

Decoding Sustainable Investment Strategies: Bridging Intentions and Actions

First draft: November 18, 2023

This draft: March 13, 2024

Keer Yang
kkeyang@ucdavis.edu
UC Davis

Ayako Yasuda
asyasuda@ucdavis.edu
UC Davis

This research has benefitted from the comments of Brad Barber, Aneesh Raghunandan, Adam Zhang, and seminar and conference participants at UC Davis and the Minnesota Family and Friends Conference. 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.

Decoding Sustainable Investment Strategies: Bridging Intentions and Actions

Abstract

Our study employs a novel machine-learning approach to uncover the underlying intentions driving U.S. mutual funds' sustainable investment strategies. The identified intentions include financial value, categorical morality, and impact. While categorical morality is implemented via exclusion, impact is generated via active investment and engagement, and they each satisfy distinct nonpecuniary preferences. We find that the majority of funds labeled as “sustainable” prioritize enhancing financial value or adhering to categorical morality over generating impact. Notably, our method reveals the prevalence of hybrid funds, blending financial and moral objectives. This blended strategy amounts to rebalancing portfolios with high ESG-rated stocks to mitigate risk, while underweighting sin industries. Impact-driven funds, though less prevalent, exhibit distinctive management practices, actively engaging with portfolio companies and supporting social and environmental shareholder proposals. Our results underscore the negative correlation between financial and impact goals. These findings illuminate the need for clear classification labels to differentiate between 'doing well and being good' versus 'doing good'. By clarifying fund objectives, our methodology enhances investor welfare by facilitating informed decision-making tailored to individual preferences. Understanding the heterogeneous underlying preferences of sustainable investors and the potential tension between them can in turn inform policy decisions, investment practices, and corporate behavior.

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

1. Introduction

The 21st century has ushered in a new era where global challenges like climate change and social inequality demand urgent attention and innovative solutions. In this context, sustainable investment has emerged as a promising mechanism to respond to these pressing issues while also delivering financial returns. The literature documents that at least three distinct goals motivate investors to allocate capital to sustainable investments: (i) to enhance the financial resilience of one's portfolio, (ii) to morally screen out unwanted categories of investments, and (iii) to induce investee companies to generate positive impact. However, despite the rapid growth of assets under management by sustainable funds, relatively little is known about which of these underlying intentions are prioritized by sustainable fund managers. This article aims to fill this gap using a novel machine-learning approach and a dataset of fund prospectuses, holdings-level ESG ratings and funds' governance practices.

Financial value investors prioritize financial returns while considering non-financial factors like environmental, social, and governance (ESG) issues for risk management. Categorical morality investors exclude morally objectionable industries from their portfolios based on ethical criteria. Impact investors seek to actively generate positive social or environmental outcomes through their investments. Which of these three intentions are prioritized the most by sustainable fund managers? Do funds pursue pure strategies that are mutually exclusive, or do they employ hybrid strategies? Which strategies are positively or negatively correlated with each other? How do these strategies translate into actions by the funds? In this article, we address these questions by using a novel model and dataset.

We are not the first to argue that lack of clarity about sustainability leads to confusion and misunderstanding. Starks (2023) in her AFA presidential address argues that lack of consensus on the meaning of sustainable finance leads to misunderstandings, and that the confusion is in large part “due to differences in whether motivation arises from *value* or *values*, that is, from regarding the ESG qualities of an investment as important to its financial *value* or, as consistent with one's *values*.” (pp.1837-8).¹ She further makes a forceful case that “finance

¹ Pedersen, Fitzgibbons, and Pomorski (2021) define “ESG-aware” investors similarly to “value” investors in Starks (2023). Chowdhry, Davies, and Waters (2019), Oehmke and Opp (2022), and Landier and Lovo (2022), among others, model impact investors who internalize the positive (negative) externalities generated (avoided) by the companies. Pastor, Stambaugh and Taylor (2021), Berk and Binsbergen (2022), and Lo and Zhang (2021), among others, model categorical moral investors whose utility depends on the greenness of their portfolios. See Section 2 for a detailed literature review for the three distinct investor motivations for sustainable investing.

researchers have an opportunity to make important contributions to the literature and society by conducting research that considers both pecuniary and nonpecuniary aspects of ESG, taking an objective stance with regard to the incentives, costs, and benefits related to sustainable finance, and that provides evidence on the associated economic implications.” We further distinguish the “values” (i.e., nonpecuniary)-driven motivation into moral exclusion and impact generation. We show that the moral exclusion and financial value motivations are often pursued together in hybrid strategies (positive correlation) whereas impact generation is rarely pursued by funds prioritizing financial value (negative correlation).

These positive and negative correlations between the strategies are likely due to an inherent tension between the private benefit of reducing ESG risk in one’s portfolio vs. the public benefit of bearing the ESG risk in one’s portfolio to generate positive externality. For example, a strategy of holding a brown firm in one’s portfolio and inducing it to turn green via engagement is beneficial to the public but is privately costly and risky. Thus, only impact investors whose nonpecuniary preferences internalize the negative externality generated by all firms are willing to bear such risk. In contrast, moral exclusion investors are willing to only hold already green firms, and this is compatible with financial value investors’ motivation of mitigating ESG risk. These differential interactions between pecuniary and nonpecuniary preferences have implications for the role of sustainable finance as solutions to the global challenges (e.g., climate change) as opposed to risk management tools to protect private wealth. Moral funds and impact funds also show other systematic differences in their portfolio and voting decisions, reflecting the distinct underlying nonpecuniary preferences (disutility derived from holding sin stocks vs. utility derived from generating positive impact).

In Section 3, we build a supervised machine-learning model that is trained to decode investor sustainability goals expressed at the sentence level. We apply this model to the section of U.S. mutual fund prospectuses that discusses fund strategies and classify each fund based on the mix of intentions that the model detects. Notably a fund can have multiple sentences that indicate multiple – and thus hybrid – sustainability intentions. This modeling approach gives us the opportunity to examine both the extensive and intensive margin of the three, potentially competing, sustainability goals. Having decoded funds’ stated sustainability goals, we then hypothesize the relationship between the stated goals and the fund managers’ actions related to portfolio holdings, cross- and within-industry composition, and voting on environmental and

social shareholder proposals. Section 4 describes the dataset we construct by merging the sustainable fund list by Morningstar with the CRSP mutual fund database, MSCI corporate ESG ratings, and ISS shareholder proposal voting database to test these hypotheses.

Section 5 provides the main empirical results of the article. Among the U.S. mutual funds identified as sustainable by Morningstar in 2018-2022, the majority prioritize enhancing financial value or adhering to categorical morality over generating impact. Specifically, 54% state they seek financial value, while 39% seek categorical morality and only 33% state they seek causal impact. We find that hybrid funds are quite popular, especially those seeking financial value and categorical morality. In contrast, there is a negative correlation between financial and impact goals. To measure relative intensity of one goal vs. the other, we construct a continuous variable based on the percentage of the total sentences classified as one of the three categories. This continuous intensity measure helps us examine the intensive margin of each of the three intentions, and to disentangle the effect of each distinct goal from one another.

Next, we assess the ESG characteristics of fund portfolios across the three investor goals using MSCI company ESG ratings to compute annual value-weighted average ESG ratings for each portfolio. MSCI's explicit definition of ESG ratings as measuring the resilience of company value against industry-specific ESG risks makes them a suitable tool for distinguishing between investor goals. Specifically, MSCI's ratings are designed for use by "financial value" investors, focusing on financial resilience rather than societal or environmental impact. As such, we hypothesize that funds primarily seeking financial value ("financial" funds) tend to hold stocks with higher MSCI ESG ratings. Our findings support this hypothesis, revealing that funds with greater "financial" intensity in their goals hold stocks with higher MSCI ESG ratings, all else being equal. Conversely, funds prioritizing impact tend to hold stocks with lower MSCI ESG ratings, reflecting their focus on generating positive societal or environmental change rather than financial resilience. Notably, moral intensity of funds shows no significant correlation with MSCI ESG ratings, consistent with the concept of categorical exclusion.

We then dissect the components of ESG ratings to discern their impact on fund portfolios. Our analysis reveals that funds with a stronger focus on financial value motivation tend to prioritize stocks with higher Environmental ratings in Climate Change, while showing lower allocation to Environmental Opportunities. Additionally, they favor stocks with superior Social ratings, particularly in Human Capital. This pattern indicates a tilt towards stocks resilient

against stranded asset or carbon regulation risk, as well as those with lower labor-related litigation or reputational risk. In contrast, funds emphasizing impact tend to hold stocks with higher Environmental ratings in Environmental Opportunities, while exhibiting lower ratings in Climate Change, Human Capital, and Stakeholder Opposition. These findings align with impact-oriented funds' preference for stocks offering significant potential for improving Environmental and Social performance. Moreover, the emphasis on Environmental Opportunities reflects their inclination towards stocks poised for impactful growth, such as those in the renewable energy sector.

MSCI ESG ratings incorporate both risks and opportunities, with different factors being considered financially material across various sectors. Consequently, average ESG ratings vary across industries. Higher overall ESG ratings could indicate that (i) financially motivated sustainable funds tilt towards sectors with high average ESG ratings, or (ii) they prefer stocks with higher ESG ratings compared to industry peers within each sector, or (iii) a combination of both factors. Our analysis reveals that financially motivated funds achieve high ESG ratings primarily through the second mechanism, systematically holding stocks with higher ESG ratings relative to their industry peers. While they also demonstrate some inclination towards sectors with high average ESG ratings, this tendency is predominantly observed in hybrid funds that pursue both financial and moral objectives.

We hypothesize that moral funds achieve their goal by categorically excluding sectors that are preordained to be misaligned with investors' values, or "sin" stocks (Hong and Kacperczyk (2009)). *Ceteris paribus*, these sectors are "controversial" and riskier – i.e., face higher risk of litigation or regulatory fines – and thus have lower average ESG ratings. Consistent with this logic, we find that moral funds have higher overall ESG ratings, and this is driven in part by sector tilt towards high average-ESG rating sectors. We also document that they significantly underweight sectors often subject to morally-based exclusion, such as Energy, Defense, Tobacco, and Utilities.

We hypothesize that impact funds seek to generate impact by engaging with company management and actively voting in support of Environmental and Social (ES) shareholder proposals. Consistent with this, we find that impact funds are more likely to vote in support of ES shareholder proposals, less likely to vote against them, and less likely to abstain or decline to vote. In contrast, financial or moral funds are not more engaged with ES proposals. Finally, we

find that impact funds improve the ESG performance of the companies they invest in during the investment holding period, relative to other sustainable funds.

Our findings shed light on why sustainable investing is often found not to be effective in generating impact: most funds labeled “sustainable” in fact do not prioritize generating impact, and instead seek to enhance financial value and/or adhere to categorical morality. Hybrid funds are common and funds combining financial and moral goals manage more assets under management than any other category. These findings illuminate the need for clear classification labels to differentiate between 'doing well and being good' versus 'doing good'. By clarifying fund objectives, our methodology enhances investor welfare by facilitating informed decision-making tailored to individual preferences. Understanding the heterogeneous underlying preferences of sustainable investors and the potential tension between them can in turn inform policy decisions, investment practices, and corporate behavior.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 presents the model and hypotheses. Section 4 describes the data. Section 5 presents our results. Section 6 concludes.

2. Related Literature

This paper relates to three strands of the literature. First, a growing theoretical and empirical literature incorporates the possibility of having investors with nonpecuniary preferences in the economy and analyzes their impact on the financial and non-financial outcomes in the equilibrium.²

Several theoretical papers examine when for-profit firms with dual objectives of both financial return and social or environmental impact emerge as a contractual solution to the problem of private provision of public goods. Morgan and Tumlinson (2019) assume all investors derive utility from public goods and show that in the absence of managerial agency cost and private benefits, corporate social responsibility emerges as a more efficient, centralized giving mechanism that solves the free-rider problem of decentralized contribution. In contrast, Chowdhry, Davies, and Waters (2019) assume only a subset of investors derive utility from

² See Starks (2023), footnote 9 for theory papers. Also see Yasuda (2023) for the literature that relates to impact investing.

social goods (prosocial investors) and analyze when their investment alongside profit-only investors improves social outcomes. When optimal, joint financing by impact and non-impact investors is mutually beneficial because impact investors are willing to give non-impact investors a subsidy in exchange for production of social goods. Oehmke and Opp (2022) similarly assume socially responsible investors who internalize the social impact of externalities generated by firms and financial investors who disregard externalities and only seek financial returns.

Further, Oehmke and Opp (2022) bring to sharper focus the disjoint nature of impact vs. moral investors' motivations: "Because avoided externalities matter, ... investments in sin industries ... can be consistent with a socially responsible investment mandate. In contrast, it is efficient to not invest in firms that are already committed to clean production ... because clean production will occur regardless of investment by socially responsible investors." (p.5) As a real-world example, Engine No. 1, an impact investor, invested in Exxon Mobil in order to install new board directors and shift its corporate strategy towards faster transition out of fossil fuel into renewable energy. The fund aims at reducing the negative externalities generated by a major oil company which, even after implementation of the recommended changes, would still produce large quantities of negative externalities in the form of greenhouse gas. In contrast, many moral funds avoid investing in fossil fuel industries and instead have positive tilts toward information technology companies such as Alphabet (Google parent) and Microsoft, which have small carbon footprint to begin with and will achieve net zero emissions with or without investments by moral funds. Moral funds' behavior is in line with deontologically-moral investors' motivation but not with that of impact (consequentialist) investors. Hartzmark and Shue (2023) document this "paradox" that many ESG funds invest in already-green firms, which generates little additional impact.³ We argue that this is because many sustainable funds seek either financial value or categorical moral goals, not an impact goal, and the lack of clarity in their goals both exacerbates the seeming paradox and potential misallocation of capital (to the extent the end investor goals and fund goals are mismatched).

While many existing papers consider the presence of investors with nonpecuniary preferences, most studies assume that all prosocial investors have one type of nonpecuniary utility. Moreover, existing studies often do not specify whether the nonpecuniary preferences

³ Atta-Darkua, Glossner, Krueger and Matos (2023) find that institutional investors that join climate pledges green their portfolios rather than holding high-emission firms and engaging with them to reduce their emissions.

assumed in the model imply deontological (categorical) morality or consequentialist (impact) morality. For example, in Pastor, Stambaugh and Taylor (2021), investors derive utility from holding green stocks and disutility from holding brown stocks. Green stocks generate positive externality whereas brown stocks generate negative externality. While firms' nonfinancial characteristics are described in terms of externalities, the form of utility assumed in the model in Pastor, Stambaugh and Taylor (2021) is consistent with categorical morality, where investors' utility increases by eliminating brown stocks from their portfolio (and concentrating in green stocks), *ceteris paribus*.⁴

Berk and Binsbergen (2022) assume that a subset of investors in the economy “will only hold clean stocks”. These investors are categorically moral (deontological) investors who impose *ex ante* ethical purity criteria on their portfolios and the paper studies the impact of deontological, categorical moral investing. In contrast, in the papers listed above (Morgan and Tumlinson (2019), Chowdhry, Davies, and Waters (2019), and Oehmke and Opp (2022)), investors are modeled to have consequentialist (impact) nonpecuniary preferences.

Our paper contributes to this literature in two ways. First, we argue that distinguishing between the two, often mutually exclusive forms of nonpecuniary moral preferences – categorical morality and impact – is important in advancing our knowledge of the role and effectiveness of sustainable investing.⁵ Second, we develop an empirical method that distinguishes whether a fund seeks a categorical moral or impact goal and document that the two goals are associated with divergent fund management practices. We also document that categorical moral funds are more common and manage more capital than impact funds in the public equity market.

Second, an increasing number of studies in the finance and economics literature adopt machine learning models in their methodological design.⁶ Bybee et al. (2021) estimate a topic model on 800,000 Wall Street Journal articles to summarize business news and measure the state

⁴ Similarly, Lo and Zhang (2021) model investors to have deontological (categorical moral) rather than consequentialist (impact) preferences. In contrast, in Landier and Lovo (2020) investors are impact investors.

⁵ Bonnefon et al. (2022) test for prevalence of deontological and consequentialist moral preferences in an experimental setting. Shanker (2023) models the effect of heterogeneous (deontological and consequentialist) moral preferences of shareholders on nonfinancial firm outcomes (e.g., emissions) in equilibrium. Also see Green and Roth (2020), Heeb et al. (2023), and Lee et al. (2020).

⁶ See, for example, Bingler et al. (2022), Rajan et al. (2022), Acikalin et al. (2022), Michaely et al. (2023), and Duchin et al. (2022). Gentzkow et al. (2017) provides a survey on use of textual data in finance and economics research, and Goldstein et al. (2021) provide an overview of the “big data” research in finance.

of the economy. Abis (2022) and Abis and Lines (2023) apply unsupervised topics models to classify mutual funds based on their overall investment strategies. Kai et al. (2021) construct a corporate culture dictionary using machine learning techniques, the word embedding model.

Our paper is the first to apply supervised machine learning methods to classify investors based on their stated sustainability goals. Topics models are good at classification when the key objectives that we wish to extract from text are nouns (= topics). In contrast, BERT and other supervised models have advantages when we wish to extract agents' intentions, i.e., what they wish to do with topic A. We contribute to the literature by demonstrating the ability of supervised natural language processing techniques to interpret agents' goals in sophisticated, reinforced learning setting and thus aid human researchers to scale up and speed up research on non-quantitative dimensions.

Third, this paper relates to the literature that studies the investor motivations for sustainable investing. Riedl and Smeets (2017) find that investors with prosocial preferences are more likely to invest in sustainable mutual funds. Bauer, Ruof, and Smeets (2021) find that Dutch retirees support allocating more of their retirement portfolios to sustainable investments even when they expect financial returns to be lower. Hartzmark and Sussman (2019) study the introduction of Morningstar Globe Ratings and find that funds categorized as low (high) sustainability resulted in net outflows (inflows). Results are consistent with either (i) investors expect high sustainability funds to outperform (though they do not) or (ii) investors derive nonpecuniary utility from holding high sustainability funds. Barber, Morse and Yasuda (2021) study venture capital impact funds and find that investors' willingness to pay for impact varies considerably across legal and regulatory environments, investor geography, and time. Diaz-Rainey et al. (2023) study shareholder proposals on climate change issues and find that while targeted firms' environmental performance rating improves significantly afterwards, emissions do not change appreciably.

As Starks (2023) emphasizes, some sustainable investors are financially motivated and use ESG characteristics of firms as financially material information to enhance financial value of their investments. Other sustainable investors are motivated by their nonpecuniary preferences with further distinction between categorical morality and impact, as discussed above. Most existing studies do not recognize this heterogeneity of sustainable investors and instead treat all sustainable funds as a homogenous fund type and investors channeling capital into sustainable

funds as also having a homogenous motive.⁷ Thus, the reported empirical results are effectively averages of a mix of investor motives and a mix of funds catering to different investor objectives.

Our paper contributes to the literature by explicitly recognizing that sustainable mutual funds are heterogeneous in their sustainability goals and classifying them into three distinct types, and then documenting how the stated goals correlate with what they do as fund managers. The method we develop help disentangle the confusion among the distinct sustainability goals and quantify the relative sizes of the three sustainable fund types in the industry. We find that financial value-oriented sustainable funds are the most common and largest in terms of the assets under management, followed by categorical moral funds, and impact funds are the least common. This raises a new question of whether the current relative sizes of the three fund types represent the relative sizes of the three preference types of end investors (households), or they reflect mismatch and misallocation due to the current confusion and lack of clarity about the funds' sustainability goals.⁸

3. Model and Hypotheses

In this section we describe our research design. First, we present our classification of investors' sustainability goals based on two dimensions – (i) whether non-financial (ESG) characteristics of the investment enters the investor's objective (utility) function directly, and (ii) if yes, whether the ESG characteristics is used to align the investments with the investors' moral values ex ante, or to measure the value of the impact generated by the investment ex post. Based on this classification system we define three investor goals – financial, moral judgment, and impact. Second, we present an objective empirical method of classifying an investor's sustainability goals based on the text description of their investment strategy. Since the three investor goals are distinguished from each other in how real (non-financial) outcomes are valued by investors, not what topics or data points the investors track, it is critical that the method can extract the

⁷ Pedersen, Fitzgibbons, and Pomorski (2021) models financially-motivated sustainable investors (ESG-aware investors) and investors who derive positive utility from holding high ESG-rated stocks (ESG-motivated investors). While their model also includes ESG-unaware traditional investors who ignore ESG information, Goldstein, Kopytov, Shen, and Xiang (2022) assume that all traditional financial investors are aware of ESG risk and differ from prosocial investors only in their preferences. Neither paper incorporates impact investors in the models.

⁸ See Giglio et al. (2023) for survey results of retail investors, Krueger et al. (2020) for survey results of institutional investors, and Bonnefon et al. (2022) for an experimental study.

investor's intent, not just keywords they use. Third, we present a machine-learning model that operationalizes this method and describe how we apply this model to classify sustainability goals of all U.S. actively-managed U.S. mutual funds and ETFs. Finally, we hypothesize how mutual funds pursuing each of the three sustainability goals would differ in their fund management activities.

3.1 Classification of Investor's Sustainability Goals

We classify investors' sustainability goals based on two dimensions in Figure 1. The first dimension is whether non-financial characteristics of the investment enters the investor's objective (utility) function directly. This dimension separates what Starks (2023) calls "Value" investing from "Values" investing. "Value" investing is motivated by using ESG information to identify and manage/pursue financial risk and opportunities, whereas "Values" investing is motivated by how the ESG characteristics of the investments impacts the investors' utility because of their nonpecuniary preferences. For "Value" investors, ESG characteristics matters if and only if it is financially material; for "Values" investors, ESG characteristics matters in and of itself.

This dimension also separates what Pedersen, Fitzgibbons and Pomorski (2021) calls "ESG-aware" investors from "ESG-motivated" investors. "ESG-aware" investors are mean-variance utility maximizers, just like traditional financial investors, but they "use assets' ESG scores to update their views on risk and expected return." In other words, they pursue the same purely financial objective as "ESG-unaware" traditional investors and use the ESG information to "do well". The focus is often on identifying and managing various ESG-related risks: "For example, ESG considerations can shed light on financially significant risks including environmental, reputation, human capital, litigation, regulatory, corruption, and climate risk, among others." (Starks (2023)). We label this investor goal as "financial value".

We further classify the investors with nonpecuniary preferences into two distinct types using the second dimension. "Values" investors derive utility from "doing good", and the second dimension captures how the investors as an actor measures the "goodness" of their investment decisions. We draw on the long-standing philosophical concepts of "Consequentialism" and "Deontology". In consequentialism, the "goodness" of an action or "... normative properties depend only on consequences. ... What is best or right is whatever makes the world best in the

future (Stanford Encyclopedia of Philosophy).⁹” Thus, to be a consequentialist investor, one has to measure the non-financial outcome of the investments (e.g., emission reduction, reduced gender pay gaps) to judge its success. This requires ex post monitoring and engagement with the portfolio companies and thus is inherently activist in nature. We label this investor goal as “impact investing”.

In contrast, in deontology, the morality of an action should be based on whether that action itself is right or wrong under a series of rules and principles, rather than based on the consequences of the action. For example, if you are against violence, then investing in a weapon manufacturer is inherently “wrong” and excluding it from your portfolio is inherently “right”, regardless of whether your exclusion decision results in less violence in the world or not. Note that the deontological concept of moral responsibility is consistent with categorical negative screening as an operationalized investment strategy. Socially-responsible investment (SRI) funds are the earliest type of sustainable fund vehicles and they predominantly employed ethical negative screening methods. Many of these SRI funds have religious orientation and offer investors an investment vehicle that is “aligned with their moral values”. We label this investor goal as “categorical morality”.

We intentionally distinguish between “impact” and “categorical morality” for a couple of reasons. First, though both concern non-pecuniary preferences of “values” investors, we argue that the ways in which sustainable funds operationalize pursuit of these objectives are diametrically opposite of one another. For example, let’s suppose that two investors are concerned about the climate change and one (Inger) is a consequentialist while the other (Sam) is a deontologist. Inger derives positive utility from investing in a polluting company (e.g., Exxon Mobil) and forcing it to reduce its GHG emissions, i.e., turning it greener. The positive externality that her investment generates makes her happy. Sam, on the other hand, believes giving capital to a company that currently derives much of its profit from fossil-fuel extraction is morally wrong, and to avoid this disutility (guilt) he excludes Exxon-Mobil from his investment portfolio. Both are nonpecuniary utility maximizers but actions that maximize their utility are mutually exclusive, and a fund cannot cater to both clienteles.

Second, not distinguishing between these goals can lead to misleading conclusions about whether the stated intent and actions are matched or not at the fund level, especially when the

⁹ Available at <https://plato.stanford.edu>.

fund’s ESG performance is measured using ESG ratings. Using the above example, a consequentialist fund should not be averse to buying a brown firm stock, whereas a deontological fund should categorically avoid investing in a brown sector. Existing studies often assume a priori that investors with nonpecuniary preferences derive values from holding green/high-ESG rating stocks and use ESG ratings as measures of “greenness” or “goodness” of a stock.¹⁰ We show that this simplification departs from practice both because the two nonpecuniary preferences map to distinctly different portfolio choices, and because the ESG ratings are not designed to capture “goodness” of a stock in ways that map to either impact or moral investors’ definition of goodness. Instead, ESG ratings – specifically, MSCI ESG ratings that we analyze – are designed to help the financially-motivated investors to identify risks and opportunities. Thus, a financial fund may systematically buy the least-brown firm (as per MSCI) within a given sector (which may be brown on average) and hold a diversified portfolio to best hedge the stranded-asset risk while not miss out on profitable opportunities. To the best of our knowledge, our paper is the first academic work to make this distinction and propose an objective method to classify and analyze the three sustainable fund types separately.

3.2 An Empirical Method of Classifying Investors’ Sustainability Goals

Currently there are no formal regulations in the U.S. that identify sustainable funds or distinguish between the three sustainable fund types.¹¹ 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.

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

¹⁰ For example, Pedersen et al. (2021) define ESG-motivated (values investors) as “hav[ing] preferences for high ESG scores”.

¹¹ See Hardy, Lambert, Yang and Yasuda (2023) for a study of EU’s Sustainable Finance Disclosure Regulation (SFDR) that requires sustainable funds to declare their funds as either Article 8 or Article 9 funds starting in 2021.

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. For example, keywords such as "labor relations" can be used by financial/ESG-aware investors to describe how they mitigate litigation risk from potential labor disputes by investing in companies with high S ratings for labor relations or used by impact investors to describe how they push portfolio companies to improve their family leave policy. 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 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.

To train the model, we need to build a sample of example sentences that sufficiently and distinctly express each of the three sustainable investing goals. 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 mentions financial performance as the motivation for monitoring externality, then such a sentence fails to meet the "impact investing" criteria. More details on the classification criteria can be found in the online appendix.

Once we have built a training sample, the BERT model is trained on the data (except for a subset reserved as a testing sample) to learn the types of sentences that qualify as "financial value", "categorical morality", and "impact investing". We leave out a subset of the manually coded sample as the "testing sample", and now ask the model to classify sentences in the testing sample into "financial", "moral", and "impact". Because we have the "true" answers for these

manually coded testing sample, we can calculate the accuracy, precision, and recall performance of the BERT model in predicting the classification correctly. The key is to provide the model with a sufficiently rich training data with sentence and phrase variations for each investment type so that the model can detect all relevant sentence types associated with the category.

Once the model can successfully mimic our human detection of “financial”, “moral”, and “impact” intent in sentences, we can then use the model to interpret any texts describing an investment strategy and predict if it pursues sustainability goals and if so, which of the three types.

3.3 A Machine-Learning Model Application to U.S. Mutual Funds

To build a training sample, we start with a list of U.S. sustainable mutual funds compiled and published by Morningstar. Morningstar started compiling the annual list in 2018, and we obtained the lists for 2018, 2019, 2020, and 2022 either directly from the company website or from other researchers who have used the list.¹² We chose this set of U.S. mutual funds deemed to be pursuing sustainability goals by Morningstar 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.

There are currently no mandatory sustainability disclosure requirements in the U.S. for mutual funds. After reading fund prospectuses of funds on the Morningstar fund list, we learned that funds typically discussed their sustainability goals in the “Principal Investment Strategy” section of the prospectus. This is a required section of the fund prospectus and is supposed to “tell you how the fund intends to achieve its investment objective” (US SEC 2016). 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. The list of ESG keywords we used for this screening step is provided in the Appendix A1.

There is a priori no fixed sample size that is appropriate to be used as a training sample. In selecting the fund-years to be included in the training sample, we opted to use a given fund only once even if the fund is selected by Morningstar for multiple years. This is because a fund

¹² We contacted Morningstar and requested the 2021 list, but it declined to provide us with the list and provided no reasons.

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. We also aimed to have broad representation of fund families in the training sample (though fund families were not included if none of their funds were selected by Morningstar to be on its sustainability fund list). All together, we used 336 fund-year observations and identified 2,834 ESG-related sentences for manual classification. These 2,834 sentences were extracted from the Principal Investment Strategy” section of the fund prospectus and contained at least 1 ESG-related keywords. They were then manually classified by the authors as “financial”, “moral”, “impact”, or “unclassified”.

We labeled a total of 668 sentences as “financial”, 181 sentences as “moral”, and 236 sentences as “impact”. Note that the majority (1,749) 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. These sentences are retained in the training sample because for the model to detect sustainability goals correctly with a low rate of false positives, it is useful to have these “fluffy” or empty ESG sentences as “unclassified”. Table [x] in the Appendix includes example sentences for each of the 4 classification categories that we manually coded.

These manually coded sentences are then split into a training and testing sample to report the BERT model performance. The model performance is reported in Section 5.1. Once the BERT model performance is confirmed, then we apply the model to classify both (i) the remaining fund-year observations of funds on the Morningstar sustainable fund list and (ii) the universe of U.S. mutual funds and ETFs from 2018-2022. The goals of this exercise is twofold. First, we want to know which of the 3 sustainability goals that the Morningstar sustainability list funds tend to pursue. Second, we want to know what percentage of funds that are not on the Morningstar list say they pursue sustainability goals, and which of the three goals is more prevalent among this set of funds (“missing sustainability funds”). These classification results are reported in Section 5.2.

3.4 Hypotheses on Fund Management Styles of the Three Fund Types

Having built a scalable method to classify all mutual funds by their sustainability goals, in this section we describe our hypotheses on how each sustainable fund type manages their funds in order to operationalize their sustainability goals.

Our first overarching hypothesis is that funds act according to their stated sustainability goals. On the one hand, this is plausible to the extent that fund managers aim to differentiate their funds from their competition by appealing to investors whose own sustainability goals match those of the funds'. In order to attract and retain investors' assets and to build their reputation as a sustainable fund manager, fund managers are incentivized to act consistent with their stated goals. On the other hand, it is also possible that in the absence of mandated and regulated disclosure standards, fund managers engage in cheap talk and their actions bear little resemblance to their stated goals. Thus, this is an open empirical question.

Hypothesis 1: Fund managers manage their funds in ways that are consistent with their stated sustainability goals.

There are two ESG or sustainability aspects of fund managers' actions that we analyze in this paper. First is the MSCI ESG ratings of the stocks they hold in their fund portfolios. MSCI ESG ratings are good markers of sustainable fund manager style, because they are designed specifically to cater to financially motivated ESG investors, as described below, and this helps us generate differentiated predictions on portfolio composition favored by each of the three investor types. Second, we analyze the fund managers' voting records for shareholder proposals that focus on social and environmental issues.

MSCI is a global provider of equity and fixed income indexes as well as ESG and climate products. On its website, it states that "MSCI ESG Ratings aim to measure a company's management of *financially relevant ESG risks and opportunities*. We use a rules-based methodology to identify industry leaders and laggards according to their exposure to ESG risks and how well they manage those risks relative to peers" (emphasis added by the authors) (<https://www.msci.com/our-solutions/esg-investing/esg-ratings>) Two things are notable. First, while it explicitly mentions financial risks and opportunities, there is no mention of alignment with moral values or the impact of company activities on the society or the environment. Thus,

MSCI ESG ratings aim to fit the needs of financially motivated sustainable investors, as opposed to those of morally motivated investors or impact investors.

Second, it also emphasizes that MSCI ratings are industry specific and thus not comparable across industries. On the same website, it states that “ESG risks and opportunities can *vary by industry* and company. Our MSCI ESG Ratings model identifies the ESG risks, (what we call Key Issues), that are *most material to a GICS® sub-industry or sector.*”

In support of our interpretations, on another webpage MSCI states that “[t]hough ESG investing has been referred to in many different ways, ... we group [them] into three approaches that investors use to achieve distinct ESG objectives: ESG integration, impact investing and values-based investing. Our ESG ratings are designed specifically for ESG integration, which uses ratings to support the building of a resilient portfolio for the specific purpose of enhancing long-term risk-adjusted returns.”¹³ What MSCI calls ESG integration matches our financially motivated sustainable investing, and values-based investing matches our morally motivated sustainable investing. Further, on the same webpage MSCI states that “Our ESG ratings assess how well companies manage risk compared with their peers, not across industries”.

Finally, MSCI ESG ratings have both “risks” and “opportunities” factors. High scores on risk-related key issues indicate that firms manage their ESG risk well, whereas high scores on opportunities-related key issues indicate that firms position themselves well to take advantage of the opportunities. Moreover, MSCI determines which risk and opportunity issues are financially material for a given industry sector, thus resulting in different average ESG performance across industries.

These features of MSCI ESG ratings help us generate predictions regarding how portfolio composition differs among the three fund types. For example, financially motivated sustainable funds favor companies with high MSCI ESG ratings relative to their industry peers. These are what MSCI recommends as “leaders” – i.e., given a set of material ESG risks and opportunities that are inherent and specific to the industry, these companies do better than their competitors in mitigating ESG risks efficiently and leveraging ESG opportunities, thus most likely to enhance the financial performance of the portfolio.

¹³ <https://www.msci.com/our-solutions/esg-investing/esg-ratings/what-esg-ratings-are-and-are-not>. See also Figure 2.

Hypothesis 2: Financially motivated sustainability funds hold stocks with higher ESG ratings relative to industry peers than other sustainability funds.

Moral motivation of sustainable investors does not render themselves to be sensitive to MSCI ESG ratings of a stock relative to its industry peers since it is not concerned about using ESG information to enhance the financial value of the portfolio. So we posit that, conditional on holding a stock in a given industry, morally motivated sustainable funds do not exhibit a significant tilt towards industry leaders. Rather, morally motivated sustainable funds categorically exclude certain sectors (typically listed in the fund prospectus) that they deem morally misaligned. A priori, it is not clear whether those categorically excluded industries tend to have significantly higher or lower MSCI ESG ratings than included industries on average. On the one hand, moral objection may be positively correlated with the frequency or magnitudes of ESG-related controversies that MSCI deems material for those industries. If so, then that implies that industries that morally motivated funds include (do not exclude) have higher industry average ESG ratings than financially motivated or impact funds. On the other hand, moral objection could be uncorrelated with financially material controversy in some cases, weakening the connection. For example, some religiously motivated funds object to and exclude abortion-related industries, but MSCI and financially motivated investors may consider this to be financially immaterial. Thus, we expect either a weak or no correlation.

Note that, although exclusion and divestment sound similar, categorical exclusion policy of morally motivated sustainable funds is not compatible with the divestment strategy of using the threat of exit to pressure companies for a change. In order to threaten to exit, a fund has to initially buy a controversial stock.

Hypothesis 3a: Morally motivated sustainability funds are less sensitive to within-sector relative ESG ratings than financially motivated funds.

Hypothesis 3b: Morally motivated sustainability funds hold stocks in less controversial industries (higher industry average ESG ratings) than other sustainability funds.

Impact motivation of sustainable investors incentivizes them to hold stocks where there is room for additional externality generation. There are two possibilities. First, impact investors may prioritize investing in specialist companies whose business model itself is tied to positive externality generation, e.g., innovation in battery efficiency. Impact funds organized as VC funds exclusively engage in this type of impact investing (Barber, Morse, and Yasuda (2021)). Second, impact investors may invest in generalist companies and pressure the management to improve the social and environmental impact of the company practice. To detect the first possibility, we exploit the disaggregated environmental rating components of the MSCI data that isolate the firm's strengths in "cleantech". We posit that impact funds have a positive tilt towards companies that score high on the "cleantech" component of its E ratings because such companies tend to specialize in products or services where positive externality generation is baked into the business model itself.

How the second type of impact investing is reflected in the average MSCI ESG ratings of the fund portfolio is more ambiguous. On the one hand, to the extent that improvement in ESG ratings within an industry is positively correlated with improved social and environmental impact of company activities, we expect that impact sustainable funds hold stocks with lower MSCI ESG ratings than other sustainable funds. On the other hand, precisely because MSCI ESG ratings measure resilience of a company's financial value with respect to its ESG risk, and not how positive the social or environmental impact of the company's activities on the society, there may not be any reliable correlation between the two. Thus, we expect either a weak or no correlation vis-à-vis traditional, non-sustainable funds. But relative to financially motivated sustainable funds, we expect that impact funds hold lower ESG rating stocks.

Hypothesis 4a: Impact sustainability funds hold stocks with lower ESG ratings relative to industry peers than other sustainability funds.

Hypothesis 4b: Impact sustainability funds overweigh stocks where externality generation is baked into the business model (e.g., cleantech).

Beyond ESG ratings of portfolio stocks, we also posit that impact funds are more activist than the other two fund types. This is because impact investing is inherently activist (Yasuda

(2023)) in order to generate positive outcomes after the investment is made. In contrast, morally motivated investing is inherently passive since the goal of being morally aligned is achieved as soon as the fund excludes objectionable investments from its portfolio. To pressure the portfolio companies to adopt company practices that generates more positive social or environmental impact, funds with stronger focus on impact are more likely to support social or environmental shareholder proposals. Financially motivated funds may actively vote to enhance the financial value of their holdings, but financial incentives alone do not predict any direction on which ways they vote on social or environmental issues.

Hypothesis 4c: Fund with stronger focus on impact are more likely to vote in support of (less likely to vote against) social or environmental shareholder proposals.

4. Data

Our study combines several data sources. For the creation of training and testing samples for our BERT model, we begin with the list of Morningstar Sustainable Funds for the years 2018, 2019, 2020, and 2022.¹⁴ 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. When a fund is selected for multiple years, we download its prospectus for each year in which it is selected, resulting in a total of 1,361 fund-year observations.

We manually code 336 of the 1,361 fund-year observations to use as the training and testing sample for the BERT model. We then apply the trained and tested BERT model to all of the 1,361 fund-year observations and generate both indicator and continuous variables for each fund-year observation according to how many of a fund's ESG-related sentences are coded as financial, moral, or impact. We are currently collecting the prospectuses of all other U.S. actively managed mutual funds and applying the BERT model to them.

Next, we construct fund-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 calendar year, which we use as the portfolio weight of the security. From the MSCI ESG ratings database, we obtain

¹⁴ We contacted Morningstar and requested the 2021 list, but it declined to provide us with the list.

MSCI's ESG ratings for stocks held by our sample mutual funds. Combining the two, we calculate the fund-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.

Finally, from the ISS database, we obtain the fund's voting records on shareholder proposals at each of the stocks it held. We calculate the fund-level ESG activism variable defined as the likelihood of supporting shareholder proposals on environmental and social issues relative to the likelihood of supporting all shareholder proposals.

5. Empirical Results

5.1 BERT Classification Model Results

Table 1, Panel A presents the performance of the BERT classification model at the sentence level in the testing sample. Out of the 336 fund-year observations, we split them into a training sample of 265 fund-year observations with 2,277 ESG-related sentences and a testing sample of 71 fund-year observations with 557 ESG-related sentences. We trained the BERT model on the 2,277 sentences from the training sample for classification. The trained BERT model can be used to classify any text into the four ESG categories. We apply the trained BERT model to classify the 557 ESG-related sentences from the testing sample. For each ESG sentence in the testing sample, we obtain the ESG classification given by the BERT model. 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 three different model performance metrics: accuracy, precision, and recall.

[Insert Table 1: BERT Model Classification 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. We computed the model performance metrics for all classes and separately for each label. For example, when calculating the BERT model's performance in predicting Financial sentences, we set all non-Financial predictions and true values as zero and calculated the three metrics only for the true values and predicted values related to Financial.

We find that the BERT model performs well in classifying the sustainability goal expressed in a sentence. Comparing the three goals, accuracy ranges from 87% (Financial) to 99% (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 81% (Financial) to 90% (Impact), meaning when a sentence is classified as a given type (positive cases), most of them are true positives. Recall ranges from 66% (Impact) to 91% (Moral), meaning most of the true positive cases are classified as such. These performance levels are in line with other studies using BERT models in the literature (e.g., see Bingler et al. (2022) and Rajan et al. (2022)). Also note that our goal is classification at the fund level rather than at the sentence level. When aggregated at the fund level, our BERT model performs at even higher levels, as we show below.

Table 1, Panel B presents the performance of the BERT model at the fund level. The dummy variable $Financial_i$ is equal to 1 if the prospectus for fund i contains at least one sentence that is classified as Financial. The dummy variable $Moral_i$ and $Impact_i$ are analogously defined. These categorical variables defined at the fund are not mutually exclusive, since a fund prospectus can contain a sentence classified as Financial as well as a sentence classified as Moral, etc. Finally, the dummy variable $Unclassified_i$ is equal to 1 if the prospectus for fund i does not contain any sentences that are classified as one of the three sustainability goals.

Comparing the three goals, accuracy ranges from 91% (Impact) to 96% (Moral), implying that when a fund is classified as either a given type (positive) or not (negative), it is correct in most cases. Precision ranges from 80% (Impact) to 97% (Financial), meaning when a fund is classified as a given type (positive cases), most of them are true positives. Recall ranges from 80% (Impact) to 100% (Moral), meaning most of the true positive cases are classified as such. For each of the sustainability goal types, the overall performance is improved at the fund level compared to the sentence level. In particular, the recall metric substantially improves for all three fund types.

In addition to the discrete (categorical) variable, we also construct continuous variables to capture the intensity or focus with which a fund pursues a given sustainability goal, while allowing a fund to pursue mixed/multiple goals. The continuous variable $Financial\ Ratio_i$ is the total number of Financial sentences divided by the total number of sentences classified as one of the three sustainability goal types in the prospectus of fund i . $Moral\ Ratio_i$ and $Impact\ Ratio_i$ are analogously defined. Note that these three ratio variables sum up to 1 for a given fund and

measure the relative intensity or focus with which a fund pursues a given sustainability goal. Comparing the fund-level sustainability goal ratios “Financial Ratio”, “Moral Ratio”, and “Impact Ratio” between the BERT model prediction and the manual coding, we find that the two are highly correlated, with the correlation of 0.942, 0.988, and 0.915 for Financial, Moral, and Impact, respectively. This suggests that the BERT model is able to reliably detect not only the presence of a sustainability goal in the fund’s stated goals but also the strength of focus on that goal relative to other, potentially competing, sustainability goals.

Finally, the continuous variable *Unclassified Ratio_i* is the total number of sentences that are not classified as any of the three sustainability goals divided by the total number of ESG-related sentences for fund *i*. This can be thought of as a measure of opacity or ambiguity in the fund’s sustainability goals.

Table 2 presents the summary statistics for our key variables. Panel A presents the results for all Morningstar funds. Panel B presents the results for a subset of Morningstar funds that (i) can be merged with the CRSP mutual fund database and (ii) hold at least one stock whose ESG rating is provided by MSCI. This conditioning is necessary for most of our analysis that rely on value weighted MSCI ESG ratings of the fund portfolios. The summary statistics are qualitatively similar in the two panels. Note that the funds’ sustainability goal variables are constructed by applying our BERT-model classifications to all Morningstar funds including those funds in the training sample. If fund *i* appears on the Morningstar sustainable fund list multiple times, we analyze the fund prospectus in each year as a new observation. In the future version of the paper, we will expand our main analysis beyond the Morningstar funds to include the universe of actively managed mutual funds in the U.S.

[Insert Table 2: Summary Statistics]

In Panel A (B), the Financial Dummy variable’s mean is 54% (54%). This indicates that the majority of Morningstar sustainable funds state they seek financial value, i.e., they incorporate ESG information to enhance the fund’s financial return and/or to manage risk. In contrast, the Moral Dummy variable’s mean is 39% (42%), while the Impact Dummy variable’s mean is 33% (29%). Note that some funds may state they pursue more than one of the three goals, while other funds may not state they pursue any of the goals. Of the 1,361 fund-year observations in Panel A, 1,090 are associated with at least one sustainability goal, while 271 are

not associated with any of the three goals. Turning to the Unclassified Ratio variable, we see that on average 63% of the ESG-related sentences associated with a fund do not meet the BERT-model criteria for a clear classification, i.e., deemed too unclear/ambiguous.

In Panel B, we report the composition of the portfolio holdings for the 892 funds that are merged with the CRSP mutual fund database and hold at least one stock whose ESG rating is provided by MSCI. About 62% of the fund portfolio (in dollar values) consist of the MSCI-rated stocks, while the other 32% consist of stocks not rated by MSCI, and 5% consist of other asset types (e.g., cash, fixed income, derivatives, other illiquid assets). MSCI-rated stocks are equal to about 66% of all stock holdings on average. Stocks in the funds' portfolios have been held for 6.4 quarters on average.

Among the classified sentences of funds associated with one or more sustainability goals (about 80% of the Morningstar fund sample), the breakdown of the fund goal focus is 51% (51%) Financial, 22% (25%) Moral, and 27% (24%) Impact in the Panel A (B) sample funds, respectively. It is interesting that the most common sustainability goal is financially driven, in contrast to the popular perception that sustainable investing is mission-driven. When the funds focus on non-pecuniary sustainability goals, the focus is split equally between Categorical Morality via exclusion and Impact Generation.

Panel C of Table 2 presents the correlation between the Financial, Moral, and Impact Dummy variables. Having a Financial sustainability goal is significantly positively correlated with having a Categorical Moral sustainability goal, with the correlation coefficient of +0.195. In contrast, having a Financial goal is significantly negatively correlated with having an Impact goal, with the correlation coefficient of -0.122. Finally, there is no correlation between having a Moral goal and an Impact goal.

Panel A of Figure 3 presents the Venn diagram of the three sustainability goals pursued by the Morningstar sustainable funds. "All Funds" include all years (2018, 2019, 2020, 2022), whereas the other figures include the funds on the list in a given calendar year. Consistent with Table 2, the Financial goal is the most commonly pursued sustainability goal, followed by Moral, with the Impact goal the least common. Hybrid funds (with overlaps between two or three circles) are also common, with 49% of the funds stating they seek more than one sustainability goal. Hybrid funds that pursue both the Financial and the Categorical Moral goals

are the most popular hybrid fund type, with 256 or 23% of all funds falling into this category. Furthermore, over time the popularity of this fund type has increased.

Panel B of Figure 3 presents similar Venn diagrams for the subsets of funds that meet our conditioning criteria, namely merging with the CRSP Mutual Funds database and further restricting to funds with at least one stock whose ESG rating is provided by MSCI. The results are qualitative similar. When we measure the Assets under Management (in \$B) for each fund type, we find that funds pursuing Financial or Moral goals (and the hybrid funds seeking both) are even more dominant compared to funds seeking an Impact goal in term of the capital they attract. Note that this sample is limited to the funds listed on the Morningstar sustainable fund list. We will extend the analysis to non-Morningstar funds in a future version of the paper.

5.2 ESG Rating Results

5.2.1 Baseline Model

Having classified the funds in terms of their sustainability goals using the BERT model, we now turn to examining the funds' portfolio ESG ratings and other characteristics.

Panel A of Table 3 presents the weighted-average ESG ratings of the fund's MSCI-rated stock holdings for funds associated with each of the three sustainability goals. Portfolio weights are used to weight-average the ESG ratings of the stocks held by each fund. Following the existing literature, we construct three variables capturing the E and S components of the ESG performance for firm j in year t . We exclude the governance component of the ESG ratings from our analysis because we think that the three sustainability goals are distinct from each other along E and S dimensions, while there are no clear a priori distinctions along the governance (G) rating dimension. The variable definitions are provided in Appendix A2. We then construct the weighted-average ESG performance measure for fund i in year t . To test if funds systematically hold stocks with different MSCI ESG ratings depending on the sustainability goals they pursue, we measure the difference in the average ESG rating between a given fund type and the remaining sustainable funds in the sample. Note that the fund classification in Panel A is done using the fund goal dummy variables, whereas in Panel B an analogous analysis is done using continuous ratio variables and a regression approach.

[Insert Table 3: Baseline Model]

We find that, consistent with Hypothesis 2, sustainable funds seeking a Financial goal systematically hold stocks with significantly higher MSCI ESG ratings. The results hold for both the Environmental and Social Key Issue ratings. Since MSCI ESG ratings are designed to measure the resilience of the company's valuation and long-term stock returns against financially material ESG-related risk, this suggests that financially motivated sustainable funds indeed use the stocks' ESG information to enhance the funds' financial performance.

Interestingly, we also find that funds seeking a Moral goal also hold stocks with significantly higher ESG ratings than other sustainable funds. Recall from Figure 3 that many of the Moral funds also seek a Financial goal. Thus, the Moral fund result may simply reflect the presence of hybrid funds seeking both a Moral and a Financial goal. We will further examine this using the continuous ratio variables in Panel B of Table 3.

In contrast to the first fund goal types, we find that funds seeking an Impact goal tend to hold stocks with lower ESG ratings. While the signs are consistently negative, the results are statistically significant, and the magnitude is much larger for the environmental key issue ratings than for social key issue ratings. Impact funds in Panel A also include pure impact funds and hybrid funds, which may moderate the results. Thus, we turn to the continuous variable-based analysis in Panel B.

Panel B, Table 3 examine the same question using the continuous ratio variables in a regression setting where the three measures of average fund ESG ratings are regressed on the continuous ratio variable indicating the intensity of the fund's focus on a given sustainability goal. Here, we find results that are consistent with Hypothesis 2, 3a, and 4a: Namely, the more intensely a fund pursues a Financial goal, the higher the ESG performance of the stocks held by the fund (2); A fund pursuing a Moral goal more intensely is not holding higher-ESG rated stocks than other funds; And finally, a fund pursuing an Impact fund more intensely is holding significantly lower ESG rated stocks than other funds. The results are also consistent with the negative correlation between the Financial and the Impact goals in Table 2, Panel C: Since impact-seeking funds and financial value-seeking funds want opposite ESG characteristics in their stock portfolios, pursuing both goals is difficult and therefore rare.

5.2.2 ESG Rating Decomposition

Next, we decompose the MSCI ESG ratings into eight environmental and social subcategories and repeat the comparison analysis in Table 3. For brevity we restrict the analysis to the *ESG Score* variable (see Appendix A for the variable definition). Table 4 and 5 presents the results.

[Insert Table 4: Environmental Rating Decomposition]

In Panel A, we confirm that Financially motivated sustainable funds hold stocks with higher ESG ratings in Climate Change (e.g., carbon emissions) and Natural Capital (e.g., water stress) subcategories. In contrast, they hold stocks with lower ESG ratings in “Environmental Opportunities” subcategory. This is interesting because the “Environmental Opportunities” subcategory measures the firm’s (costly) commitment to invest in areas such as clean tech, green buildings, and renewable energy, whereas the other subcategories measure the resilience of the firm’s financial value against the environmental risk, such as physical, regulatory, and transition risk. The results suggest that financially motivated funds prioritize holding companies that are good at managing their environmental risk, not companies that position themselves to invest in emerging technologies that contribute toward decarbonization and other environmental goals. The results are consistent with Hypothesis 2.

In Panel A, moral funds’ results are qualitatively similar to those of financial funds. But in Panel B, where we examine the intensity of the goal focus using the continuous ratio variables, we find that funds with higher categorical morality focus do not hold stocks with higher environmental ratings in the Climate Change or Natural Capital subcategories. They hold stocks with lower ratings in the Environmental opportunities. It is possible that the positive correlation between financial and moral goals drive some of the results for moral funds reported in Panel A.

In contrast to the financial and moral funds, impact-seeking funds behave in opposite ways: they hold stocks with high performance in the “Environmental Opportunities” subcategory, such as investment in clean technology, and low ESG ratings in areas of Climate Change and Natural Capital subcategories (see Panels A and B). The Environmental Opportunities results are consistent with Hypothesis 4b: Funds pursuing an Impact goal seek to generate impact in part by investing in firms for which impact generation is inherently baked into its core business model. For solar panel or EV manufacturers, for example, growth in the firm’s operations itself implies

impact generation via replacement of fossil-fuel based activities with renewable energy-based activities.

In Table 5, we present the social performance decomposition results. Similar to Table A, we find that financial funds hold stocks with higher performance in Human Capital (e.g., labor management), Product Liability, and Social Opportunities (e.g., access to healthcare and finance). In contrast, more-impact-focused funds hold stocks with lower social performance in Human Capital and Stakeholder Opposition (e.g., Community Relations) subcategories. The results are consistent with Hypothesis 4a in that impact funds selectively hold firms with more room for improvement in their social performance and engage with them to generate positive impact.

[Insert Table 5: Social Rating Decomposition]

5.3 Cross-Industry Allocation

Next, we analyze portfolio holding decisions of funds pursuing a Categorical Morality goal. These funds categorically exclude stocks that belong to industries that are considered morally unsuitable. To the extent that “sin”-fulness of industries is positively correlated with the financial materiality of controversies, those excluded industries could have lower MSCI ESG ratings on average. This then implies that Moral funds tend to concentrate their holdings in *industries* with higher average MSCI ESG ratings. Note that this portfolio strategy is distinct from the Financial funds’ strategy of buying stocks with high ESG ratings relative to their industry peers within each industry, and holding a diversified portfolio across all industries.

We operationalize this task in three steps. First, we calculate the average MSCI ESG score for each industry. Second, we replace the individual firm’s ESG score with the industry average ESG score of the industry that the firm belongs to for each stock held by the sample mutual funds. Third, we calculate the weighted average fund ESG performance measure but this time using the industry average ESG score instead of the firm’s actual ESG score. In other words, we isolate the fund’s cross-industry tilting decisions and suppresses its within-industry individual stock-picking decisions.

[Insert Table 6: Industry Tilting Decisions by Moral Funds]

Table 6 presents the results. Consistent with Hypothesis 3b, we find in Panel A that Funds seeking a Categorical Morality goal hold stocks belonging to industries with higher average ESG ratings than other sustainable funds. Results remain when we move hybrid funds in Panel B. In contrast, Financial funds do not tilt their holdings toward industries with higher average ESG ratings once hybrid funds are removed (see Panel B). We interpret this to be driven by Moral funds' aversion to "sin" or controversies. Impact funds, on the other hand, appear to tilt toward industries with lower average ESG ratings (see Panel C), which is the opposites of Moral funds' tendencies.

Which industries do Moral funds tend to tilt towards/away from? We compare the industry weights of Moral and other sustainable funds in five sector subcategories that are frequently subject to moral-based exclusion: (fossil-fuel) energy, defense (which is part of Aerospace), Casino, Tobacco, and Utilities (based on fossil-fuel energy). Table 7 presents the results.

[Insert Table 7: What Industries do Moral Funds Hold?]

In Panel A, we find that Moral Funds have lower portfolio weights for stocks in all selected sin sectors except Casino. Note that some moral funds are hybrid funds pursuing more than one sustainability goals. When we exclude hybrid funds (see Panel B), the portfolio weights on Casinos drop substantially, from 0.47% to 0.04%, and this is lower than the average sustainable funds, though the difference is not statistically different. For other four sectors, the results are qualitatively the same – i.g., moral funds place significantly lower weights in sin stocks. As a reference point, we report in Panel C the portfolio weights of these sectors among the S&P 500 firms. With the exception of utilities, portfolio weights of these sectors among sustainable fund holdings are substantially lower than in S&P500. Note that utilities can include power plants powered by renewable energy. Taken together, we find that moral funds significantly underweight those sectors that they say they morally object to.

5.4 Within-Industry Allocation

We now suppress the funds' cross-industry tilting decisions and isolate its within-industry individual stock-picking decisions. We operationalize this in two steps. First, we industry-adjust each firm's ESG rating by calculating its percentile ranking within the industry that the firm belongs to (normalized to scale from 0 to 10). Second, using this industry adjusted ESG rating, we calculate the fund's weighted average percentile rank. Third, we compare this for each of the three fund goals.

[Insert Table 8: ESG Rating Percentile Ranks Within Industry]

Table 8 presents the results. Consistent with Hypothesis 2, Financially motivated funds hold stocks with high ESG ratings relative to industry peers. Moral funds also hold stocks with high ESG ratings relative to industry peers; however, the magnitudes are smaller compared to those for the Financial funds. The results are qualitatively similar in Panel B.

In sharp contrast, we find that Impact funds hold stocks with significantly lower ESG ratings relative to the industry peers. Panel B confirms this result in a regression setting with the continuous Impact Ratio variable. The results are consistent with Hypothesis 4a again, suggesting that Impact funds seek out “laggards” firms within an industry, presumably since those firms have the most room for improvement in their environmental or social practice, and thus the most potential for impact generation.

5.5 Voting Analysis

We turn to the ISS mutual funds voting data to test whether funds pursuing an Impact goal vote differently from sustainable funds. Following He, Kahraman and Lowry (forthcoming), we construct the ES dummy variable for Environmental and Social shareholder proposals. Then we calculate the percentage of the ES shareholder proposals that fund i voted in support of in year t , and divide this by the percentage of the *all* shareholder proposals that fund i voted in support of in year t . A ratio greater than 1 indicates that fund i is more likely to support an ES shareholder proposal than other, non-ES shareholder proposals. We further construct probability ratios for “vote against”, “abstain from voting”, and “do not cast a vote” using analogous methods.

[Insert Table 9: Environmental and Social Shareholder Proposal Voting]

Table 9 presents the results. In Panel A, we find that Impact-seeking funds are more likely to support an ES shareholder proposal than other, non-ES shareholder proposals (ratio = 1.101). This ratio for Impact funds is significantly higher than for other sustainable funds. In contrast, both Financial funds and Moral funds are less likely to support an ES shareholder proposal than other, non-ES shareholder proposals (ratio = 0.821 and 0.838).

Similarly, while Impact funds are more likely to vote against ES shareholder proposals than other, non-ES proposals (ratio = 3.251), this ratio is still significantly smaller than for other sustainable funds (ratio = 5.712). Impact funds' ratios for abstention and non-voting are also significantly smaller than for other sustainable funds. Overall, Impact funds are comparatively more activist, more likely to vote for ES shareholder proposals and more likely to support them. These results are consistent with Hypothesis 4c.

Panel B presents the regression results. Funds with a stronger focus on Impact generation are more likely to vote, less likely to abstain, less likely to vote against, and more likely to vote for ES shareholder proposals. In contrast, Financial and Moral goals are not associated with more engagement with ES shareholder proposals.

5.6 ESG Performance Changes during Investment Holding Period

We now examine whether companies' ESG performance improves during the investment holding period by impact funds. If impact funds are willing to invest in firms that are laggards for the purpose of improving their ESG practice, and if they vote more frequently in support of shareholder proposals, then we expect them to sell the stock after the firms' ESG performance has improved on average. We measure the ESG performance change during holding periods in two ways. First, for stocks that exit a fund's portfolio in a given quarter, we calculate the difference between the stock's (within-industry) ESG percentile ranking as of the last quarter before it exits and the average ESG percentile ranking during the holding period. In other words, this captures the within-stock performance improvement. Second, for stocks that exit a fund's portfolio in a given quarter, we calculate the difference between the ESG ratings of the exiting stock and the average ESG ratings of the remaining portfolio in that quarter. This captures the within-portfolio performance differential at the time of exit.

Table 10 reports the results using the first, within-stock improvement measure. First, we find that stocks that are exited on average deteriorate in their ESG performance during the holding period. This is not surprising, since exiting decisions are made endogenously and it is natural for all fund managers to sell stocks with disappointing performance, financial or sustainable. Second, consistent with our conjecture, we find that stocks that are exited by impact funds show better performance change during the holding period compared to other sustainable funds. The results are driven by Pollution & Waste Management, Environmental Opportunities, Product Liability, Stakeholder Opposition, and Social Opportunities subcategories. In Panel B, the results are qualitatively unchanged.

[Insert Table 10: ESG Performance Changes During Investment Holding Periods]

Table 11 presents the ESG performance of stocks at exit relative to the portfolio. Similar to Table 10, we find that stocks that are exited on average have worse ESG rating than the stocks remaining in the portfolios. And we also find that stocks that are exited by impact funds show better relative performance compared to the portfolio. In sum, we find some suggestive evidence that impact funds improve the ESG performance of the companies they invest in during the investment holding period, relative to other sustainable funds.

[Insert Table 11: ESG Rating of Exited Companies Relative to the Portfolio]

5.7 Expansion of the BERT-Model Analysis

We are currently expanding our BERT-model analysis to include non-Morningstar U.S. mutual funds. While the efforts are ongoing, we present preliminary classification results in this section using a 2022 subsample.

While over 7,000 U.S. mutual funds filed fund prospectus with SEC in 2022, not all these fund prospectuses contain sentences with ESG keywords. Thus, conditioning on having at least one sentence with ESG keywords, the BERT model is applied to approximately 5,000 mutual funds. Panel A of Figure 4 presents the results of the BERT model analysis in a Venn diagram.

[Insert Figure 4: Expansion of the BERT-Model Analysis]

The BERT model classifies 747 non-Morningstar funds as pursuing one or more of the three sustainability goals, while 4,245 non-Morningstar funds are not classified as sustainable. Combined with the 461 Morningstar-list funds that the BERT model classifies, the BERT model identifies and classifies 1,208 funds as meeting the criteria for one or more of the three sustainability goals. 66 of the Morningstar funds are not classified as sustainable by the BERT model. Overall, the Venn diagram indicates that there is a substantial overlap between Morningstar’s sustainable fund selection and the BERT model’s classification criteria. But there is also a significant non-overlap. We plan to examine the determinants of these patterns and report the results in a future version of the paper.

Panel B of Figure 4 presents the classification of the 747 non-Morningstar funds that the BERT model identifies as sustainable into the three sustainable goals. By far the most common fund type is the pure Financial fund (573, or 77% of total), and the Impact fund is the least common.

5.8 Robustness

As a robustness check, we exclude funds that hold less than 30% of their portfolio in stocks that receive the MSCI ESG ratings and recreate Table 3 with the new sample. Table 12 presents the results, and they are qualitatively similar.

[Insert Table 12: Robustness]

6. Conclusion

Sustainable investment has grown exponentially in popularity in recent years as a potentially powerful channel through which private investors respond to the world’s most pressing problems, such as the climate change and social inequality. However, sustainable investment is not a cohesive investing principle catering to a uniform type of investors; instead, in its evolution it has fused three distinct investor goals under a single umbrella term, resulting in confusing, sometimes conflicting investing principles being marketed under the same name to end investors, who are also distinctly heterogeneous in their preferences.

We propose to alleviate this confusion by developing a supervised machine-learning model-based method that classifies investment managers' stated goals on sustainability into three distinct objectives: financial value, categorical morality, and impact. This is achieved by evaluating two dimensions of investor preferences: (i) whether investors have nonpecuniary preferences or not (*value* vs. *values*) and (ii) whether investors have *ex ante*, categorical moral preferences or *ex post*, consequentialist impact preferences. We are one of the first papers to make the latter distinction, and the first to apply the distinction to classify sustainable mutual funds.

We apply the model to prospectuses of U.S. mutual funds. Among the funds identified as sustainable by Morningstar, 54% state they incorporate ESG to enhance financial performance, while 39% practice categorical morality via exclusion and only 33% state they seek to generate impact. Stated goals meaningfully correlate with how the funds are managed. Financially motivated funds systematically hold stocks with high MSCI ESG ratings relative to industry peers, which is consistent with ESG risk management. Morally motivated funds categorically tilt away from companies in controversial industries (e.g., mining), but are otherwise insensitive to relative ESG ratings. Impact funds hold stocks with lower ESG performance than the others, which is consistent with them engaging with laggard firms to generate positive impact. Impact funds are also more likely to support social and environmental shareholder proposals.

Hybrid funds are common. Funds combining financial and moral goals are the largest category and are growing the fastest. These funds are “doing well” *and* “doing good (ethically pure)” side-by-side, not “doing well *by* doing good (generating impact)”. Our findings shed light on why sustainable investing is often found not to be effective in generating impact: most funds labeled “sustainable” in fact do not seek to generate impact, and instead seek to enhance financial value and/or satisfy ethical purity criteria. We argue that introduction of commonly accepted and mandated classification labels that distinguish among the financial, moral, and impact sustainability goals and using them in place of the all-too-broad “sustainability” term would improve investor welfare and advance research.

References

- Abis, Simona, 2022, “Man vs. Machine: Quantitative and Discretionary Equity Management”, unpublished working paper.
- Abis, Simona, and Anton Lines, 2023, “Broken Promises, Competition, and Capital Allocation in the Mutual Fund Industry”, unpublished working paper.
- Acikalin, Utku and Caskurlu, Tolga and Hoberg, Gerard and Phillips, Gordon M. (2022) “Intellectual Property Protection Lost and Competition: An Examination Using Large Language Models”, unpublished working paper.
- Atta-Darkua, Vaska, Glossner, Simon, Krueger, Philipp and Matos, Pedro, “Decarbonizing Institutional Investor Portfolios: Helping to Green the Planet or Just Greening Your Portfolio?”, unpublished working paper.
- Baker, Malcolm, Mark L. Egan, and Suproteem K. Sarkar (2022) *How Do Investors Value ESG?*. No. w30708. National Bureau of Economic Research.
- Barber, Brad, Adair Morse, and Ayako Yasuda (2021) “Impact Investing” *Journal of Financial Economics* 139, 162-185
- Berk, Jonathan, and Jules H. Van Binsbergen. "The impact of impact investing." (2021), unpublished working paper.
- Bingler, Julia and Kraus, Mathias and Leippold, Markus and Webersinke, Nicolas (2022) “How Cheap Talk in Climate Disclosures relates to Climate Initiatives, Corporate Emissions, and Reputation Risk”, unpublished working paper.
- Bonnefon, J-F, A. Landier, P. Sastry, and D. Thesmar, 2023, “The moral preferences of investors: Experimental evidence” (2022), unpublished working paper.
- Bybee, Leland, Bryan T. Kelly, Asaf Manela, and Dacheng Xiu. (2021) “Business news and business cycles”, unpublished working paper.
- Chowdhry, Bhagwan, Shaun William Davies, and Brian Waters (2019) “Investing for Impact” *The Review of Financial Studies* 32: 864–904
- 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.
- Duchin, R., J. Gao, and Q. Xu, “Sustainability or Greenwashing: Evidence from the Asset Market for Industrial Pollution”, unpublished working paper.

- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. (2019) "Text as data." *Journal of Economic Literature* 57, 535-574.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, Zhenhao Tan, Stephen Utkus, and Xiao Xu, 2023, Four facts about ESG beliefs and investor portfolios, unpublished working paper.
- 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, Daniel, and Benjamin Roth. "The allocation of socially responsible capital." Unpublished working paper.
- Hardy, C., M. Lambert, K. Yang, and A. Yasuda (2023) "What Sustainability Goals Does EU's Sustainable Finance Disclosure Regulation Promote? A Machine-Learning Model Approach", Unpublished working paper.
- 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, Yazhou Ellen, Bige Kahraman, and Michelle Lowry, "ES Risks and Shareholder Voice", forthcoming in *The Review of Financial Studies*.
- Heeb, Florian, Julian F. Kölbel, Falko Paetzold and Stefan Zeisberger (2023) "Do Investors Care about Impact?" *Review of Financial Studies* 36, 1737-1787.
- Hong, Harrison, and Marcin Kacperczyk, 2009, The price of sin: The effects of social norms on markets, *Journal of Financial Economics* 93, 15–36.
- Johnson, Steve, 2023 "Companies with good ESG scores pollute as much as low-rated rivals" *Financial Times*, July 30, 2023.
- 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 (2020) "ESG Investing: How to Optimize Impact?" Unpublished working paper

- Lee, Matthew, Arzi Adbi, and Jasjit Singh (2020) "Categorical cognition and outcome efficiency in impact investing decisions" *Strategic Management Journal* 41:86–107
- Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan. (2021) "Measuring corporate culture using machine learning." *The Review of Financial Studies* 34, 3265-3315.
- Lo, Andrew and Ruixun Zhang, 2021 "Quantifying the Impact of Impact Investing", Unpublished working paper.
- Lins, K.V., Servaes, H. and Tamayo, A. (2017), "Social Capital, Trust, and Firm Performance: The Value of Corporate Social Responsibility during the Financial Crisis" *The Journal of Finance*, 72: 1785-1824.
- 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.
- Oehmke, Martin and Opp, Marcus M. (2022) "A Theory of Socially Responsible Investment" Unpublished working paper
- Pástor, Lubos, Robert Stambaugh, and Lucian Taylor (2021) "Sustainable Investing in Equilibrium" *Journal of Financial Economics*, 142(2), 550-571
- Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski (2021) "Responsible Investing: The ESG-efficient Frontier" *Journal of Financial Economics* 142(2) 572–597
- Rajan, Raghuram G. and Ramella, Pietro and Zingales, Luigi (2022) "What Purpose Do Corporations Purport? Evidence from Letters to Shareholders", Unpublished working paper.
- Shanker, Harshini (2023) "Ethics and Externalities in Equilibrium", unpublished working paper.
- Starks, Laura (2023) "Presidential Address: Sustainable Finance and ESG Issues—*Value* versus *Values*" *Journal of Finance* 78, 1837-1872.
- U.S. Securities and Exchange Commission. (2016) "How to Read a Mutual Fund Prospectus" from the SEC website.
- Yasuda, Ayako (2023) "Impact Investing", 2023, In: Cumming, D., Hammer, B. (eds.) The Palgrave Encyclopedia of Private Equity, Palgrave Macmillan.

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','discrimination','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
$Financial_i$	equal to 1 if the prospectus for fund i contains at least one sentence that is classified as Financial / ESG-aware
$Moral_i$	equal to 1 if the prospectus for fund i contains at least one sentence that is classified as Moral / SRI
$Impact_i$	equal to 1 if the prospectus for fund i contains at least one sentence that is classified as Impact investing
$Unclassified_i$	equal to 1 if the prospectus for fund i does not contain any sentences that are classified as one of the three sustainability goals
$Financial\ Ratio_i$	the total number of $Financial$ sentences divided by the total number of sentences classified as one of the three sustainability goal types in the prospectus of fund i
$Moral\ Ratio_i$	the total number of $Moral$ sentences divided by the total number of sentences classified as one of the three sustainability goal types in the prospectus of fund i
$Impact\ Ratio_i$	the total number of $Impact$ sentences divided by the total number of sentences classified as one of the three sustainability goal types in the prospectus of fund i
$Unclassified\ Ratio_i$	total number of sentences that are not classified as any of the three sustainability goals divided by the total number of ESG-related sentences for fund i . It measures the opacity/ambiguity of the fund's sustainability goals.
$ Holding-quarters_{it}$	the average number of quarters for which the stocks in the portfolio have been held by fund i as of year t .
$ESG\ Score_{jt}$	the average of the evaluated MSCI Environmental and Social key issues scores that firm j received from MSCI in year t
$E\ Score_{jt}$	the average of the evaluated MSCI Environmental key issues scores that firm j received from MSCI in year t
$S\ Score_{jt}$	the average of the evaluated MSCI Social key issues scores that firm j received from MSCI in year t
$Climate\ Change_{jt}$	the average of the evaluated MSCI Climate Change theme key issues scores that firm j received from MSCI in year t
$Natural\ Capital_{jt}$	the average of the evaluated MSCI Natural Capital theme key issues scores that firm j received from MSCI in year t
$Pollution\ \&\ Waste_{jt}$	the average of the evaluated MSCI Pollution and Waste theme key issues scores that firm j received from MSCI in year t
$Envir.\ Opp_{jt}$	the average of the evaluated MSCI Environmental Opportunities theme key issues scores that firm j received from MSCI in year t
$Human\ Capital_{jt}$	the average of the evaluated MSCI Human Capital theme key issues scores that firm j received from MSCI in year t

<i>Product Liability</i> _{jt}	the average of the evaluated MSCI Product Liability theme key issues scores that firm <i>j</i> received from MSCI in year <i>t</i>
<i>Stakeholder Opposition</i> _{jt}	the average of the evaluated MSCI Stakeholder Opposition theme key issues scores that firm <i>j</i> received from MSCI in year <i>t</i>
<i>Social Opp.</i> _{jt}	the average of the evaluated MSCI Social Opportunities theme key issues scores that firm <i>j</i> received from MSCI in year <i>t</i>
<i>ESG Ranking</i> _{jt}	the percentile ranking of its <i>ESG Score</i> _{jt} within the MSCI industry for firm <i>j</i> in year <i>t</i> , normalized to scale from 0 to 10,
<i>ESG Industry</i>	the value-weighted industry-year level MSCI ESG rating using stock-year level ESG Score based on the average market value of each stock in that year from CRSP
Energy	equal to 100 if firm belongs to MSCI industry “Oil & Gas”
Aerospace	equal to 100 if firm belongs to MSCI industry “Aerospace & Defense”
Casino	equal to 100 if firm belongs to MSCI industry “Casinos & Gaming”
Tobacco	equal to 100 if firm belongs to MSCI industry “Tobacco”
Utilities	equal to 100 if firm belongs to MSCI industry “Utilities”
ES "For" Vote/Total "For" Vote _{it}	The percentage of Environmental and Social shareholder proposals that fund <i>i</i> voted for in year <i>t</i> , divided by the percentage of all shareholder proposals that fund <i>i</i> voted for in year <i>t</i> .
ES "Against" Vote/Total "Against" Vote	The percentage of Environmental and Social shareholder proposals that fund <i>i</i> voted against in year <i>t</i> , divided by the percentage of all shareholder proposals that fund <i>i</i> voted against in year <i>t</i> .
ES "Abstain" Vote/Total "Abstain" Vote	The percentage of Environmental and Social shareholder proposals that fund <i>i</i> abstained from voting in year <i>t</i> , divided by the percentage of all shareholder proposals that fund <i>i</i> abstained from voting in year <i>t</i> .
ES "Do Not Vote"/Total "Do Not Vote"	The percentage of Environmental and Social shareholder proposals that fund <i>i</i> did not cast a vote in year <i>t</i> , divided by the percentage of all shareholder proposals that fund <i>i</i> did not cast a vote in year <i>t</i> .

Table 1: BERT Model Classification Performance

This table reports the performance of the BERT model trained to classify ESG sentences. In Panel A, 2,277 ESG-related sentences in the training sample and 557 ESG-related sentences in the testing sample are manually classified into four different types: Financial, Moral, Impact, or none. The BERT model is trained on the training sample for classification. The trained BERT model is applied on the testing sample to obtain the ESG classification given by the BERT model. Three 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. 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. The model performance measures for all classifications are calculated separately. For example, when calculating the BERT model's performance in predicting Financial sentences, we set all non-Financial-related predictions and true values as zero and calculated the three metrics only for Financial-related true values and predicted values. 'All' means the sentence is Financial, Moral, or Impact. Panel B reports the performance of the BERT model trained to classify ESG sentences at the fund level. A fund is classified as “Financial” if it contains at least one Financial sentence. Similar classification applies to “Moral” and “Impact” funds. A fund could be classified into multiple types. If none of the firms' ESG sentences are classified as Financial, Moral, or Impact, the fund is classified as “Unclassified”. Both BERT model classification and manual classification of ESG sentences are used to classify funds' objectives. Three performance metrics, including accuracy, precision, and recall, on the categorical classification are reported. The last column, 'Ratio', reports the correlation between the BERT-predicted ratio of each category and the manually coded ratio for each category.

Panel A			
	Accuracy	Precision	Recall
All	0.832	0.795	0.795
Financial	0.869	0.812	0.723
Moral	0.978	0.857	0.909
Impact	0.905	0.900	0.659

Panel B				
	Accuracy	Precision	Recall	Ratio
Financial	0.941	0.974	0.927	0.942
Moral	0.956	0.900	1.000	0.988
Impact	0.912	0.800	0.800	0.915
Unclassified				0.838

Table 2: Summary Statistics

This table reports the summary statistics of key variables. Panel A include fund-year level measure for all morningstar funds in 2018, 2019, 2020, and 2022. Panel B include fund-year level measure for morningstar funds in 2018, 2019, 2020, and 2022 that are rated by MSCI at least once after 2014. Panel C report the correlation between each categorical variable, using the classification based on BERT model predictions.

Panel A: All Funds						
	Mean	Median	Std Dev	25	75	Obs.
Financial Dummy	0.545	1	0.498	0	1	1361
Moral Dummy	0.386	0	0.487	0	1	1361
Impact Dummy	0.331	0	0.471	0	1	1361
Financial Ratio	0.506	0.5	0.406	0	1	1090
Moral Ratio	0.223	0	0.313	0	0.33	1090
Impact Ratio	0.271	0	0.38	0	0.5	1090
Unclassified Ratio	0.634	0.667	0.292	0.455	0.87	1361

Panel B: MSCI Funds						
	Mean	Median	Std Dev	25	75	Obs.
Financial Dummy	0.536	1.000	0.499	0.000	1.00	892
Moral Dummy	0.415	0.000	0.493	0.000	1.00	892
Impact Dummy	0.290	0.000	0.454	0.000	1.00	892
Financial Ratio	0.510	0.500	0.405	0.000	1.00	700
Moral Ratio	0.247	0.111	0.322	0.000	0.40	700
Impact Ratio	0.243	0.000	0.367	0.000	0.50	700
Unclassified Ratio	0.632	0.667	0.300	0.429	0.88	892
Per. of MSCI Stocks	0.618	0.733	0.338	0.339	0.93	892
Per. of Other Stocks	0.322	0.189	0.347	0.008	0.58	892
Per. of Other Assets	0.052	0.042	0.045	0.024	0.07	892
Per of MSCI Stocks over All stocks	0.661	0.803	0.359	0.370	0.99	892
Holding Quarters	6.434	5.985	3.815	3.693	8.65	892

Panel C: Fund Classifications Correlations using Bert Classification			
	Financial	Moral	Impact
Financial	1.000		
Moral	0.195***	1.000	
Impact	-0.122***	-0.027	1.000

Table 3: ESG Rating Baseline Model

This table reports the fund's portfolio companies' MSCI ESG ratings across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the fund's portfolio companies' MSCI ESG ratings, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. Five different stock-year level ESG ratings are constructed based on the MSCI ESG rating dataset. ESG Score is the average of the evaluated MSCI Environmental and Social Key Issues scores. E Score is the average of the evaluated MSCI Environmental Key Issues scores. S Score is the average of the evaluated MSCI Social Key Issues scores. For each fund, we calculate the value-weighted ESG ratings for the calendar years 2018, 2019, 2020, and 2022. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. Standard deviations are clustered at fund level.

Panel A: Average ESG Ratings

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
ESG Score	6.179	6.362	0.712	5.726	6.68	892		
E Score	6.954	7.180	0.895	6.404	7.58	892		
S Score	5.405	5.501	0.736	5.042	5.80	892		
Financial Funds								
ESG Score	6.342	6.513	0.660	6.067	6.77	478	0.351***	(5.345)
E Score	7.140	7.372	0.830	6.705	7.69	478	0.401***	(4.708)
S Score	5.545	5.609	0.665	5.288	5.90	478	0.302***	(4.577)
Moral Funds								
ESG Score	6.325	6.494	0.639	5.994	6.79	370	0.250***	(3.799)
E Score	7.091	7.301	0.801	6.564	7.70	370	0.235***	(2.789)
S Score	5.559	5.615	0.656	5.222	5.91	370	0.264***	(4.023)
Impact Funds								
ESG Score	6.104	6.306	0.738	5.579	6.66	259	-0.106	(-1.385)
E Score	6.814	6.982	0.894	6.256	7.47	259	-0.197**	(-2.068)
S Score	5.394	5.495	0.805	4.863	5.85	259	-0.015	(-0.194)

Panel B: Regression Results

	(1) ESG Score	(2) E Score	(3) S Score
Financial Ratio	0.284*** (3.072)	0.398*** (3.549)	0.169* (1.696)
Moral Ratio	0.134 (1.355)	0.075 (0.626)	0.194 (1.576)
Impact Ratio	-0.448*** (-4.464)	-0.541*** (-4.532)	-0.354*** (-3.045)
Observations	700	700	700

Table 4: Environmental Rating Decomposition

This table reports the fund's portfolio companies' MSCI ESG ratings for each Environmental Ratings category across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the fund's portfolio companies' MSCI ESG ratings, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. Four different stock-year level ESG ratings for each category are constructed based on the MSCI ESG rating dataset. For the Climate Change category, the average of the evaluated MSCI Climate Change Key Issues scores. The same definition applies to the Natural Capital, Pollution & Waste, and Environmental Opportunities categories. For each fund, we calculate the value-weighted ESG ratings for the calendar years 2018, 2019, 2020, and 2022. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. Standard deviations are clustered at fund level.

Panel A: Average Environmental Ratings

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
Climate Change	8.001	8.241	1.009	7.469	8.72	892		
Natural Capital	6.763	7.175	1.319	5.980	7.66	882		
Pollution & Waste	5.606	5.480	1.364	4.892	6.26	834		
Envir. Opp.	5.068	4.961	1.027	4.492	5.46	854		
Financial Funds								
Climate Change	8.219	8.456	0.925	7.755	8.85	478	0.470***	(4.927)
Natural Capital	7.085	7.416	1.179	6.518	7.89	476	0.699***	(5.437)
Pollution & Waste	5.677	5.543	1.341	4.967	6.25	447	0.154	(1.388)
Envir. Opp.	4.926	4.958	0.890	4.436	5.37	458	-0.306***	(-3.135)
Moral Funds								
Climate Change	8.177	8.406	0.917	7.656	8.87	370	0.301***	(3.150)
Natural Capital	6.980	7.295	1.200	6.286	7.83	369	0.372***	(2.934)
Pollution & Waste	5.528	5.452	1.228	4.893	6.04	350	-0.135	(-1.242)
Envir. Opp.	4.940	4.908	0.951	4.403	5.39	354	-0.219**	(-2.422)
Impact Funds								
Climate Change	7.908	8.112	1.033	7.363	8.68	259	-0.131	(-1.216)
Natural Capital	6.562	6.726	1.266	5.713	7.58	257	-0.283**	(-2.016)
Pollution & Waste	5.609	5.499	1.269	4.895	6.31	240	0.003	(0.029)
Envir. Opp.	5.234	5.152	1.042	4.642	5.77	250	0.235**	(2.131)

Panel B: Regression Results

	(1) Climate Change	(2) Natural Capital	(3) Pollution & Waste	(4) Envir. Opp.
Financial Ratio	0.355*** (2.820)	0.657*** (3.956)	0.188 (1.292)	-0.334** (-2.497)
Moral Ratio	0.121 (0.881)	0.143 (0.815)	-0.189 (-1.120)	-0.369** (-2.507)
Impact Ratio	-0.524*** (-3.904)	-0.914*** (-5.238)	-0.085 (-0.514)	0.680*** (4.169)
Observations	700	696	649	672

Table 5: Social Rating Decomposition

This table reports the fund's portfolio companies' MSCI ESG ratings for each Social Ratings category across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the fund's portfolio companies' MSCI ESG ratings, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. Four different stock-year level ESG ratings for each category are constructed based on the MSCI ESG rating dataset. For the Human category, the average of the evaluated MSCI Human Key Issues scores. The same definition applies to the Product Liability, Stakeholder Opposition, and Social Opportunities categories. For each fund, we calculate the value-weighted ESG ratings for the calendar years 2018, 2019, 2020, and 2022. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. Standard deviations are clustered at fund level.

Panel A: Average Social Ratings

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
Human Capital	5.201	5.272	0.857	4.791	5.70	892		
Product Liability	5.919	6.045	1.125	5.124	6.69	882		
Stakeholder Opposition	7.093	7.267	1.193	6.445	7.83	759		
Social Opp.	4.661	4.900	0.909	4.261	5.23	737		
Financial Funds								
Human Capital	5.346	5.373	0.755	5.005	5.76	478	0.312***	(3.849)
Product Liability	6.013	6.175	1.051	5.286	6.74	477	0.203**	(2.410)
Stakeholder Opposition	7.128	7.279	1.184	6.484	7.83	415	0.077	(0.655)
Social Opp.	4.808	4.941	0.810	4.495	5.30	424	0.345***	(3.466)
Moral Funds								
Human Capital	5.347	5.413	0.777	4.984	5.77	370	0.249***	(3.090)
Product Liability	6.069	6.214	1.030	5.528	6.80	370	0.257***	(3.162)
Stakeholder Opposition	7.164	7.266	0.991	6.524	7.76	329	0.125	(1.174)
Social Opp.	4.668	4.891	0.857	4.258	5.24	344	0.014	(0.142)
Impact Funds								
Human Capital	5.117	5.207	0.933	4.551	5.71	259	-0.119	(-1.263)
Product Liability	5.990	6.102	1.196	5.174	6.81	258	0.100	(1.058)
Stakeholder Opposition	7.038	7.245	1.285	6.282	7.91	231	-0.079	(-0.566)
Social Opp.	4.556	4.899	1.077	3.707	5.24	206	-0.145	(-1.176)

Panel B: Regression Results

	(1)	(2)	(3)	(4)
	Human Capital	Product Liability	Stakeholder Opposition	Social Opp.
Financial Ratio	0.235* (1.896)	0.044 (0.390)	0.091 (0.553)	0.349** (2.288)
Moral Ratio	0.259* (1.718)	0.111 (0.876)	0.351** (2.071)	-0.173 (-1.130)
Impact Ratio	-0.483*** (-3.599)	-0.139 (-1.076)	-0.373* (-1.968)	-0.319 (-1.626)
Observations	700	698	607	602

Table 6: Industry Tilting Decisions by Moral Funds

This table reports the fund's portfolio industry' MSCI ESG ratings across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the fund's portfolio industries' MSCI ESG ratings, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. For each MSCI industry, we compute the total market value of the industry's stock held by each fund for each calendar year. Two different industry-year level ESG ratings are constructed based on the MSCI ESG rating dataset. First, we calculate two different stock-year level ESG ratings based on the MSCI ESG rating dataset. Then, we calculate the value-weighted industry-year level MSCI ESG rating using stock-year level ESG ratings based on the average market value of each stock in that year from CRSP. For each fund, we calculate the value-weighted industry ESG ratings for the calendar years 2018, 2019, 2020, and 2022. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, we redo all analyses as in Panel A, but we exclude hybrid funds - those with more than one sustainable investment objective. In Panel C, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. In columns (2), we exclude hybrid funds. Standard deviations are clustered at fund level.

Panel A: Average Ratings

	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
ESG Industry	6.080	6.087	0.331	5.914	6.31	892		
Financial Funds								
ESG Industry	6.119	6.121	0.292	5.967	6.34	478	0.086***	(3.124)
Moral Funds								
ESG Industry	6.145	6.132	0.282	5.985	6.36	370	0.112***	(4.399)
Impact Funds								
ESG Industry	6.096	6.089	0.341	5.907	6.33	259	0.023	(0.748)

Panel B: Average Ratings Excluding Hybrid Funds

	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
ESG Industry	6.036	6.045	0.349	5.872	6.24	551		
Financial Funds								
ESG Industry	6.061	6.093	0.302	5.907	6.30	183	0.038	(1.083)
Moral Funds								
ESG Industry	6.116	6.128	0.278	5.957	6.31	76	0.093**	(2.100)
Impact Funds								
ESG Industry	6.017	6.010	0.369	5.774	6.26	100	-0.023	(-0.425)

Panel C: Regression Results

	(1)	(2)
	ESG Industry	ESG Industry
	Exclude Hybrid Funds	
Financial Ratio	0.021 (0.587)	0.002 (0.036)
Moral Ratio	0.066 (1.580)	0.070 (1.532)
Impact Ratio	-0.076* (-1.691)	-0.060 (-1.126)
N	700	359

Table 7: What industries do Moral Funds hold?

This table reports the fund's portfolio companies' industry holdings across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the fund's portfolio companies' industry distribution, we construct indicator variables that equal one hundred if the fund's portfolio companies belong to each of the following five industries: Energy, Aerospace, Casino, Tobacco, and Utilities. For each fund, we calculate the value-weighted indicator variables for the calendar years 2018, 2019, 2020, and 2022. The value-weighted indicator variables are our measures of fund's portfolio companies' industry distribution. In Panel A, summary statistics are provided for Moral funds, all sustainable funds, and all CRSP mutual funds. The column labeled 'Difference' calculates the average difference between Moral funds and non-Moral belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, we redo all analyses as in Panel A, but we exclude hybrid funds - those with more than one sustainable investment objective. In Panel C, we report the summary statistics of the value-weighted industry distribution of these five industries among S&P 500 firms from 2018 to 2022.

Panel A: Summary Statistics for Moral Funds, All Sustainable Funds, and All CRSP Funds

	Moral Funds		All Sustainable Funds				All CRSP Mutual Funds			
	Mean	Std Dev	Mean	Std Dev	Difference	P-Value	Mean	Std Dev	Difference	P-Value
Energy	1.573	2.859	1.978	4.853	-0.691*	(-1.956)	2.198	6.138	-0.757**	(-2.401)
Aerospace	0.412	0.921	0.898	4.204	-0.832*	(-1.881)	0.785	3.207	-0.462**	(-2.379)
Casino	0.470	4.447	0.232	2.894	0.406	(1.143)	0.209	2.504	0.316	(0.978)
Tobacco	0.046	0.346	0.089	0.699	-0.074*	(-1.654)	0.171	1.224	-0.155**	(-2.303)
Utilities	3.257	7.425	5.553	12.480	-3.923***	(-3.936)	5.663	13.282	-2.978***	(-3.768)

Panel B: Summary Statistics for Moral Funds and All Sustainable Funds Excluding Hybrid Funds

	Moral Funds Excluding Hybrid Funds		All Sustainable Funds Excluding Hybrid Funds			
	Mean	Std Dev	Mean	Std Dev	Difference	P-Value
Energy	1.356	2.269	2.225	5.723	-1.008**	(-2.172)
Aerospace	0.487	0.813	1.201	5.271	-0.828*	(-1.654)
Casino	0.040	0.117	0.065	0.487	-0.029	(-0.800)
Tobacco	0.022	0.113	0.117	0.843	-0.110**	(-2.455)
Utilities	3.015	9.652	6.709	14.470	-4.284**	(-2.551)

Panel C: SP 500 Firms 2018 to 2022

	Mean	Median	Std Dev	25	75	Obs.
Energy	3.590	3.981	1.116	2.473	4.15	5
Aerospace	2.548	2.561	0.326	2.287	2.84	5
Casino	0.208	0.195	0.058	0.185	0.25	5
Tobacco	1.092	1.018	0.236	0.919	1.22	5
Utilities	3.027	3.104	0.324	2.940	3.13	5

Table 8: ESG Rating Percentile Ranks Within Industry

This table reports the fund's portfolio companies' MSCI ESG ratings' percentile rankings across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the fund's portfolio companies' MSCI ESG ratings, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. The ESG ranking of a firm is the percentile ranking of its ESG score within the MSCI industry for each year, normalized to scale from 0 to 10. ESG Score is the average of the evaluated MSCI Environmental and Social Key Issues scores. For each fund, we calculate the value-weighted ESG ratings for the calendar years 2018, 2019, 2020, and 2022. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, we redo all analyses as in Panel A, but we exclude hybrid funds - those with more than one sustainable investment objective. In Panel C, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. In columns (2), we exclude hybrid funds. Standard deviations are clustered at fund level.

Panel A: Average Rankings

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
ESG Ranking	7.616	8.022	1.288	6.909	8.52	892		
Financial Funds								
ESG Ranking	7.937	8.302	1.164	7.601	8.64	478	0.692***	(5.629)
Moral Funds								
ESG Ranking	7.908	8.248	1.108	7.445	8.67	370	0.499***	(4.160)
Impact Funds								
ESG Ranking	7.411	7.878	1.360	6.647	8.45	259	-0.288**	(-1.992)

Panel B: Average Rankings Excluding Hybrid Funds

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
ESG Ranking	7.443	7.876	1.312	6.647	8.39	551		
Financial Funds								
ESG Ranking	7.840	8.218	1.166	7.364	8.58	183	0.594***	(3.798)
Moral Funds								
ESG Ranking	7.925	8.141	1.039	7.408	8.62	76	0.559***	(3.086)
Impact Funds								
ESG Ranking	6.978	6.935	1.261	6.108	8.05	100	-0.568***	(-2.846)

Panel C: Regression Results

	(1) ESG Ranking	(2) ESG Ranking
		Exclude Hybrid Funds
Financial Ratio	0.597*** (3.541)	0.453** (2.336)
Moral Ratio	0.326* (1.808)	0.390** (2.103)
Impact Ratio	-0.975*** (-5.490)	-0.887*** (-4.483)
Observations	700	359

Table 9: Environmental and Social Shareholder Proposal Voting

This table reports the fund's shareholders proposals across the fund's sustainable investment objectives. For each fund year, we calculate the percentages of "For," "Against," "Abstain," and "Do Not Vote" among all shareholder proposals and among Environmental and Social (ES) proposals. We then calculate the ratio of the percentages of "For," "Against," "Abstain," and "Do Not Vote" for ES proposals to those for all proposals. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. Standard deviations are clustered at fund level.

Panel A: Average Ratios

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	P-Value
ES "For" Vote/Total "For" Vote	0.81	0.772	0.748	0.28	1.12	271		
ES "Against" Vote/Total "Against" Vote	5.712	2.903	9.047	0.543	8.48	269		
ES "Abstain" Vote/Total "Abstain" Vote	3.056	0	8.111	0	1.76	137		
ES "Do Not Vote"/Total "Do Not Vote"	7.107	0	22.047	0	4.67	104		
Financial Funds								
ES "For" Vote/Total "For" Vote	0.821	0.689	0.895	0.231	1.12	151	0.011	(0.082)
ES "Against" Vote/Total "Against" Vote	6.749	4.613	10.554	0.828	8.56	151	2.618**	(2.073)
ES "Abstain" Vote/Total "Abstain" Vote	4.187	0	7.946	0	6.12	73	2.423*	(1.761)
ES "Do Not Vote"/Total "Do Not Vote"	7.6	1.222	18.925	0	4.89	49	1.863	(0.442)
Moral Funds								
ES "For" Vote/Total "For" Vote	0.838	0.612	1.026	0.228	1.03	101	0.035	(0.182)
ES "Against" Vote/Total "Against" Vote	7.692	5.157	11.934	0.095	10.14	101	3.345**	(2.185)
ES "Abstain" Vote/Total "Abstain" Vote	1.432	0	4.268	0	0	43	-2.367**	(-2.075)
ES "Do Not Vote"/Total "Do Not Vote"	10.924	2.172	29.933	0	6.25	43	7.458	(1.613)
Impact Funds								
ES "For" Vote/Total "For" Vote	1.101	1.054	1.06	0.318	1.29	83	0.414**	(2.008)
ES "Against" Vote/Total "Against" Vote	3.251	0.828	5.139	0.031	4.63	81	-3.415***	(-3.092)
ES "Abstain" Vote/Total "Abstain" Vote	0.631	0	1.605	0	0	57	-4.151***	(-3.281)
ES "Do Not Vote"/Total "Do Not Vote"	2.894	0	5.709	0	3.95	31	-5.406*	(-1.709)

Panel B: Regression Results

	(1)	(2)	(3)	(4)
	ES "For" Vote/Total "For" Vote	ES "Against" Vote/Total "Against" Vote	ES "Abstain" Vote/Total "Abstain" Vote	ES "Do Not Vote"/Total "Do Not Vote"
Financial Ratio	-0.264 (-1.482)	3.079 (1.523)	7.083*** (3.864)	2.168 (0.238)
Moral Ratio	-0.188 (-0.975)	3.692 (1.208)	-2.573 (-0.844)	13.434 (0.97)
Impact Ratio	0.368** (2.314)	-5.269*** (-3.289)	-5.698*** (-3.949)	-10.692** (-2.148)
Observations	209	207	109	75

Table 10: ESG Performance Changes During Investment Holding Periods

This table reports the MSCI ESG percentile rankings changes of the companies when they exit the fund's portfolio across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the MSCI ESG percentile rankings changes of the companies when they exit the fund's portfolio, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. For stocks that exit the portfolio in that quarter, we calculate the difference between the ESG percentile rankings and the average ESG percentile rankings of the stock from the first quarter when it enters the portfolio and the last quarter before it exits during the holding period. A holding period is defined as a period when the stock appears in the fund's portfolio without a gap. ESG Percentile Rankings are defined the same as in Table 8. For each fund, we calculate the ESG percentile rankings changes for the calendar years 2018, 2019, 2020, and 2022. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. Standard deviations are clustered at fund level.

Panel A: Average Ratings

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
Climate Change	-0.689	-0.418	1.813	-1.535	0.47	2458		
Natural Capital	-0.300	-0.197	1.956	-1.426	0.89	2263		
Pollution & Waste	-0.401	-0.233	2.267	-1.712	1.06	1837		
Envir. Opp.	-0.347	-0.069	2.145	-1.548	1.06	1793		
Human Capital	-0.265	-0.115	1.959	-1.199	0.93	2458		
Product Liability	-0.380	-0.353	1.990	-1.604	0.84	2278		
Stakeholder Opposition	-0.428	-0.130	2.404	-1.789	1.29	920		
Social Opp.	-0.510	-0.064	2.305	-1.685	1.07	1402		
Financial Funds								
Climate Change	-0.697	-0.418	1.826	-1.475	0.43	1290	-0.057	(-0.663)
Natural Capital	-0.363	-0.200	1.984	-1.472	0.83	1215	-0.149	(-1.471)
Pollution & Waste	-0.337	-0.228	2.265	-1.627	1.11	935	0.100	(0.791)
Envir. Opp.	-0.390	-0.197	2.161	-1.538	1.05	884	-0.103	(-1.021)
Human Capital	-0.274	-0.127	1.918	-1.196	0.92	1290	-0.046	(-0.506)
Product Liability	-0.472	-0.455	1.972	-1.624	0.73	1233	-0.211**	(-2.075)
Stakeholder Opposition	-0.463	-0.196	2.287	-1.706	1.22	486	-0.111	(-0.563)
Social Opp.	-0.482	0.000	2.291	-1.644	1.13	760	0.014	(0.097)
Moral Funds								
Climate Change	-0.639	-0.448	1.702	-1.437	0.41	1042	0.056	(0.641)
Natural Capital	-0.425	-0.293	1.868	-1.499	0.71	989	-0.234**	(-2.206)
Pollution & Waste	-0.489	-0.281	2.212	-1.798	1.00	794	-0.189	(-1.458)
Envir. Opp.	-0.347	-0.126	2.020	-1.523	0.95	736	-0.014	(-0.127)
Human Capital	-0.245	-0.116	1.833	-1.066	0.88	1042	0.014	(0.164)
Product Liability	-0.437	-0.425	1.914	-1.577	0.71	1016	-0.109	(-1.059)
Stakeholder Opposition	-0.566	-0.273	2.328	-1.932	1.12	439	-0.302*	(-1.757)
Social Opp.	-0.475	-0.041	2.404	-1.725	1.22	686	0.027	(0.186)
Impact Funds								
Climate Change	-0.725	-0.387	1.976	-1.672	0.60	670	-0.076	(-0.780)

Natural Capital	-0.283	-0.134	2.079	-1.493	1.15	593	0.019	(0.148)
Pollution & Waste	-0.216	-0.009	2.326	-1.572	1.39	495	0.234*	(1.728)
Envir. Opp.	-0.197	0.100	2.119	-1.287	1.16	483	0.198*	(1.797)
Human Capital	-0.194	-0.127	2.004	-1.186	1.11	672	0.082	(0.881)
Product Liability	-0.202	-0.145	2.096	-1.545	1.09	609	0.243*	(1.934)
Stakeholder Opposition	-0.164	0.000	2.349	-1.496	1.46	226	0.335*	(1.671)
Social Opp.	-0.230	0.032	2.319	-1.295	1.22	359	0.350**	(2.319)

Panel B: Regression Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Climate Change	Natural Capital	Pollution & Waste	Envir. Opp.	Human Capital	Product Liability	Stakeholder Opposition	Social Opp.
Financial Ratio	-0.017 (-0.145)	0.040 (0.275)	0.155 (0.829)	-0.101 (-0.702)	-0.070 (-0.566)	-0.192 (-1.317)	-0.051 (-0.188)	-0.113 (-0.573)
Moral Ratio	0.134 (0.850)	-0.349** (-2.112)	-0.619*** (-2.827)	-0.150 (-0.942)	-0.073 (-0.475)	-0.263* (-1.719)	-0.540* (-1.957)	-0.292 (-1.109)
Impact Ratio	-0.096 (-0.697)	0.276 (1.582)	0.368* (1.876)	0.257* (1.811)	0.157 (1.068)	0.525*** (2.646)	0.626** (2.138)	0.502** (2.062)
Observati ons	1873	1744	1383	1323	1875	1769	709	1098

Table 11: ESG Rating of Exited Companies Relative to the Portfolio

This table reports the MSCI ESG Rating of the companies when they exit the fund's portfolio relative to the portfolio across the fund's sustainable investment objectives. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the ESG Rating of the companies when they exit the fund's portfolio relative to the portfolio, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. For stocks that exit the portfolio in that quarter, we calculate the difference between the ESG ratings and the average ESG ratings of the portfolio in that quarter. For each fund, we calculate the relative ESG rating for the calendar years 2018, 2019, 2020, and 2022. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. Standard deviations are clustered at fund level.

Panel A: Average Ratings

All Funds	Mean	Median	Std Dev	25	75	Obs.	Difference	T-stats
Climate Change	-0.639	-0.488	1.661	-1.586	0.46	2407		
Natural Capital	-0.571	-0.651	1.673	-1.598	0.40	2216		
Pollution & Waste	-0.486	-0.474	1.724	-1.487	0.57	1780		
Envir. Opp.	-0.260	-0.279	1.319	-1.080	0.51	1746		
Human Capital	-0.173	-0.140	1.359	-0.886	0.61	2407		
Product Liability	-0.230	-0.274	1.560	-1.123	0.68	2233		
Stakeholder Opposition	-0.529	-0.413	2.141	-1.713	1.02	860		
Social Opp.	-0.366	-0.407	1.359	-1.280	0.43	1340		
Financial Funds								
Climate Change	-0.656	-0.504	1.633	-1.496	0.41	1272	-0.074	(-0.913)
Natural Capital	-0.588	-0.634	1.642	-1.563	0.34	1197	-0.050	(-0.549)
Pollution & Waste	-0.446	-0.387	1.649	-1.349	0.55	911	0.044	(0.437)
Envir. Opp.	-0.294	-0.289	1.311	-1.056	0.47	868	-0.081	(-1.151)
Human Capital	-0.160	-0.110	1.331	-0.856	0.64	1272	-0.000	(-0.001)
Product Liability	-0.284	-0.328	1.533	-1.166	0.58	1216	-0.142*	(-1.896)
Stakeholder Opposition	-0.592	-0.520	2.092	-1.590	0.87	463	-0.185	(-1.083)
Social Opp.	-0.403	-0.434	1.344	-1.293	0.43	737	-0.116	(-1.252)
Moral Funds								
Climate Change	-0.602	-0.461	1.547	-1.369	0.35	1028	0.036	(0.448)
Natural Capital	-0.594	-0.646	1.598	-1.566	0.32	974	-0.053	(-0.588)
Pollution & Waste	-0.428	-0.419	1.602	-1.352	0.56	774	0.069	(0.691)
Envir. Opp.	-0.298	-0.347	1.304	-1.075	0.47	721	-0.075	(-1.083)
Human Capital	-0.141	-0.099	1.273	-0.799	0.60	1028	0.033	(0.528)
Product Liability	-0.189	-0.211	1.496	-1.076	0.63	1002	0.060	(0.785)
Stakeholder Opposition	-0.636	-0.448	2.140	-1.723	0.87	417	-0.253	(-1.621)
Social Opp.	-0.373	-0.358	1.379	-1.316	0.43	665	-0.042	(-0.455)
Impact Funds								
Climate Change	-0.629	-0.428	1.746	-1.547	0.48	654	-0.010	(-0.116)
Natural Capital	-0.487	-0.589	1.773	-1.548	0.54	579	0.108	(1.005)

Pollution & Waste	-0.303	-0.280	1.712	-1.287	0.79	477	0.226**	(2.122)
Envir. Opp.	-0.153	-0.152	1.370	-1.042	0.61	469	0.141*	(1.898)
Human Capital	-0.165	-0.142	1.386	-0.975	0.69	656	-0.008	(-0.126)
Product Liability	-0.088	-0.153	1.610	-1.033	0.87	594	0.184**	(2.102)
Stakeholder Opposition	-0.164	-0.251	2.103	-1.277	1.30	203	0.459***	(2.620)
Social Opp.	-0.239	-0.278	1.369	-1.186	0.55	337	0.154	(1.575)

Panel B: Regression Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Climate Change	Natural Capital	Pollution & Waste	Envir. Opp.	Human Capital	Product Liability	Stakeholder Opposition	Social Opp.
Financial Ratio	-0.093 (-0.833)	-0.009 (-0.068)	-0.084 (-0.606)	-0.094 (-0.947)	0.003 (0.036)	-0.250** (-2.311)	-0.282 (-1.178)	-0.118 (-0.898)
Moral Ratio	0.141 (1.051)	-0.211 (-1.297)	-0.098 (-0.567)	-0.148 (-1.339)	-0.036 (-0.329)	0.024 (0.204)	-0.428* (-1.744)	-0.101 (-0.580)
Impact Ratio	-0.001 (-0.007)	0.215 (1.419)	0.205 (1.285)	0.250** (2.287)	0.028 (0.279)	0.334** (2.374)	0.893*** (3.620)	0.310* (1.861)
Observations	1840	1712	1344	1292	1842	1738	665	1051

Table 12: Robustness

This table reports the fund's portfolio companies' MSCI ESG ratings across the fund's sustainable investment objectives, considering only funds with significant holdings of MSCI stocks. For each fund's security holdings, we calculate the average market value of the securities for the calendar year using the CRSP mutual fund database. To calculate the fund's portfolio companies' MSCI ESG ratings, we ignore non-stock holdings and only include stock holdings that have at least one MSCI ESG rating from 2014 onwards. Five different stock-year level ESG ratings are constructed based on the MSCI ESG rating dataset. ESG Score is the average of the evaluated MSCI Environmental and Social Key Issues scores. E Score is the average of the evaluated MSCI Environmental Key Issues scores. S Score is the average of the evaluated MSCI Social Key Issues scores. For each fund, we calculate the value-weighted ESG ratings for the calendar years 2018, 2019, 2020, and 2022. We are only including funds that have over 30% in market value of their stock holdings covered by MSCI ESG ratings from 2014 onwards. In Panel A, summary statistics are provided for all funds, funds that are classified as Financial, Moral, and Impact. The column labeled 'Difference' calculates the average difference between funds within each group and funds not belonging to that group. The column 'T-stats' reports the t-statistics for the difference. In Panel B, the regression results are reported for regressing the average ESG ratings on the continuous measure of funds' sustainable investment types. Standard deviations are clustered at fund level.

Panel A: Average ESG Ratings

All Funds	Mean	Median	StdDev	25	75	Obs.	Difference	T-stats
ESG Score	6.331	6.491	0.644	5.970	6.78	703		
E Score	7.237	7.453	0.824	6.791	7.80	703		
S Score	5.424	5.521	0.584	5.148	5.77	703		
Financial Funds								
ESG Score	6.514	6.626	0.579	6.315	6.89	379	0.397***	(5.804)
E Score	7.458	7.639	0.730	7.151	7.96	379	0.481***	(5.303)
S Score	5.569	5.612	0.519	5.351	5.87	379	0.313***	(5.227)
Moral Funds								
ESG Score	6.462	6.648	0.640	6.153	6.91	295	0.227***	(3.175)
E Score	7.369	7.627	0.818	6.942	7.97	295	0.229**	(2.448)
S Score	5.555	5.629	0.554	5.313	5.90	295	0.225***	(3.686)
Impact Funds								
ESG Score	6.262	6.439	0.696	5.737	6.77	216	-0.098	(-1.221)
E Score	7.106	7.355	0.862	6.579	7.74	216	-0.189*	(-1.854)
S Score	5.419	5.523	0.670	4.975	5.82	216	-0.007	(-0.105)

Panel B: Regression Results

	(1) ESG Score	(2) E Score	(3) S Score
Financial Ratio	0.342*** (3.434)	0.480*** (3.773)	0.205** (2.382)
Moral Ratio	0.055 (0.525)	0.026 (0.186)	0.084 (0.904)
Impact Ratio	-0.438*** (-4.131)	-0.579*** (-4.419)	-0.298*** (-2.898)
Observations	553	553	553

Figure 1: The Three Sustainability Goals

		Goodness of investment	
		Consequentialism Outcome-based	Deontology Principle-based
Investor Utility Function	Non-pecuniary preferences	Impact Investing	Categorical Morality
	Financial Value only	Financial Value	

Figure 2: The Three Sustainability Goals from MSCI website

Source: <https://www.msci.com/our-solutions/esg-investing/esg-ratings/what-esg-ratings-are-and-are-not>.

Understanding the most common ESG objectives

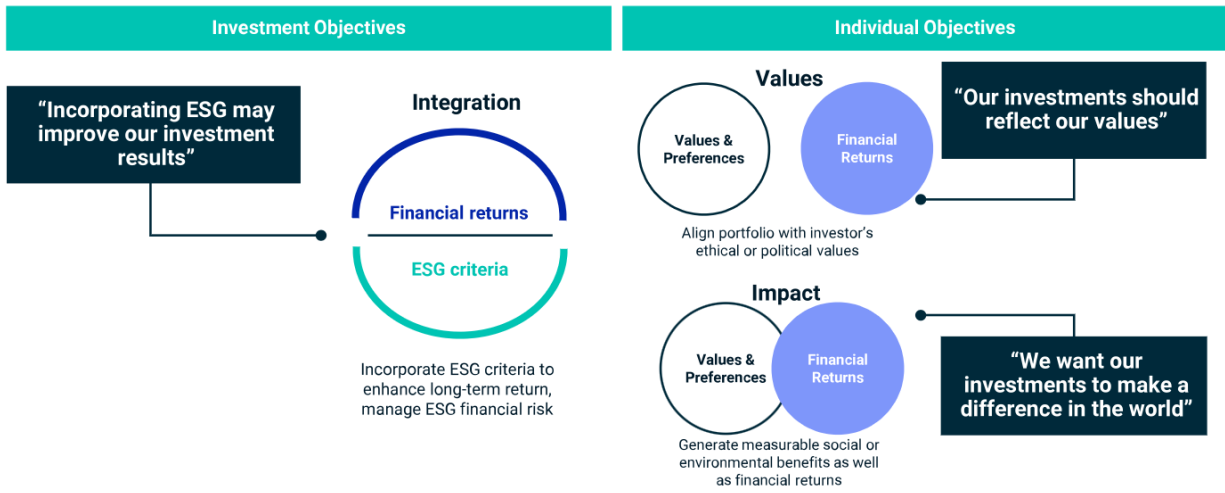


Figure 3: Dynamics Change of Funds Type

This figure displays the number of funds that belongs to each sustainable investment types. A fund is classified as “Financial” if it contains at least one Financial sentence. A fund is classified as “Moral” if it contains at least one Moral sentence. A fund is classified as “Impact” if it contains at least one Impact sentence. A fund could be classified into multiple types. BERT model classification which is trained using our manually-classified training sample are used to classify all funds’ objectives. Panel A displays the classification results are reported for all year and each year in 2018, 2019, 2020, and 2022 for all funds. Panel B displays the classification results that are reported for funds that can be linked to the CRSP mutual fund database, for funds that held at least one stock that was rated by MSCI, and for funds that can be linked to the CRSP mutual fund database and are weighted by total assets under management.

Panel A: All Funds



Panel B:



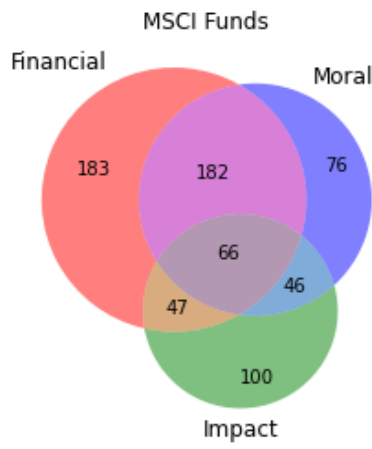
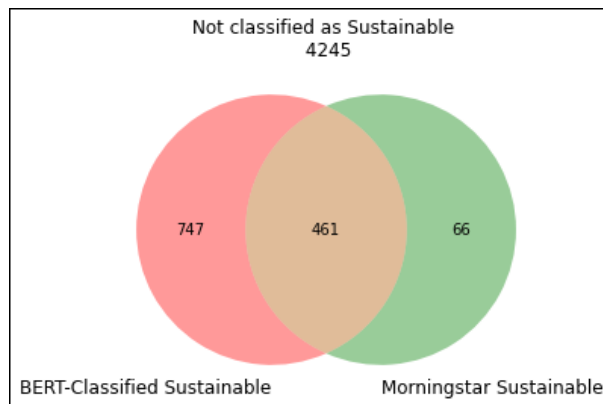


Figure 4: Expansion of the BERT-Model Analysis

This table reports the number of funds that are not classified by Morningstar as Sustainable Funds. Panel A presents the total number of funds, both those on and not on the 2022 Morningstar Lists, with ESG keywords. Among the funds with ESG keywords, we report the number of funds classified as Sustainable Funds by the BERT model. Panel B displays the number of funds that belong to each sustainable investment type for funds not classified by Morningstar as Sustainable Funds in 2022. A fund is classified as "Financial" if it contains at least one financial sentence. A fund is classified as "Moral" if it contains at least one moral sentence. A fund is classified as "Impact" if it contains at least one impact sentence. A fund could be classified into multiple types.

Panel A:



Panel B:

