact evaluation considers environmental, social, and governance criteria such as corporate governance practices, product quality and safety, workplace diversity practices, environmental impact and sustainability, community investment and development, and human rights record."

"Strictly in accordance with its guidelines and mandated procedures, WilderShares, LLC (the "Index Provider") compiles and maintains the Underlying Index, which is composed of stocks of publicly traded companies listed on a major exchange in the United States that are engaged in the business of the advancement of cleaner energy and conservation or are important to the development of clean energy."

(2) directly stating the creation of positive impact or the measurement of impact, without providing specific details.

" The Fund may invest a significant amount of its assets in taxable and tax-exempt municipal bonds that finance community projects whose primary purpose, in the Adviser's view, is a positive impact to the community in which the project is located."

"The Adviser researches the use of proceeds of each bond invested in by the Fund utilizing both quantitative metrics and qualitative details to measure the impact achieved."

Unclassified Sentences

The Unclassified category includes sentences that are too ambiguous to be clearly assigned to a single category, as well as sentences containing generic ESG content that does not reflect any specific investor philosophy. A few typical examples are provided below.

"The Fund seeks to achieve its investment objective while applying an environmental, social and governance (ESG methodology developed by the Fund's subadviser in the selection of portfolio investments.)"

"ESG considerations are fully integrated across all asset classes."

"Final investment decisions are made by the portfolio managers typically on the basis of market conditions as well as technical and ESG considerations with respect to both specific instruments and the overall composition of the portfolio."

2. Manually Classified Sample Construction

In this section, we provide a detailed explanation of the process used to construct the manually classified training and testing sample. We begin with a list of U.S. sustainable mutual funds compiled and published by Morningstar. Morningstar began compiling this annual list in 2018, and we obtained the lists for 2018, 2019, 2020, and 2022 either directly from the company's website or from other researchers who have used these data. We focus on this set of U.S. mutual funds identified by Morningstar as pursuing sustainability goals because we expect these funds to be more likely than other funds to describe their sustainability objectives in the investment strategy section of their prospectuses. Oversampling true positives, defined as sentences classified as Financial, Moral, or Impact, from imbalanced samples is a common strategy in machine learning and textual analysis to improve classification balance and performance (He and Garcia, 2009). Using a random number generator, we select 2,834 sentences for manual classification.

To ensure that our manually classified sentences also include potential sustainable funds outside the Morningstar list, we supplement the sample with 741 sentences drawn from funds not listed by Morningstar. This additional sample from Non-Morningstar funds also oversamples true positives and follows standard practices in machine learning and textual analysis.

In total, our manually classified sample is drawn from 362 fund prospectuses and includes 3,575 ESG-related sentences. We then randomly divide the sample into 2,950 sentences for training and 625 sentences for testing.

3. Comparison with GenAI

In this section, we compare the predictions of our BERT model with those of a generative AI (GenAI) model. One natural methodological question, given that GenAI models are widely used for daily tasks, is why we do not rely on a GenAI model for classification. To address this question, we first briefly discuss the methodological differences among three potential textual analysis approaches, with a primary focus on comparing the BERT model and the generative AI model. Table IA1 provides a summary of the key differences across model types.

[Insert Table IA1: Language Model Comparison]

One possible classification approach is topic modeling, such as Latent Dirichlet Allocation (LDA). These methods are conceptually similar to principal component analysis in that they allow the data itself to be represented by a set of latent topics (i.e., they are unsupervised), with each topic potentially corresponding to a class. However, models of this type are limited in their ability to capture complex relationships between words. Moreover, they are not trained on ground-truth labels, meaning we cannot directly assess whether the learned topics correspond to correct human coding or generate classifications that align with the true labels. Topic models are not well suited to our setting because the subtlety of sustainable investment motives requires a model that can understand and process complex relationships between words and can be trained to learn the ground-truth classifications identified by researchers.

The BERT model is the approach used in our paper. It has the advantage of being able to capture complex relationships between words through its transformer architecture. BERT is pre-trained so that the base model can already infer word meanings from their full context. Moreover, it is a supervised model and can be trained on ground-truth labels provided by researchers. By fine-tuning the model using researcher-labeled data, BERT can generate classifications that align closely with the ground truth, combining the benefits of rich contextual language representations from pre-training with task-specific learning from supervised re-training.

One alternative classification approach is to use a Generative AI (GenAI) model. Modern GenAI models—such as ChatGPT, Google's Gemini series, and Claude—are built on transformer architectures. While these models are substantially larger than BERT, they share a similar architectural foundation and are capable of modeling complex relationships between words and phrases. Because these models are pretrained on massive text corpora, one could directly prompt them to classify sustainable investment goals into financial, moral, and impact categories based on the model's semantic understanding of these concepts. In this setting, however, the training data underlying the model implicitly serves as the "ground truth," rather than labels explicitly defined and provided by the researcher. For example, when ChatGPT is asked to classify funds based on its reading of fund prospectuses, the classification is generated without supervision and without being trained or calibrated on researcher-provided ground-truth labels. This methodology is problematic because retail investors and the media have been documented to often misunderstand sustainable goals, and such misunderstandings reflected in texts appearing in news articles and online forum discussions may be transmitted to the model through its pretraining data. In the next subsection, we present the results obtained using this approach and compare them with our BERT-based classification.

One way to "force" a GenAI model to produce outputs closer to researcher-defined ground truth is through carefully designed prompting or through Retrieval-Augmented Generation (RAG). Prompt engineering can steer the model's responses toward the categories intended by the researcher; for example, when interacting with ChatGPT, prompts are often structured to constrain the model to answer a specific question in a desired format. RAG further extends this idea by augmenting prompting with an explicit retrieval step, in which researcher-provided ground-truth examples are supplied to the model at inference time, thereby encouraging classifications that align more closely with the reference labels.

However, this approach is not an ideal implementation of GenAI for classification tasks. If the objective is to ground the model in human researchers' labeled training samples, a supervised model such as BERT can directly learn this mapping through fine-tuning, achieving satisfactory performance with substantially greater stability and without hallucinations. Consequently, when researcher-defined ground truth is available for training, GenAI models become a less suitable choice relative to BERT-based classifiers.

3.1. GPT 5 Results

In this section, we compare the classifications generated by the BERT model with those produced by a generative AI model. Specifically, we use OpenAI's GPT-5 model for comparison. We adopt "gpt-5-2025-08-07" as a representative snapshot to benchmark BERT model performance against that of a ChatGPT-style model. The “gpt-5-2025-08-07” model is a fixed snapshot of GPT-5 released on August 7, 2025. It is designed for general-purpose coding and chatbot tasks and is widely used and regarded as highly successful. We use the following prompt for classification:

"You are an investor reading a mutual fund prospectus. Given a sentence related to ESG (Environmental, Social, and Governance), classify the fund's investment approach based on its primary intent: Financial – primarily focused on generating financial returns; Moral – guided by ethical or moral values; Impact – aiming to achieve measurable environmental or social outcomes; Unclear – if there isn’t enough information to make a clear classification. Respond with only one word: 'Financial', 'Moral', 'Impact', or 'Unclear'."

[Insert Table IA2: ChatGPT Model Performance]

Table IA2 provides the ChatGPT model prediction performance. We report the performance of ChatGPT-5 prediction performance on the 625 sentences in the testing sample. We find that ChatGPT-5 prediction performance is significantly lower than BERT model. The accuracy for financial sentences is 0.72 with ChatGPT-5, compared to 0.91 in BERT model, which is a 19-percentage points decreases. The Moral and Impact accuracy are also smaller. We find that the decrease in performance is mainly due to the low precision of ChatGPT-5 prediction, in Financial, Moral, and Impact prediction, most strongly for financial sentences. The Precision decreases from 0.85 to 0.47, from 0.85 to 0.62, from 0.83 to 0.61, for Financial, Moral, and Impact Respective. Too many false positives are given with ChatGPT prediction. Overall, we find that ChatGPT-5 underperforms our BERT model.

3.2. Conceptual Background and Additional Textual Analysis Materials

BERT is able to capture relationships between words because it is trained to model contextual dependencies through self-attention. BERT uses a Transformer architecture in which each word attends to every other word in the sentence. This self-attention mechanism allows the model to weight the relevance of surrounding words when forming a representation for a given token, enabling it to capture semantic, syntactic, and long-range relationships. During pretraining, BERT is optimized using a masked language modeling objective, which forces the model to infer a missing word from its full context. As a result, the learned representations encode rich information about how words relate to each other across different contexts, allowing BERT to distinguish meaning based on usage rather than relying solely on individual word frequencies.

BERT fine-tuning and retrieval-augmented generation (RAG) represent fundamentally different approaches to textual analysis. Fine-tuning a BERT model involves supervised learning in which model parameters are updated using labeled data to directly optimize performance on a specific task, such as text classification. As a result, task-relevant information is embedded in the model's parameters, yielding stable, reproducible predictions that are well suited for empirical analysis and hypothesis testing. In contrast, RAG-based large language models do not learn task-specific decision rules from labeled data. Instead, they retrieve relevant external documents at inference time and generate free-form responses conditioned on the retrieved text and the prompt. While RAG systems are flexible and effective for open-ended question answering, their outputs are inherently prompt-dependent and stochastic, and they lack a fixed decision boundary. Consequently, RAG is not a substitute for fine-tuned discriminative models when the research objective requires consistent classification, interpretability, and reproducibility.

Prompting and retrieval-augmented generation (RAG) are related but conceptually distinct. A prompt refers to the textual input used to condition a large language model's behavior at inference time and relies entirely on the model's pretrained, parametric knowledge. In contrast, RAG is a system-level architecture that augments prompting with an explicit retrieval step, in which relevant external documents are retrieved and incorporated into the prompt before generation. While both approaches operate without updating model parameters, RAG enables the model to condition its output on non-parametric, externally stored information, whereas prompting alone does not involve retrieval or access to external knowledge. Consequently, prompting is a component of AG, but RAG represents a broader framework designed to extend language models beyond their fixed internal knowledge.

4. Discussion of Additional Results

4.1. Sector Exclusion by Moral Funds

We examine whether moral funds that express exclusionary investment goals in their prospectuses indeed underweight sectors they claim to avoid. In the manual coding of prospectus sentences, we find that alignment of investment goals with moral or ethical values is frequently expressed as prohibition against certain business activities and exclusion of companies engaged significantly in such activities. Some of the excluded activities reflect traditional religious values (e.g., alcohol, gambling, weapons, and abortion), while others are grounded in modern ethical or scientific concerns (e.g., fossil fuels and climate impact, tobacco and health risks). Moral funds vary in orientation—some are religious (e.g., Catholic, Presbyterian), while others are secular—leading to heterogeneity in exclusion practices. For example, some funds may avoid casinos but not coal, while others exclude fossil fuels and weapons but not abortion providers. As a result, average portfolio weights in excluded sectors are expected to be lower than those of benchmark funds, but not zero.

We use the following NAICS sectors corresponding to commonly excluded categories for this analysis: "Tobacco", "Oil & Gas Extraction", "Natural Gas Distribution", "Coal Mining", "Casino & Gambling", and "Aerospace" (including weapons manufacturers). These variables are defined in Appendix A2.

Panel A of Table IA3 presents the sector weights of fund portfolios in these categories by benchmark, sustainable, and moral funds. As expected, moral funds have substantially lower portfolio weights in these sectors than benchmark funds.

[Insert Table IA3: Sector Exclusion by Moral Funds]

In Panel B, we report the results of regression analysis. We find that moral funds hold significantly smaller portion of their portfolios in tobacco, gas, casino and aerospace stocks relative to benchmark funds in the same quarter-year. The coefficients for oil & gas and coal mining are also negative but insignificant, possibly because traditional religious funds—which make up a substantial share of our sample—typically do not exclude fossil fuels. We plan to collect fund-specific exclusion categories to enable more granular analysis and sharper inferences.

4.2.Environmental Rating Decomposition

To further examine the heterogeneity of sustainable funds' investment strategies, we decompose the overall environmental (E) ratings into four underlying MSCI subcategories and repeat the comparison analysis from Table IA4.

[Insert Table IA4: Environmental Rating Decomposition]

Panel A of Table IA4 reports the value-weighted average scores for the four MSCI environmental subcategories—**Climate Change**, **Natural Capital**, **Pollution & Waste**, and **Environmental Opportunities**—for benchmark, sustainable, financial, and impact funds. Among benchmark funds, the highest scores are observed in the Climate Change category, while the lowest appear in Environmental Opportunities.

According to MSCI's ESG Ratings Methodology Manual, the *Climate Change* subcategory captures companies' exposure to operational and supply chain risks stemming from carbon pricing, regulatory emissions limits, and physical climate hazards. *Natural Capital* assesses exposure to biodiversity loss, deforestation, unsustainable raw material sourcing (e.g., palm oil), and water stress. *Pollution & Waste* reflects potential liabilities from regulatory actions related to e-waste, packaging, toxic emissions, and contamination. Finally, *Environmental Opportunities* evaluates how well-positioned firms are to benefit from the transition to a low-carbon economy—particularly in clean technology, green buildings, and renewable energy.

Comparing fund types, we find that financial funds generally hold stocks with higher environmental subcategory scores than benchmark funds across all four dimensions except *Pollution & Waste*. In contrast, impact funds hold stocks with higher *Pollution & Waste* and *Environmental Opportunities* scores relative to financial funds.

Panel B presents regression results controlling for time fixed effects. In Panel B-1, financial funds are shown to hold firms with significantly higher scores in all four subcategories compared to benchmark funds, consistent with their preference for ESG "leaders" as a means of managing material risk and enhancing returns. Panel B-2 compares impact funds to financial funds. We find that impact funds hold stocks with significantly higher *Pollution & Waste* and *Environmental Opportunities* scores—consistent with investing in companies whose core businesses (e.g., renewable energy, green buildings, and clean tech) contribute directly to environmental solutions. These represent a class of investments where the business model itself serves as a generator of positive environmental impact.

However, we also find that impact funds hold companies with significantly lower *Natural Capital* scores relative to financial funds. These firms are more likely to be associated with operations that generate negative externalities—such as deforestation, water depletion, or unsustainable sourcing—and thus face regulatory, physical, or reputational risks. While MSCI ratings may only imperfectly capture the full extent of these externalities, this finding is consistent with the interpretation that impact funds allocate capital to relatively poor environmental performers, potentially as part of an engagement-based strategy aimed at improvement. Notably, the higher *Environmental Opportunities* scores among impact fund holdings are also consistent with our earlier finding that these funds tilt more heavily toward the utilities sector (see Table 2 and Figure 3), where opportunities for energy transition are more prevalent.

Taken together, the decomposition analysis reinforces the view that impact funds differ fundamentally from financial funds in how they approach sustainability. Impact funds allocate capital to firms that are both positioned to reduce environmental harms (e.g., by improving supply chains or reducing emissions) and to deliver environmental solutions through their core products and services. In contrast, financial funds predominantly allocate to firms that already score well on environmental metrics across the board. These differences reflect divergent approaches: one emphasizing risk-adjusted return enhancement through ESG integration, the other emphasizing real-world impact generation through improvement and innovation.

4.3. Robust Flow Performance Results

This section examines whether the results on the association between mutual funds' stated sustainable investment strategies and flow–performance sensitivity are robust to the use of alternative decay functions. We show that the results remain robust when using different decay functions and alternative fixed effects.

Table IA5 investigates how flows respond to performance. Following Barber, Huang, and Odean (2016), we construct alphas using an exponential decay function applied to prior monthly returns, with the CAPM alpha serving as the performance metric given its stronger predictive power for flows. Panel A compares moral and benchmark funds. Across specifications, we find that moral funds earn significantly higher unconditional flows (roughly 0.3% per month). More notably, moral funds exhibit significantly attenuated flow–performance sensitivity. For example, in column (2), a 1 percentage point decline in alpha reduces flows by 0.388% for benchmark funds, but by only 0.195% (= 0.388 − 0.193) for moral funds. Using the 18-month decay in column (4), the analogous decline is 0.823% for benchmark funds versus 0.499% for moral funds. These findings are consistent with moral funds attracting investors who derive utility from ethical exclusionary screens and who are less responsive to performance.

[Insert Table IA5: Flow Performance Sensitivity]

Panel B compares impact and benchmark funds. Unlike moral funds, impact funds do not exhibit statistically different flow–performance sensitivity relative to benchmarks. Although impact funds display significantly higher unconditional flows in columns (1) and (2), the difference becomes insignificant when using the 18-month alpha decay (columns (3)–(4)). Overall, we find no evidence that impact funds are systematically matched with less performance-sensitive investors.

Panel C compares financial and benchmark funds. As expected, we find no statistically significant difference in flow–performance sensitivity between the two groups. This is consistent with financial funds competing for flows from traditional pecuniary investors.

Our findings underscore the importance of distinguishing among sustainable funds by stated objectives. Our results are robust when using an alternative three-month decay function and when excluding fixed effects

4.4. Voting Behavioral of Moral Funds

To further unpack the heterogeneity among moral funds, we classify them into subgroups based on the presence of hybrid objectives: *pure moral*, *financial-moral*, *impact-moral*, and *financial-impact-moral*.[1] Panel A-2 presents the average support rates for each subgroup. We find that *pure moral* funds—those with no hybrid classification—exhibit voting behavior nearly identical to financial funds. In contrast, moral funds with impact-related language (e.g., *impact-moral*, *financial-impact-moral*) support ES proposals at substantially higher rates. These findings suggest that moral funds with impact elements adopt more activist voting behavior, while pure moral funds behave similarly to financially motivated funds in the proxy voting context.

[Insert Table IA6: Moral Fund Votes on Environmental and Social Shareholder Proposals]

[1] As described in the methodology section, funds are classified into one of three types—financial, moral, or impact—based on the most frequent classification among their sustainability-related sentences. For example, a fund is classified as "moral" if the number of sentences labeled as moral exceeds those labeled as financial or impact. Within the set of moral funds, we further define subgroups such as *financial-moral* (moral funds with at least one sentence also classified as financial) and *impact-moral* (moral funds with at least one sentence also classified as impact), and so on.

4.5. Emission Intensity Changes: Difference-in-Differences Analysis

We report the coefficients used to construct Figure 4, which presents the difference-in-differences analysis. We also provide results for Scope 1&2&3 emission intensity. The results are reported in Table IA7. We find very similar results, with no evidence of pretrends, and treated firms exhibiting significantly lower emission intensity after treatment.

[Insert Table IA7: Difference-in-Differences Analysis of Emission Intensity]

In our main specification, if two funds begin holding firm i at the same time, we still set the indicator equal to one. In Figure IA1, we present robustness results in which, in such cases, we instead replace the indicator with the total number of simultaneous holding events that share the same entry date. We find very similar results, with no evidence of pretrends and treated firms exhibiting significantly lower emissions intensity after treatment.

[Insert Figure IA1: Robust DiD Analysis of Firm Emission Intensity during Holding Periods]

5. Shareholder Proposal Measures

This section describes our methodology for constructing environmental and social (ES) shareholder proposal measures. We identify ES proposals using ISS category codes based on Table A1 in He, Kahraman, and Lowry (2023), supplemented by manual review. To construct outcome-focused proposals, we examine the resolution type and agenda general description for each ISS category code and review ISS's detailed descriptions. Because many ISS categories include both disclosure-oriented and outcome-oriented components, we identify six ISS categories that are exclusively outcome-oriented and use these as our measure of pure outcome-oriented proposals.

6. Additional Tables and Figure

Table IA1: Language Model Comparison

Model Type	Topics (LDA) model	BERT model	Generative AI
Process complex relationship between words	No	Yes	Yes
Supervised	No	Yes	No (but prompt and RAG)
Trained on "Ground truth"	No	Yes	No (but prompt and RAG)
Model output stable over time	N/A	Yes	Yes/No
Prone to deviation and hallucination from "ground truth"	N/A	No	Yes

Table IA2: ChatGPT Model Performance

This table reports the performance of the BERT model trained to classify ESG-related sentences. We report sentence-level prediction performance using the "gpt-5-2025-08-07" model with the following prompt: "You are an investor reading a mutual fund prospectus. Given a sentence related to ESG (Environmental, Social, and Governance), classify the fund's investment approach based on its primary intent: Financial—primarily focused on generating financial returns; Moral—guided by ethical or moral values; Impact—aiming to achieve measurable environmental or social outcomes; Unclear—if there is insufficient information to make a clear classification. Respond with only one word: 'Financial,' 'Moral,' 'Impact,' or 'Unclear.'" We compare the ChatGPT-generated classifications with manual classifications from a testing sample of 625 sentences. Four different model performance measures are calculated to measure the accuracy of BERT classification. Accuracy is the ratio of (true positives + true negatives) divided by the total number of observations (fraction of correct classifications). Precision is the ratio of true positives divided by the sum of true positives and false positives. Recall is the ratio of true positives divided by the sum of true positives and false negatives. *f1* is defined as $[\frac{2}{\frac{1}{Precision}+\frac{1}{Recall}}]$. We also provide the results from BERT model for comparison.

	ChatGPT-5				BERT			
	Accuracy	Precision	Recall	f1	Accuracy	Precision	Recall	f1
Financial	0.7216	0.4749	0.764	0.5857	0.9088	0.8467	0.7888	0.8167
Moral	0.96	0.6182	0.8947	0.7312	0.9840	0.8500	0.8947	0.8718
Impact	0.912	0.6125	0.6712	0.6405	0.9440	0.8276	0.6575	0.7328

Table IA3: Sector Exclusion by Moral Funds

This table reports the portfolio weights of moral funds in sectors commonly excluded on ethical grounds. In Panel A presents mean sector weights for benchmark, sustainable, and moral funds. In Panel B, the regression results are reported for regressing the fund sector weights on the moral dummy variables for a combined sample of benchmark funds and moral funds with year-quarter fixed effects and standard errors clustered at the fund level. *t*-stats are reported in parentheses.

Panel A: Excluded Sector Weights

	Benchmark	Sustainable	Moral
Tobacco	0.38%	0.13%	0.09%
Oil & Gas Extraction	1.11%	0.96%	0.85%
Gas Distribution	0.35%	0.42%	0.25%
Coal Mining	0.03%	0.02%	0.03%
Casino & Gambling	0.05%	0.03%	0.01%
Aerospace	0.61%	0.40%	0.31%
Observations	59423	16463	3492

Panel B: Regression Analysis

	(1)	(2)	(3)	(4)	(5)	(6)
	Tobacco	Oil & Gas	Gas Distribution	Coal Mining	Casino & Gambling	Aerospace
Moral Fund	-0.247***	-0.162	-0.094***	-0.013	-0.028***	-0.315***
	(-6.420)	(-1.572)	(-2.762)	(-0.969)	(-8.477)	(-4.763)
Observations	62915	62915	62915	62915	62915	62915
Adjusted R^2	0.013	0.042	0.004	0.005	0.022	0.007
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Table IA4: Environmental Rating Decomposition

This table reports the decomposition of the environmental (E) ratings of portfolio holdings by fund type. Subcategory ratings are from MSCI and cover Climate Change, Natural Capital, Pollution & Waste, and Environmental Opportunities. For each fund-quarter, we calculate the value-weighted Subcategory E ratings for all MSCI-rated stock holdings. In Panel A, the mean Climate Change, Natural Capital, Pollution & Waste, and Environmental Opportunities ratings are provided for the benchmark funds, sustainable funds, financial funds, and impact funds. In Panel B, the regression results are reported for regressing the fund-quarter Subcategory E ratings on the financial dummy and impact dummy variables. *t*-stats are reported in parentheses. All regressions include year–quarter fixed effects, with standard errors clustered at the fund level where indicated.

Panel A: Average Environmental Ratings

	(1) Benchmark	(2) Sustainable	(3) Financial	(4) Impact
Climate Change	7.251	7.750	7.804	7.676
Natural Capital	6.272	6.611	6.692	6.323
Pollution & Waste	5.333	5.094	5.019	5.505
Environmental Opportunities	4.479	4.729	4.678	5.299
Observations	59413	16459	12245	1789

Panel B: Regression Results

Panel B-1: Financial Funds vs Benchmark Funds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Climate Change		Natural Capital		Pollution & Waste		Env. Opportunities	
Financial	0.209***	0.209***	0.199***	0.199***	0.278***	0.278***	0.104***	0.104***
Fund	(15.850)	(4.148)	(13.588)	(3.484)	(20.654)	(5.949)	(12.331)	(3.172)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes	No	Yes	No	Yes
Observations	71510	71510	71145	71145	67715	67715	67321	67321
Adjusted R^2	0.137	0.137	0.068	0.068	0.377	0.377	0.042	0.042

Panel B-2: Impact Funds vs Financial Funds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Climate Change		Natural Capital		Pollution & Waste		Env. Opportunities	
Impact	0.029	0.029	-0.297***	-0.297**	0.214***	0.214**	0.678***	0.678***
Fund	(0.996)	(0.243)	(-8.700)	(-2.295)	(6.672)	(2.009)	(27.508)	(6.018)
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustering at Fund	No	Yes	No	Yes	No	Yes	No	Yes
Observations	13789	13789	13706	13706	12490	12490	12867	12867
Adjusted R^2	0.107	0.107	0.050	0.050	0.282	0.282	0.077	0.077

Table IA5: Flow Performance Sensitivity

This table analyzes the sensitivity of fund flows to performance. For each year-month, fund alpha is calculated as the difference between realized excess return and predicted excess return based on the CAPM model, where the market factor loading is estimated using the previous 60 months of observations. Monthly alpha for each period is then computed accordingly. The independent variable, Alpha at time t, is calculated as a weighted average of returns from t – 1 to t – 3 (t – 18) using an exponential decay function with a decay parameter λ = 0.2. $\mathrm{Alpha}_{jt} = \frac{\sum_{s=1}^{3(18)} e^{-\lambda(s-1)}\tilde{\alpha}_{t-s}}{\sum_{s=1}^{3(18)} e^{-\lambda(s-1)}}$. Panel A reports results for Moral Funds and Benchmark Funds. Panel B reports results for Impact Funds and Benchmark Funds. Panel C reports results for Financial Funds and Benchmark Funds. Control variables include lagged log total AUM, log age, lagged expense ratio, and a load dummy. Standard errors are double clustered by fund and year-month.

Panel A: Moral Funds vs Benchmark Funds

	(1) Flow Rate	(2) Flow Rate	(3) Flow Rate	(4) Flow Rate
Alpha*Moral Fund	-0.197***	-0.193***	-0.324***	-0.324***
	(-3.882)	(-3.839)	(-2.712)	(-2.751)
Alpha	0.371***	0.388***	0.815***	0.823***
	(10.368)	(10.808)	(12.039)	(12.303)
Moral Fund	0.301**	0.301***	0.269**	0.267**
	(2.620)	(2.628)	(2.069)	(2.059)
Flow Rate at t-4	0.023***	0.023***		
	(3.914)	(3.963)		
Flow Rate at t-19			0.049***	0.051***
			(4.963)	(5.272)
Year Month FE	No	Yes	No	Yes
Double Clustering at Fund Year Month	Yes	Yes	Yes	Yes
Alpha Decay Month	3	3	18	18
Observations	124968	124968	100845	100845
Adjusted R2	0.005	0.007	0.006	0.007

Panel B: Impact Funds vs Benchmark Funds

	(1) Flow Rate	(2) Flow Rate	(3) Flow Rate	(4) Flow Rate
Alpha*Impact Fund	-0.003	0.003	-0.010	-0.003
	(-0.038)	(0.030)	(-0.041)	(-0.014)
Alpha	0.371***	0.389***	0.814***	0.823***
	(10.360)	(10.817)	(12.033)	(12.315)
Impact Fund	0.503**	0.513**	0.385	0.394
	(2.148)	(2.185)	(1.503)	(1.521)
Flow Rate at t-4	0.023***	0.023***		
	(3.915)	(3.967)		
Flow Rate at t-19			0.049***	0.051***
			(4.869)	(5.161)
Year Month FE	No	Yes	No	Yes
Double Clustering at Fund Year Month	Yes	Yes	Yes	Yes
Alpha Decay Month	3	3	18	18
Observations	121202	121202	97446	97446
Adjusted R2	0.005	0.007	0.005	0.007

Panel C: Financial Funds vs Benchmark Funds

	(1) Flow Rate	(2) Flow Rate	(3) Flow Rate	(4) Flow Rate
Alpha*Financial Fund	-0.068	-0.049	-0.059	-0.049
	(-0.728)	(-0.525)	(-0.376)	(-0.313)
Alpha	0.371***	0.391***	0.814***	0.831***
	(10.323)	(10.985)	(11.974)	(12.391)
Financial Fund	0.167	0.186	0.136	0.150
	(1.358)	(1.578)	(1.119)	(1.243)
Flow Rate at t-4	0.022***	0.022***		
	(4.120)	(4.170)		
Flow Rate at t-19			0.050***	0.052***
			(5.544)	(5.821)
Year Month FE	No	Yes	No	Yes
Double Clustering at Fund Year Month	Yes	Yes	Yes	Yes
Alpha Decay Month	3	3	18	18
Observations	139755	139755	114384	114384
Adjusted R2	0.005	0.007	0.005	0.006

Table IA6: Moral Fund Votes on Environmental and Social Shareholder Proposals

This table reports the analysis of moral funds' votes on environmental and social shareholder proposals. For each fund-year, we calculate the percentages of Environmental and Social (ES) proposals received by the companies held by the fund for which the fund voted "For," "Against," "Abstain," and "Do Not Vote". Mean values are provided for pure moral, financial-moral, moral-impact, and financial-moral-impact funds.

	(1)	(2)	(3)	(4)
	Pure Moral	Financial-Moral	Moral-Impact	Financial-Moral-Impact
Vote "For" ES	40.5%	54.2%	62.7%	79.2%
Vote "Against" ES	53.9%	41.4%	29.2%	10.5%
Abstain ES Vote	2.7%	1.7%	1.9%	2.5%
Do not vote	2.0%	2.0%	2.6%	0.1%
Observations	761	250	108	45

Table IA7: Difference-in-Differences Analysis of Emission Intensity

This table reports the dynamic treatment effect of being held by sustainable funds. We estimate a dynamic treatment effect model.

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

Where $Holding_{i,t}^{k}$ is a dummy variable equal to one for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$. For each treated firm—defined as a stock newly purchased by a sustainable fund—we identify a matched control firm based on ROA and log(total assets) within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between *t–24* and *t+24*. We define a fund's holding session as the uninterrupted period during which a stock appears in the fund's portfolio and the fund remains classified as a given type (impact, financial, or moral) at the time of entry. The session ends when either the stock disappears from holdings or the fund changes classification. Investment sessions lasting six months or fewer are excluded. Similar definitions apply across all fund types. The outcome variable is Scope 1&2 and Scope 1&2&3 carbon intensity. Firm and year-month fixed effects are included, and standard errors are clustered at the firm level. Columns (1) and (2), (3) and (4), (5) and (6), and (7) and (8) report results for impact, environmental impact, financial, and moral fund holdings, respectively. Environmental impact funds are defined as impact funds in the top quintile based on their support for environmental related shareholder proposals.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Impact		Env. Impact		Financial		Moral	
	1&2	1&2&3	1&2	1&2&3	1&2	1&2&3	1&2	1&2&3
t-18	2.815	1.954	2.678	2.485	0.873	0.791	2.509**	2.786**
	(0.914)	(0.607)	(0.648)	(0.598)	(1.069)	(0.923)	(2.101)	(2.213)
t-17	2.938	1.963	3.081	2.554	0.431	0.357	2.149*	2.716**
	(0.979)	(0.628)	(0.772)	(0.634)	(0.567)	(0.446)	(1.813)	(2.144)
t-16	2.836	1.934	2.378	2.166	-0.073	-0.164	2.179*	2.630**
	(0.972)	(0.640)	(0.614)	(0.549)	(-0.103)	(-0.218)	(1.877)	(2.123)
t-15	1.970	1.104	1.825	1.765	-0.503	-0.573	1.449	2.079*
	(0.679)	(0.369)	(0.492)	(0.464)	(-0.739)	(-0.775)	(1.262)	(1.679)
t-14	1.571	0.792	1.031	1.129	-0.727	-0.715	1.316	2.016
	(0.546)	(0.267)	(0.283)	(0.300)	(-1.081)	(-0.981)	(1.075)	(1.531)
t-13	-0.155	-0.713	0.757	0.993	-1.030	-1.081	1.081	1.666
	(-0.059)	(-0.259)	(0.209)	(0.264)	(-1.640)	(-1.554)	(0.895)	(1.285)
t-12	-0.372	-0.805	0.438	0.934	-1.315**	-1.317*	0.695	1.022
	(-0.139)	(-0.289)	(0.122)	(0.251)	(-2.089)	(-1.886)	(0.562)	(0.779)
t-11	-0.697	-1.311	-0.512	-0.064	-1.617**	-1.611**	0.649	0.962
	(-0.273)	(-0.489)	(-0.152)	(-0.018)	(-2.508)	(-2.265)	(0.561)	(0.781)
t-10	-1.023	-1.370	-1.262	-0.401	-1.780***	-1.789**	0.686	0.800

	(-0.398)	(-0.507)	(-0.381)	(-0.114)	(-2.705)	(-2.486)	(0.580)	(0.642)
t-9	-1.700	-2.084	-2.802	-1.846	-1.950***	-2.035***	0.145	0.409
	(-0.657)	(-0.768)	(-0.806)	(-0.498)	(-2.925)	(-2.835)	(0.121)	(0.325)
t-8	-1.762	-2.325	-2.880	-1.785	-2.235***	-2.193***	-0.217	-0.112
	(-0.707)	(-0.884)	(-0.846)	(-0.487)	(-3.295)	(-3.015)	(-0.177)	(-0.087)
t-7	-2.115	-2.726	-3.166	-1.878	-2.398***	-2.402***	-0.677	-0.442
	(-0.831)	(-1.013)	(-0.938)	(-0.513)	(-3.514)	(-3.301)	(-0.560)	(-0.350)
t-6	-2.696	-3.410	-3.834	-2.544	-2.449***	-2.523***	-1.107	-0.962
	(-1.070)	(-1.275)	(-1.133)	(-0.688)	(-3.533)	(-3.410)	(-0.897)	(-0.748)
t-5	-3.510	-4.212	-4.292	-3.084	-2.492***	-2.611***	-0.925	-0.814
	(-1.306)	(-1.486)	(-1.176)	(-0.779)	(-3.653)	(-3.565)	(-0.748)	(-0.631)
t-4	-3.730	-4.276	-5.054	-3.628	-2.794***	-2.780***	-1.204	-1.220
	(-1.394)	(-1.520)	(-1.381)	(-0.917)	(-4.034)	(-3.758)	(-0.980)	(-0.953)
t-3	-3.966	-4.485	-5.427	-4.187	-2.888***	-2.835***	-1.388	-1.369
	(-1.465)	(-1.574)	(-1.458)	(-1.039)	(-4.105)	(-3.766)	(-1.121)	(-1.063)
t-2	-4.007	-4.522	-5.559	-4.417	-2.715***	-2.564***	-1.276	-1.472
	(-1.476)	(-1.587)	(-1.508)	(-1.105)	(-3.806)	(-3.327)	(-0.993)	(-1.110)
t-1	-4.147	-4.584	-5.795	-4.510	-2.888***	-2.749***	-1.328	-1.418
	(-1.522)	(-1.593)	(-1.571)	(-1.122)	(-4.045)	(-3.563)	(-1.027)	(-1.055)
t	-3.349	-3.459	-5.540	-4.354	-3.028***	-2.785***	-1.682	-1.744
	(-1.237)	(-1.210)	(-1.453)	(-1.060)	(-4.122)	(-3.475)	(-1.299)	(-1.296)
t+1	-3.878	-4.067	-5.905	-4.896	-2.987***	-2.751***	-1.560	-1.500
	(-1.430)	(-1.426)	(-1.447)	(-1.125)	(-3.963)	(-3.366)	(-1.190)	(-1.102)
t+2	-3.668	-3.962	-5.833	-4.690	-2.935***	-2.750***	-1.627	-1.497
	(-1.368)	(-1.397)	(-1.425)	(-1.081)	(-3.849)	(-3.368)	(-1.244)	(-1.095)
t+3	-3.354	-3.838	-5.475	-4.627	-2.907***	-2.798***	-1.810	-1.680
	(-1.254)	(-1.359)	(-1.312)	(-1.050)	(-3.708)	(-3.323)	(-1.367)	(-1.209)
t+4	-3.424	-3.864	-6.113	-5.029	-2.790***	-2.616***	-1.609	-1.500
	(-1.275)	(-1.356)	(-1.453)	(-1.135)	(-3.403)	(-2.968)	(-1.235)	(-1.098)
t+5	-3.487	-3.749	-6.763	-5.440	-2.804***	-2.651***	-1.661	-1.574
	(-1.286)	(-1.310)	(-1.579)	(-1.213)	(-3.373)	(-2.971)	(-1.250)	(-1.138)
t+6	-3.995	-4.223	-7.183	-5.902	-2.499***	-2.337***	-1.808	-1.598
	(-1.461)	(-1.471)	(-1.634)	(-1.286)	(-2.966)	(-2.593)	(-1.367)	(-1.162)
t+7	-4.288	-4.559	-8.129*	-6.843	-2.093**	-1.798*	-1.427	-0.994
	(-1.449)	(-1.487)	(-1.851)	(-1.508)	(-2.227)	(-1.772)	(-1.016)	(-0.686)
t+8	-4.367	-4.657	-7.704*	-6.391	-1.996**	-1.563	-1.269	-0.768
	(-1.461)	(-1.510)	(-1.751)	(-1.399)	(-2.117)	(-1.539)	(-0.881)	(-0.516)
t+9	-4.737	-4.843	-8.821*	-7.273	-1.602	-1.177	-1.340	-0.945
	(-1.539)	(-1.534)	(-1.931)	(-1.550)	(-1.639)	(-1.119)	(-0.909)	(-0.623)
t+10	-5.338*	-5.829*	-8.714*	-7.274	-1.275	-0.719	-1.242	-0.920
	(-1.738)	(-1.822)	(-1.908)	(-1.531)	(-1.308)	(-0.679)	(-0.838)	(-0.604)
t+11	-5.259	-5.778*	-9.426**	-7.604	-1.122	-0.646	-1.002	-0.718
	(-1.632)	(-1.724)	(-1.998)	(-1.554)	(-1.111)	(-0.590)	(-0.658)	(-0.460)
t+12	-5.654*	-6.148*	-9.953**	-8.299*	-1.130	-0.669	-1.139	-0.982

	(-1.702)	(-1.781)	(-2.055)	(-1.662)	(-1.064)	(-0.583)	(-0.712)	(-0.599)
t+13	-7.187**	-7.470**	-10.544**	-9.118*	-0.825	-0.170	-0.698	-0.034
	(-2.084)	(-2.100)	(-2.152)	(-1.802)	(-0.778)	(-0.149)	(-0.434)	(-0.020)
t+14	-7.096**	-7.289*	-11.145**	-9.044*	-0.685	-0.132	-0.760	0.072
	(-1.972)	(-1.957)	(-2.326)	(-1.840)	(-0.616)	(-0.109)	(-0.449)	(0.041)
t+15	-6.428*	-6.392*	-10.756**	-8.455*	-0.483	0.102	-0.878	-0.117
	(-1.725)	(-1.656)	(-2.146)	(-1.676)	(-0.414)	(0.081)	(-0.501)	(-0.065)
t+16	-6.614	-6.998*	-11.989**	-9.637*	-1.833*	-1.896	-1.500	-0.370
	(-1.615)	(-1.657)	(-2.389)	(-1.863)	(-1.678)	(-1.597)	(-0.875)	(-0.210)
t+17	-6.554	-6.989	-11.544**	-8.912*	-1.591	-1.488	-1.207	-0.160
	(-1.509)	(-1.567)	(-2.229)	(-1.681)	(-1.366)	(-1.173)	(-0.666)	(-0.086)
t+18	-6.283	-6.595	-12.036**	-8.945	-0.903	-0.708	-1.078	0.359
	(-1.380)	(-1.409)	(-2.183)	(-1.596)	(-0.717)	(-0.518)	(-0.610)	(0.196)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster at Firm	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	111661	111661	65717	65717	230796	230796	254848	254848
Adjusted R2	0.953	0.959	0.956	0.964	0.930	0.936	0.930	0.937

Figure IA1: Robust DiD Analysis of Firm Emission Intensity during Holding Periods

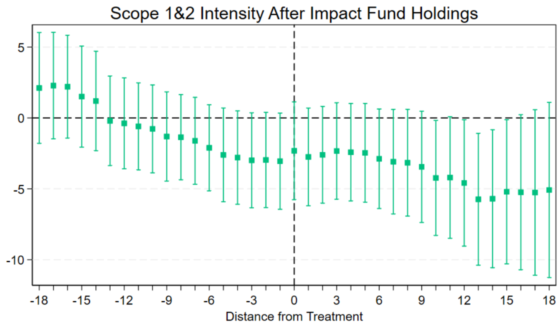

Panel A: Impact Fund Holdings

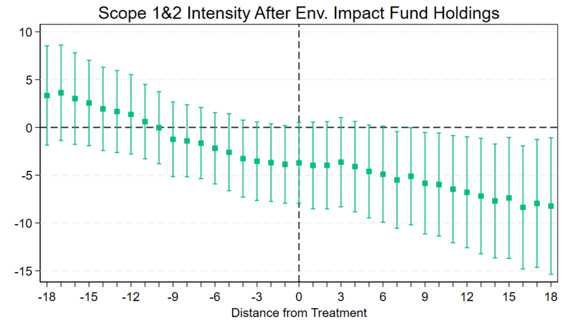

Panel B: Environmental Impact Fund Holdings

This figure shows the dynamic treatment effect on firms' Scope 1 & 2 carbon intensity after acquisition by sustainable funds. We estimate a dynamic treatment effect model.

$$y_{i,t} = \sum_{k=-18}^{18} Holding_{i,t}^{k} + FEs + \varepsilon_{i,t}$$

where $Holding_{i,t}^{k}$ is a the total number of simultaneous holding events for firm i at calendar time t if t is k months before/after the first month in which firm i becomes held by a fund of a given sustainability type for $k \in [-18,18]$. For each treated firm—defined as a stock newly purchased by a sustainable fund—we identify a matched control firm based on ROA and log(total assets) within the same 2-digit NAICS sector. Control firms are further required to have received no sustainable fund investment between t–24 and t+24. The outcome variable is Scope 1&2 carbon intensity. Firm and year-month fixed effects are included, and standard errors are clustered at the firm level. Confidence intervals are reported at the 90% level. Panel A reports results for impact fund holdings, Panel B reports environmental impact fund holdings, focusing on funds in the top quintile based on their support for environmental related shareholder proposals.