Different Shades of ESG Funds

SIMONA ABIS
University of Colorado Boulder

ANDREA M. BUFFA
University of Colorado Boulder

MEHA SADASIVAM
Columbia University

April 6, 2026*

Abstract

We classify US active equity mutual funds by the objective for which they use ESG information. Impact funds, which incorporate non-pecuniary ESG preferences into their portfolio decisions, account for less than 5% of ESG assets under management; the remainder is overwhelmingly managed by opportunity funds using ESG data for financial gain. Impact funds tilt strongly toward high-externality stocks, show no greenwashing or green window dressing, and avoid firms that commit ESG violations. Opportunity funds favor high-ESG stocks only when financially material and inflate their ESG exposure before disclosure dates. Treating ESG funds as a monolithic category obscures fundamentally different economic motives.

*Contacts: `simona.abis@colorado.edu`, `buffa@colorado.edu` and `meha.sadasivam@gsb.columbia.edu`. For helpful comments and discussions, we thank Laurent Barras, Bruno Gerard, Itay Goldstein, Pedro Matos, Lakshmi Naaraayanan, Marcus Opp, Enrichetta Ravina, Daniel Schmidt, Clemens Sialm, Laura Starks, Ivo Welch, Costas Xiouros, as well as seminar and conference participants at the Federal Reserve Bank of Chicago, BI Norwegian Business School, Goethe University Frankfurt, University of Mannheim, Tilburg University, Erasmus University Rotterdam, and at the Conference in Sustainable Finance at the University of Luxembourg.

ESG encompasses a wide variety of investments and strategies. I think investors should be able to drill down to see what's under the hood of these strategies. This gets to the heart of the SEC's mission to protect investors, allowing them to allocate their capital efficiently and meet their needs.

Gary Gensler, SEC Chair, 2022 [1]

1 Introduction

Not all ESG funds are created equal. A fund that excludes tobacco stocks, one that channels capital toward firms with a positive impact on society, and one that leverages ESG data to sharpen its return forecasts are all routinely labeled "ESG." Yet they pursue fundamentally different objectives, hold different portfolios, and are likely to serve different investor clienteles. In her 2023 AFA Presidential Address, Starks (2023) argues that the central fault line in sustainable finance runs between "value" and "values": between investors who incorporate ESG information to improve financial performance and those who do so to express non-pecuniary preferences. Survey evidence from Edmans, Gosling, and Jenter (2026) confirms that this diversity of motives is pervasive in practice, with institutional investors reporting objectives that range from purely financial integration to impact-oriented mandates. Yet the empirical literature has largely treated ESG funds as a monolithic category. This is problematic: in empirical work, pooling economically distinct strategies can reverse conclusions about ESG fund behavior; and in policy, the design of disclosure mandates and taxonomies hinges on understanding what different funds actually promise and deliver. This paper provides the first comprehensive classification of US active equity mutual funds by the *objective* for which they use ESG information, and traces the consequences of this heterogeneity across fundamental dimensions of ESG fund behavior.

We classify funds into four mutually exclusive types based on a careful, manual reading of the Principal Investment Strategy (PIS) section of their SEC-filed prospectuses. *Exclusionary* funds apply negative screens to restrict their investment universe but do not otherwise use ESG criteria to guide portfolio allocation. *Impact* funds directly incorporate non-pecuniary ESG preferences into their portfolio decisions, seeking to channel capital toward firms that generate positive externalities. *Opportunity* funds use ESG information to enhance risk-adjusted returns, without expressing a preference over ESG outcomes per se. *Mention-only* funds reference ESG terminology without explaining how it is used. Each of the first three types has well-established roots in the theoretical literature on ESG investing.[2] Our contribution is to nest all of these within a unified

[1] SEC Press Release on new ESG disclosure proposed rule.

[2] Exclusionary funds correspond to the exclusionary investors in Heinkel, Kraus, and Zechner (2001) and Zerbib (2022); impact funds to the ESG-motivated (Type-M) investors in Pedersen, Fitzgibbons, and Pomorski (2021) and

empirical framework organized around a single principle: the objective for which ESG information enters the fund's optimization. Extracting this objective from prospectus text is a demanding task: the key distinction between, say, an impact fund and an opportunity fund often hinges on a few words that reveal whether ESG information serves a non-pecuniary preference or a financial objective. Automated methods, including large language models, struggle to reliably detect these nuances, which is why we classify the entire corpus manually before using the resulting labeled sample to train a language model for scalability (Internet Appendix VI).

Our first set of findings challenges the prevailing narrative about the growth of ESG investing. ESG funds grew from 4% of active equity funds in 2015 to 25% by 2023, but this expansion was driven overwhelmingly by opportunity funds, whose share of ESG fund AUM rose from 9% to 79%. Impact funds moved in the opposite direction, with their share falling from 33% to just 4%. Most opportunity funds were not created as ESG vehicles: 82% began as non-ESG funds and later revised their prospectuses to incorporate ESG language, consistent with the integration of increasingly material ESG factors into existing return-maximizing mandates. Impact funds, by contrast, are predominantly "born ESG" (70% entered the sample with an impact classification) and are the only category to experience sustained positive flow-driven AUM growth. The rapid expansion of ESG investing, at least in the active equity space, thus reflects the mainstreaming of ESG as a financial input far more than the growth of non-pecuniary demand.

Our second set of findings establishes that these different objectives translate into starkly different portfolio allocations. Using multiple measures of portfolio ESG orientation that capture not only the average rating of a fund's holdings but also which stocks it overweights, underweights, and excludes relative to its benchmark, we show that impact funds display by far the strongest tilt toward high-ESG stocks. Their portfolio ESG ratings are roughly three times higher than those obtained by pooling all ESG funds into a single category, and significantly higher than those of any other fund type. Exclusionary funds show a smaller tilt, consistent with the mechanical effect of removing low-ESG stocks from the investible universe. Opportunity and mention-only funds display only modest differences relative to non-ESG funds, suggesting that their use of ESG information does not translate into a directional preference for high-ESG holdings.

Since overall ESG ratings aggregate dimensions that serve very different economic purposes, a portfolio tilt toward high-ESG stocks need not reflect a non-pecuniary preference. To disentangle the sources of the documented tilts, we decompose firms' ESG ratings into an *ESG-externality* component, capturing a firm's impact on society and the environment, and an *ESG-materiality* component, measuring how ESG factors affect the firm's financial performance. This decompo-

the green investors in Pástor, Stambaugh, and Taylor (2021); and opportunity funds to the ESG-aware (Type-A) investors in Pedersen, Fitzgibbons, and Pomorski (2021) and the return-maximizing agents in Goldstein, Kopytov, Shen, and Xiang (2022).

sition reveals a sharp distinction: impact funds are the only type that tilts toward stocks with higher ESG-externality ratings, consistent with their non-pecuniary preference for firms that generate positive externalities. All other fund types, to the extent that they tilt at all, do so exclusively through the materiality channel. We confirm these patterns in a distributional analysis across ESG-rating quintiles, which shows that impact funds systematically shift portfolio weight from the bottom to the top quintile on both dimensions, while opportunity funds display a flat profile on externality quintiles but a positive slope on materiality quintiles when actively conditioning on ESG-materiality information. Finally, we formalize these patterns in a parsimonious theoretical framework where exclusionary, impact, and opportunity funds solve distinct portfolio optimization problems, and show that the model's numerical predictions align closely with the empirical evidence across quintiles of both rating components.

Our third set of findings exploits the granularity of our classification to deliver novel insights on three questions of direct relevance for investors and regulators. First, we ask whether ESG funds greenwash. We argue that greenwashing, properly defined, requires a gap between what a fund *promises* and what it *delivers*, and therefore applies most directly to impact funds. We flag funds as potential greenwashers if their portfolio ESG ratings persistently fall in the bottom of the non-ESG distribution. Impact funds have the lowest greenwashing rates across nearly all criteria, indicating that the vast majority of funds claiming non-pecuniary ESG objectives hold portfolios consistent with those claims.

Second, we investigate ESG window dressing, the strategic inflation of ESG exposure around quarterly disclosure dates. Following Parise and Rubin (2025), we estimate ESG factor loadings in control, pre-disclosure, and post-disclosure windows, and decompose the aggregate result by fund type. Window dressing is entirely absent among impact and exclusionary funds and concentrated in opportunity and mention-only funds. Opportunity funds exhibit a textbook inverted-V pattern: their ESG loading nearly triples before disclosure and reverts precisely to baseline afterward. Impact funds, by contrast, maintain the highest and most stable ESG loading across all windows. The aggregate ESG window dressing documented by Parise and Rubin (2025) is, in our sample, an average of a null effect for funds with binding ESG mandates and a large effect for those without.

Third, we ask whether impact funds are simply chasing high ESG ratings or whether they also avoid firms that cause actual harm. ESG ratings measure how well firms manage ESG issues, but a firm can score well on ratings while still committing violations, and in our data the two dimensions are indeed largely uncorrelated. To provide an independent test, we construct a novel dataset of ESG violations from Violation Tracker, covering legally adjudicated instances of environmental contamination, labor abuses, fraud, and other regulatory breaches, and classify each offense as ESG-externality or ESG-materiality related, mirroring our rating decomposition. Impact funds display the largest reduction in violation exposure (16% less than non-ESG funds), and their avoidance is disproportionately concentrated in externality-related offenses. Opportunity funds also avoid violator stocks, but their avoidance is more evenly distributed across externality and materiality violations, consistent with treating violation history as a signal of financial risk rather than as a screen motivated by social concerns.

This convergence of evidence reinforces the paper's central finding: the different shades of ESG funds matter. Our classification captures economically meaningful differences in investment objectives that manifest consistently across different aspects of fund behavior, and unlocks novel insights about ESG investing that are invisible when funds are treated as a homogeneous group.

Related literature. Our paper contributes to several streams of research on ESG investing.

Theory of ESG investing. A large and rapidly growing theoretical literature studies how non-pecuniary ESG preferences and ESG information affect portfolio allocation, asset prices, and real outcomes. Pástor, Stambaugh, and Taylor (2021) and Pedersen, Fitzgibbons, and Pomorski (2021) develop equilibrium models in which green or ESG-motivated investors coexist with conventional ones, generating a "greenium" and affecting price informativeness. Avramov, Lioui, Liu, and Tarelli (2024) extend this to a dynamic setting with ESG uncertainty. Goldstein, Kopytov, Shen, and Xiang (2022) study how the coexistence of profit-motivated and responsible investors affects price informativeness across multiple equilibria. Heinkel, Kraus, and Zechner (2001) and Zerbib (2022) model how exclusionary strategies affect firms' cost of capital and incentives for green investment. On the normative side, Oehmke and Opp (2024) develop a theory of socially responsible investment and derive a micro-founded ESG investment criterion, while Inderst and Opp (2025) analyze whether a taxonomy for sustainable investment products is warranted in the presence of conventional environmental policy tools. Berk and van Binsbergen (2024) and Broccardo, Hart, and Zingales (2022) compare the real effects of exit (divestment) versus voice (engagement), arguing that direct engagement may be more effective at improving firms' ESG practices. Dasgupta, Jenter, Mathews, and Voss (2026) study the viability of socially responsible funds when investors can free-ride, showing that exclusionary funds endogenously create the conditions under which impact funds become viable through a symbiotic buy-low-sell-high mechanism. Our paper provides the first comprehensive empirical mapping between the investor types featured in these models and the actual behavior of mutual funds, using a classification grounded in the same economic primitives (objectives and information) that organize the theoretical literature.

Identifying and classifying ESG funds. Prior empirical work has relied on proxies that do not capture the heterogeneity in ESG fund objectives: fund names (Cremers, Riley, and Zambrana, 2024; Farroukh, Harford, and Shin, 2023; van der Beck, 2023), Morningstar sustainability classifications (Baker, Egan, and Sarkar, 2022; Raghunandan and Rajgopal, 2022), PRI signatory status (Gibson Brandon, Glossner, Krueger, Matos, and Steffen, 2022; Ceccarelli, Glossner, and Homanen, 2022), and holdings-based measures (Lowry, Wang, and Wei, 2025; Ghoul and Karoui, 2022). These approaches identify whether a fund is ESG but do not distinguish the objective for which ESG information is used. Two recent papers also analyze fund prospectuses: Andrikogiannopoulou, Krueger, Mitali, and Papakonstantinou (2022) use text analysis to construct a green intensity measure and a greenwashing indicator, distinguishing ESG from non-ESG funds but not among ESG types; Frankel, Martin, Wang, and Yu (2024) classify funds into ESG-integration and ESG-focused categories following the SEC's proposed disclosure framework, which focuses on the *importance* of ESG information rather than the *objective* for which it is used. In contemporaneous work, Yang and Yasuda (2025) use a BERT model trained on a hand-labeled subset of prospectuses to classify the funds labeled as sustainable by Morningstar into financial, moral, and impact categories, and study the link between fund intent and real-world outcomes. Our approach is complementary: we classify the full corpus manually, which allows us to capture fine distinctions that require close reading of the text, and yields a comprehensive training sample for language models (described in the Internet Appendix VI). In a related vein, Lowry, Wang, and Wei (2025) distinguish "committed" from other ESG funds based on their incentives to engage, and find that only committed funds contribute to real ESG improvements. Their classification is based on portfolio characteristics, while ours is based on disclosed objectives; the two approaches are complementary and yield reinforcing conclusions about the importance of distinguishing among ESG fund types.

ESG fund holdings. The empirical evidence on whether ESG funds hold higher-ESG stocks is mixed. Pastor, Stambaugh, and Taylor (2024), Lowry, Wang, and Wei (2025), and Raghunandan and Rajgopal (2022) find that ESG funds do tilt toward higher-ESG holdings, whereas Gibson Brandon, Glossner, Krueger, Matos, and Steffen (2022), Farroukh, Harford, and Shin (2023), and Andrikogiannopoulou, Krueger, Mitali, and Papakonstantinou (2022) find little difference between ESG and non-ESG funds in the US. A key factor behind these conflicting results is that each study uses a different method to identify ESG funds, leading to samples with different compositions of fund types. Our granular classification resolves this tension: once ESG funds are decomposed by objective, impact funds display large and significant tilts while opportunity and mention-only funds do not, and the aggregate result depends on the relative weight of each type in the sample.

Greenwashing and window dressing. Several papers examine whether ESG funds deliver on their promises. Andrikogiannopoulou, Krueger, Mitali, and Papakonstantinou (2022) develop a text-based greenwashing measure and find that investors cannot distinguish genuinely green funds from greenwashers. Parise and Rubin (2025) document aggregate evidence of green window dressing around disclosure dates. Our paper contributes to this literature by showing that greenwashing and window dressing are not general properties of ESG funds but are concentrated in specific fund types with the weakest ESG mandates. Impact funds, for which the greenwashing concern is most economically relevant, display the lowest incidence of such behavior. More broadly, our findings suggest that the widespread perception of greenwashing in the ESG fund industry may be an artifact of pooling fund types with fundamentally different objectives.

Paper layout. Section 2 describes the data, presents our classification methodology, and compares it with alternative approaches. Section 3 documents the ESG investment landscape, including the growth, composition, and capital flows of ESG funds by type. Section 4 analyzes the investment behavior of ESG fund types through portfolio-level ESG ratings, the externality-materiality decomposition, the distribution of holdings across ESG-rating quintiles, and a theoretical framework that formalizes the predictions. Section 5 presents novel insights on greenwashing, green window dressing, and ESG violations. Section 6 concludes.

2 Classification of ESG Funds

We classify active equity mutual funds in the US based on how ESG information enters their investment process. Our categorization distinguishes among three core strategies—exclusionary, impact, and opportunity—that differ not in *whether* a fund considers ESG factors, but in the *objective* for which it does so. This section describes the data, the classification methodology, and its validation against alternative approaches.

2.1 Data

Our analysis covers the period from March 2015 to June 2023 and draws on four main data sources: fund characteristics and returns, portfolio holdings, fund prospectuses, and firm-level ESG ratings.

Fund characteristics and returns. We use the CRSP Survivorship-Bias-Free Mutual Fund dataset to identify our sample and extract information on fund characteristics and returns. Starting from the universe of equity funds, we exclude international funds, sector funds, index funds, and variable annuities. We drop observations preceding each fund's first offer date to account for incubation bias (Evans, 2010), funds with less than $5 million in total net assets (TNA) (Kacperczyk, Sialm, and Zheng, 2008), and funds with fewer than 12 months of observations. For each period, we aggregate observations across share classes: we retain the first offer date of the oldest share class, sum TNA across all share classes, and compute TNA-weighted averages of all other variables (fees, returns, turnover).[3] Finally, we follow the data appendix of Pástor, Stambaugh, and Taylor (2015) to match our sample to the Morningstar Mutual Fund dataset, from which we obtain each fund's Morningstar category, star rating, and sustainability ratings. We also use the CRSP monthly security database and Fama–French factors.

Fund holdings. We obtain portfolio holdings from the CRSP Mutual Fund Holdings dataset. Based on holdings data, we apply two additional filters: we exclude funds holding fewer than 10 stocks and funds holding, on average, less than 80% of their assets (excluding cash) in common stock (Kacperczyk, Sialm, and Zheng, 2008). When holding observations are missing or available only at quarterly frequency, we forward-fill them to monthly frequency.

Fund prospectuses. We download the Mutual Fund Prospectus Risk/Return Summary Data Set from the SEC's EDGAR platform. We extract the Principal Investment Strategy (PIS) section of each prospectus and match it to the CRSP Mutual Fund dataset using the link between SEC CIK identifiers and fund identifiers. Since any material change to a fund's management must be reported to both the SEC and fund investors, for months in which a PIS section is unavailable, we fill it with the most recently available version.

ESG ratings. We acquire firm-level ESG ratings from MSCI, Sustainalytics, and Refinitiv. We use both overall ratings and their most disaggregated sub-components. To harmonize the data across providers, we forward-fill all scores up to 12 months and re-scale all sub-components to a $[0, 100]$ range.[4] We merge ESG scores to fund holdings using a mapping between each provider's firm identifiers and a CUSIP-to-PERMCO link. MSCI ratings offer the broadest coverage for our sample period, so we use them in our baseline analysis and employ a combined rating across providers in robustness tests reported in the Internet Appendix IV.

2.2 Classification Methodology

Our classification leverages the narrative strategy descriptions that funds are required to provide in the PIS section of their prospectuses. In this section, funds describe their main investment

[3] We construct a comprehensive share-class-to-fund mapping using the CRSP Class Group identifier (`crsp_cl_grp`), the WFICN identifier in MFLinks, and fund names. This approach ensures broader coverage than relying on MFLinks alone, which excludes many recently launched funds (Shive and Yun, 2013; Zhu, 2020).

[4] MSCI ratings are available monthly and each sub-component is scaled on a $[0, 10]$ range. Sustainalytics ratings are available monthly with most sub-components scaled on a $[0, 100]$ range, except for a few event-based sub-components scaled on a $[0, 5]$ range. Refinitiv scores are available annually; most sub-components are scaled on $[0, 1]$ or $[0, 100]$, with a few having an arbitrary range that we winsorize at the 0.01% level before re-scaling.

methodology, including the types of securities they hold and the primary criteria used to select them. If a fund incorporates ESG considerations into its asset selection process, this is where it must disclose that information. The narrative nature of these disclosures allows us to infer the objective with which ESG information enters a fund's strategy, which in turn allows us to map each fund to a specific ESG type.

2.2.1 ESG fund types

Our categorization builds on a rich theoretical literature on ESG investing. Each of the fund types we define has well-established roots in existing models; however, these models typically feature only a subset of the investor types that coexist in practice. Our contribution is to nest all of them within a unified framework organized around a single guiding principle: the objective for which ESG information enters a fund's portfolio optimization.

Most theoretical models of ESG investing also feature a baseline group of investors who do not incorporate ESG information in any capacity, e.g., the ESG-unaware (Type-U) investors in Pedersen, Fitzgibbons, and Pomorski (2021), or the conventional investors in Pástor, Stambaugh, and Taylor (2021), Avramov, Lioui, Liu, and Tarelli (2024), and Zerbib (2022). In our empirical setting, these correspond to *non-ESG* funds, which we identify through the dictionary-based filtering described in Section 2.2.2 below. Among those funds whose prospectuses do reference ESG information, we distinguish three core types—exclusionary, impact, and opportunity—based on the objective for which that information is used, plus a residual category.

- *Exclusionary* funds apply negative screens based on ESG criteria to restrict their investment universe, but do not otherwise use ESG information to guide portfolio allocation within the investible set. For instance, a fund may exclude companies involved in tobacco, firearms, or thermal coal, and then optimize its portfolio purely on the basis of risk and return. Exclusionary strategies have been extensively studied in the sin stocks literature (Hong and Kacperczyk, 2009), and are explicitly modeled by Heinkel, Kraus, and Zechner (2001), Zerbib (2022), and Broccardo, Hart, and Zingales (2022). As an example, one fund in our sample states that it "*generally avoid[s] investments in issuers [it] deem[s] to have significant alcohol, gaming or tobacco business*."
- *Impact* funds directly incorporate non-pecuniary ESG preferences into their portfolio allocation decisions, alongside risk-return considerations. These funds seek to channel capital toward companies that score well on ESG dimensions—not because doing so improves financial performance, but because they value those characteristics in their own right. This type of investor is central to most theoretical frameworks on ESG investing: they correspond

to the ESG-motivated (Type-M) investors in Pedersen, Fitzgibbons, and Pomorski (2021), to the green investors in Pástor, Stambaugh, and Taylor (2021) and Avramov, Lioui, Liu, and Tarelli (2024), and to the responsible investors in Goldstein, Kopytov, Shen, and Xiang (2022) and Zerbib (2022). A representative disclosure reads: "*the fund will invest in companies that [the adviser] believes have strong environmental and social profiles [...] impact investors seek to use their investments to create a more fair and sustainable world.*"

- *Opportunity* funds use ESG information to enhance risk-adjusted financial performance, without expressing a non-pecuniary preference over ESG characteristics. They treat ESG factors as material inputs into their assessment of a company's future risks and returns. In existing theoretical models, these investors are typically aware of how ESG considerations affect asset prices and optimize accordingly: they correspond to the ESG-aware (Type-A) investors in Pedersen, Fitzgibbons, and Pomorski (2021), and to the return-maximizing investors in Goldstein, Kopytov, Shen, and Xiang (2022) and Heinkel, Kraus, and Zechner (2001) who exploit the price effects generated by the presence of exclusionary or impact investors. A typical example: "*ESG integration is driven by taking into account material ESG risks which could impact investment returns, rather than being driven by specific ethical principles or norms.*"

- *Mention-only* is a residual empirical category for funds whose prospectuses contain a brief reference to ESG terminology without explaining how it is used. Unlike the three core types above, this category does not have a theoretical counterpart—it arises purely from the limitations of the disclosure. These funds cannot be assigned to any of the core categories based on their prospectuses alone; their true intent is an empirical question that we explore in subsequent sections.

While each of the theoretical models cited above features a subset of these investor types, no single framework encompasses all of them. Our classification is the first to provide a comprehensive empirical counterpart for the full spectrum of ESG investment objectives identified in the theoretical literature. A more formal characterization of how each strategy maps to a fund's optimization problem is provided in Section 4.4. Internet Appendix II reports additional examples of PIS sections assigned to each category.

2.2.2 Implementation

Our classification proceeds in three stages.

Stage 1: Dictionary-based filtering. We first separate candidate ESG funds from definite non-ESG funds by selecting all prospectuses that mention ESG-related terminology. We employ an extensive dictionary (reported in Internet Appendix I) that includes generic ESG-related terms (e.g., "social responsibility," "environmental"), common acronyms (e.g., "ESG," "SRI," "SDG"), and specific ESG-related issues (e.g., "climate change," "child labor," "firearms"). The dictionary is designed to be over-inclusive: its sole purpose is to ensure that no ESG-related fund is inadvertently excluded from further analysis. The underlying assumption is that a fund implementing any ESG-related strategy would mention at least one of the terms in the dictionary in its prospectus.

Stage 2: Manual classification. We then manually read each filtered PIS section and assign it to one or more of the following labels: exclusionary, impact, opportunity, and mention-only. Labels are not required to be mutually exclusive at this stage. For instance, a fund may disclose both an impact objective and the use of exclusionary screens; a reader would assign both the impact and exclusionary labels.

To ensure reliability, we split the universe of unique PIS sections into six random subsets, each assigned to two independent readers (three readers in total, the co-authors of this paper). When both readers agree on a classification, that label is automatically assigned. In cases of disagreement, the section is referred to a third reader for adjudication.[5]

Stage 3: Pecking order. To assign each fund-month observation to a single mutually exclusive category, we apply the following pecking order. First, if a PIS section receives an *impact* label, it is assigned to the impact category regardless of any additional labels. This reflects the fact that impact funds may also employ exclusions or consider how ESG factors affect financial performance, but their defining feature is their non-pecuniary preference.[6] Second, conditional on not being classified as impact, if a section receives an *opportunity* label, it is assigned to the opportunity category regardless of whether it also mentions exclusions.[7] Third, sections labeled only as *exclusionary* are assigned to the exclusionary category. Fourth, remaining sections with a *mention-only* label are assigned to the mention-only category. In total, our classification produces four mutually exclusive fund types—exclusionary (Exc), impact (Imp), opportunity (Opp), and mention-only (Men).

The advantage of our manual classification is that it allows us to capture subtle differences in intent that automated methods struggle to detect. Indeed, the key distinction between impact

[5]We also record the specific text excerpts that motivate each classification decision. These annotations serve as training data for the large language model described in the Internet Appendix VI.

[6]Among sections labeled as impact, 65% also mention exclusions and 24% mention opportunity motives.

[7]Among sections labeled as opportunity (and not impact), 17% also mention exclusions.

and opportunity funds often hinges on a few words—whether a fund uses ESG information "*to create a more fair and sustainable world*" or "*to better assess strategic business issues that may impact the performance of a company.*" These differences are small in the text but, as we show in subsequent sections, translate into starkly different portfolio allocation behavior. We initially considered alternative approaches, including training large language models (LLMs) on a small set of manually labeled examples, with limited success.[8] Having classified the full corpus, we now have an extensive sample of manually labeled text that can be used to train a language model for future applications; we describe such a model in the Internet Appendix VI.

2.3 Comparison with Alternative Methodologies

To highlight the distinctive features of our classification, we compare it to three popular approaches used in the academic literature to identify ESG funds. Unlike ours, these methods do not distinguish among different ESG objectives; instead, they aim to separate ESG from non-ESG funds based on coarser criteria. As we show below, this coarseness leads to substantial misclassification—both by missing many funds with genuine ESG objectives and by labeling as sustainable funds that use ESG information purely for financial purposes.

The first approach searches for sustainability-related terms in the fund's name ("Name ESG"), following Cremers, Riley, and Zambrana (2024), Farroukh, Harford, and Shin (2023), and van der Beck (2023). The rationale is that if a fund brands itself with ESG-related terminology, this must be a central component of its strategy. The second approach uses Morningstar's "Sustainable Investment Overall" classification ("Mstar ESG"), which identifies funds for which one or more sustainable investing approaches are central to the investment process, based on regulatory filings (Morningstar, 2022). This variable is populated starting in 2020, limiting the overlap with our sample. The third approach classifies a fund as sustainable if its family is a signatory to the Principles for Responsible Investment ("PRI ESG"), following Ceccarelli, Glossner, and Homanen (2022), Gibson Brandon, Glossner, Krueger, Matos, and Steffen (2022), and Kim and Yoon (2023).

Figure 1 displays, for each ESG type in our classification, the percentage of fund-month observations that would be classified as sustainable under each of the three alternative methods. Three key differences emerge.

First, both fund-level methods (Name ESG and Mstar ESG) correlate most strongly with our impact category: funds with non-pecuniary ESG objectives are the most likely to brand themselves

[8]Transformer-based models require substantial training data to reliably distinguish among six labels; in our setting, the required sample size approached the full corpus. Generative AI models required fewer examples but performed poorly, as ESG terminology carries different meanings across the sources on which these models are trained—precisely the ambiguity that our classification seeks to resolve.

accordingly or to be recognized as sustainable by Morningstar. However, these methods miss a large share of funds with genuine impact objectives. The Name ESG approach identifies only 18% of impact funds and 3% of exclusionary ones, while Mstar ESG identifies 83% and 19%, respectively (the latter over a shorter sample period). Funds that pursue ESG goals without prominently advertising them in their names, or that do so through exclusions rather than positive screening, are largely invisible to these approaches.

Second, a substantial fraction of funds that we classify as opportunity also appear as sustainable under the alternative methods: 19% under Mstar ESG and 4% under Name ESG. This reflects a fundamental difference in what these methods capture. Both Name ESG and Mstar ESG focus on how *important* ESG factors are to a fund's strategy, whereas our method focuses on the *objective* for which ESG information is used. An opportunity fund for which ESG analysis is a core part of its risk-return assessment may be classified as sustainable by Morningstar, even though it expresses no non-pecuniary preference over ESG outcomes. Pooling such funds together with impact funds, as these coarser methods implicitly do, conflates two economically distinct investment strategies that, as we demonstrate in subsequent sections, produce very different portfolio allocations.

Third, the PRI signatory approach classifies a large fraction of funds across *all* categories as sustainable, indicating that PRI signatory families offer a wide variety of fund types (including non-ESG ones), rather than specializing in sustainable strategies. This makes PRI status a particularly noisy proxy for ESG investment objectives.

Table 1 quantifies these discrepancies by reporting Type 1 (false positive) and Type 2 (false negative) error rates for each alternative method, benchmarked against our classification.[9] The Name ESG approach has a high false negative rate (82–89%) but a more moderate false positive rate (36–46%). The Morningstar approach produces a lower false negative rate (20–48%) but a higher false positive rate (41–50%). The PRI approach generates very high error rates in both dimensions (52–64% false negatives, 95–97% false positives). These error rates underscore that existing methods are not simply noisier versions of our classification—they capture fundamentally different dimensions of ESG investing.

Our classification also differs from the SEC's proposed disclosure rule, which distinguishes between "ESG Integration" and "ESG Focused" funds based on the importance of ESG factors, rather than the objective for which they are used.[10] By providing a finer partition organized around economic objectives, our method recovers distinctions that are invisible to approaches

[9]We define as "correctly classified as sustainable" any fund that our method categorizes as impact (Panel A) or as impact or exclusionary (Panel B).

[10]Under the SEC's proposal, "ESG Integration" funds consider ESG factors jointly with other factors, while "ESG Focused" funds treat ESG as the primary criterion in asset selection.

based solely on the prominence of ESG in a fund's strategy. As we show in the following sections, these distinctions map to meaningfully different investment behaviors.

3 The ESG Investment Landscape

Having classified active equity mutual funds by the objective for which they use ESG information, we now document key trends in the allocation of capital to ESG funds. Our granular classification reveals that the widely documented growth in ESG investing masks substantial heterogeneity: the composition of ESG funds has shifted dramatically over our sample period, with implications for how we interpret the rise of sustainable finance.

3.1 Growth of ESG Funds

The top-left panel of Figure 2 shows the rapid increase in both the number of ESG funds and the capital they manage over the period 2015–2023. ESG funds represented 4% of all active equity mutual funds in 2015, managing roughly 1% of the capital invested in this space. By 2023, they accounted for 25% of active equity funds and 20% of total assets under management—a sixfold increase in less than a decade. The top-right panel shows that this trend is not confined to a few large fund families: the share of families offering at least one ESG fund grew from 11% to 39% over the same period.

The common perception is that this growth reflects a surge in demand for funds that allocate capital based on non-pecuniary ESG preferences. Our classification reveals a starkly different picture. The bottom panels of Figure 2 decompose the growth of ESG funds by category, in terms of both the number of funds (bottom-left) and assets under management (bottom-right). The number of impact funds increased modestly, from 23 to 46, while the number of opportunity funds grew from just 5 to 239. As a result, funds with an impact motive went from representing 38% of all ESG funds in 2015 to only 13% by 2023. Conversely, funds with an opportunity motive went from 9% to 66%. The number of exclusionary funds, instead, remained essentially flat throughout the period. The shift is even more pronounced when measured by assets under management. Funds with an opportunity motive went from managing 9% of total ESG fund AUM to 79%, while funds with an impact motive went from 33% to just 4%. Exclusionary funds, despite their flat count, saw their share decline from 57% of AUM (49% of ESG funds) in 2015 to 2% of AUM (9% of funds) by 2023, simply because other categories grew around them. Mention-only funds also grew, moving from 1% of AUM (5% of funds) to 14% of AUM (12% of funds).

These patterns suggest that the growth of ESG investing, at least within the active equity mutual fund space, has been driven predominantly by funds that use ESG information to improve risk-adjusted returns, rather than by funds that pursue non-pecuniary ESG objectives. This distinction has direct implications for how we interpret the demand for sustainable investment products and for the potential real-economy effects of ESG capital flows.

3.2 New Funds versus Existing Funds

A natural question is whether the increase in ESG funds comes primarily from newly created funds or from existing funds that modify their prospectuses to incorporate ESG information. Figure 3 addresses this question by displaying, for each ESG type, the percentage of funds that were classified as such when they first appeared in our dataset versus those that transitioned from a different classification.

A sharp distinction emerges between fund types that are "born ESG" and those that become ESG over time. Among exclusionary funds, 76% were classified as exclusionary from the start; similarly, 70% of impact funds entered our sample with an impact classification. In contrast, only 11% of opportunity funds and 16% of mention-only funds were classified as such at their first appearance. The vast majority of opportunity and mention-only funds—82% and 84%, respectively—were initially categorized as non-ESG, having later revised their prospectuses to incorporate ESG language. Transitions between different ESG types are rare: the off-diagonal entries outside the first column are negligible.

These findings are consistent with the view that opportunity funds are not fundamentally changing their investment objectives; rather, they are integrating ESG dimensions, which have become increasingly material for assessing risks and returns, into their existing return-maximizing mandates. Crucially, these funds explicitly disclose their intent to use ESG information as a tool for improving financial performance. They do not claim to be impact funds that subsequently fall short of their promises. The distinction between a fund that *becomes* ESG by adding material considerations to an existing strategy and one that is *born* ESG with a non-pecuniary mandate is important for evaluating claims of greenwashing, a point we return to in Section 5.1.

Regarding mention-only funds, the predominance of transitions from non-ESG status makes it more likely that they behave either as opportunity funds—integrating ESG considerations into their investment process without making this explicit—or as non-ESG funds that reference ESG terminology without meaningfully acting on it.

3.3 Capital Flows

The growth in ESG fund AUM documented in Figure 2 reflects three sources: funds switching type, investment returns, and net investor flows. Having shown in Figure 3 that type-switching is a substantial driver—particularly for opportunity and mention-only funds—we now isolate the contribution of net flows.

Figure 4 displays the cumulative percentage change in AUM attributable solely to net flows, by ESG type. This measure captures the growth in AUM for each category that would have been realized absent returns or type-switching, and thus directly reflects investor demand.

As a benchmark, non-ESG funds experience a steady decline in cumulative flow-driven growth over our sample period, consistent with the broader trend of capital migrating toward passive investment vehicles. Exclusionary and opportunity funds follow a similar trajectory, suggesting that investor demand for these categories does not differ markedly from that for non-ESG funds over the full sample period. Impact funds stand out as the only category to experience a sustained increase in flow-driven AUM growth for nearly the entire sample period, ending with a net positive cumulative growth by 2023. This indicates that, despite their small and declining share of the ESG fund universe, impact funds continue to attract disproportionate investor demand, consistent with a genuine appetite for non-pecuniary ESG investing among a subset of investors. Finally, mention-only funds experienced an early increase in flow-driven growth (2015–2018), followed by a sharper decline than non-ESG funds.[11]

Taken together, these flow patterns reinforce two conclusions. First, the rapid growth in ESG fund AUM is not primarily driven by investor demand for ESG products; for most categories, net flows are comparable to, or worse than, those of non-ESG funds. The growth instead reflects the reclassification of existing non-ESG funds, as documented in the previous subsection. Second, among the categories that do attract differential investor demand, impact funds are the clearest beneficiaries, suggesting that investors who actively seek out ESG funds tend to favor those with explicit non-pecuniary objectives.

4 Investment Behavior of ESG Funds

The previous section documented that different types of ESG funds vary substantially in their prevalence, growth trajectories, and ability to attract investor flows. In this section, we ask whether these different types also differ in how they actually invest. We study the relationship

[11]The large amplitude of changes for mention-only funds reflects the small number of funds in this category during the early years of our sample (3–7 funds in 2015–2018).

between our ESG fund classification and the ESG ratings of fund holdings, progressively refining the analysis from overall ESG ratings to their externality and materiality components, and from average portfolio ratings to the full distribution of holdings across ESG quintiles.

4.1 Fund-Level ESG Ratings

We begin by investigating how the average ESG ratings at the fund portfolio level vary across the identified categories of ESG funds. To this end, we aggregate monthly firm-level ESG ratings at the fund portfolio level, using different weighting schemes to isolate distinct aspects of investment behavior.

Denoting firm i's ESG rating in month t by s_{it}^{ESG}, we first cross-sectionally standardize each rating within a given month to make them comparable across time periods, obtaining the standardized rating

$$z_{it}^{ESG} \equiv \frac{s_{it}^{ESG} - E_t^{cs}[s_{it}^{ESG}]}{\sqrt{Var_t^{cs}[s_{it}^{ESG}]}}, \tag{1}$$

and then construct four measures of fund-level ESG ratings in a given month:

- *Value-weighted* (Z_{jt}^{VW}): standardized ESG ratings are TNA-weighted based on all stocks in a fund's portfolio,

$$Z_{jt}^{VW} = \frac{\sum_i^{N_{jt}} w_{jit} \cdot z_{it}^{ESG}}{\sum_i^{N_{jt}} w_{jit}}, \tag{2}$$

  where w_{jit} corresponds to the weight of asset i in the portfolio of fund j in month t.

- *Over-weighted* (Z_{jt}^{b+}): standardized ESG ratings are weighted by the excess portfolio weights of the subset of stocks that the fund overweights relative to its benchmark,

$$Z_{jt}^{b+} = \frac{\sum_i^{M_{jt}} \left(w_{jit} - w_{jit}^b\right)^+ \cdot z_{it}^{ESG}}{\sum_i^{M_{jt}} \left(w_{jit} - w_{jit}^b\right)^+}, \tag{3}$$

  where $(w_{jit} - w_{jit}^b)^+$ is the difference in weight of asset i in the portfolio of fund j and that of its benchmark in month t if positive, and is zero otherwise. M_{jt} is the total number of assets in the union of the portfolio of fund j and its benchmark. We follow Cremers and Petajisto (2009) and identify a fund's benchmark as the index that minimizes the fund's active share in a given month.

- *Under-held* ($Z_{jt}^{b^-}$): standardized ESG ratings are weighted by the excess portfolio weights of the subset of stocks that the fund underweights relative to its benchmark but for which it holds positive weights,

$$Z_{jt}^{b^-} = \frac{\sum_i^{M_{jt}} \left(w_{jit} - w_{jit}^b\right)^- \mathbf{1}_{\{w_{jit}>0\}} \cdot z_{it}^{ESG}}{\sum_i^{M_{jt}} \left(w_{jit} - w_{jit}^b\right)^- \mathbf{1}_{\{w_{jit}>0\}}}, \tag{4}$$

where $\left(w_{jit} - w_{jit}^b\right)^-$ is the difference in weight of asset i in the portfolio of fund j and that of its benchmark in month t if negative, and is zero otherwise.

- *Not-held* ($Z_{jt}^{b_0^-}$): standardized ESG ratings are weighted by the excess portfolio weights of the subset of benchmark stocks that the fund does not hold at all,

$$Z_{jt}^{b_0^-} = \frac{\sum_i^{M_{jt}} \left(w_{jit} - w_{jit}^b\right)^- \mathbf{1}_{\{w_{jit}=0\}} \cdot z_{it}^{ESG}}{\sum_i^{M_{jt}} \left(w_{jit} - w_{jit}^b\right)^- \mathbf{1}_{\{w_{jit}=0\}}}. \tag{5}$$

To analyze mean differences in these four measures across ESG fund types, we estimate the following panel regression:

$$Z_{jt}^p = \sum_{\theta\in\Theta} \beta_\theta \, \mathbf{1}_{\{\theta_{jt}=\theta\}} + \eta' X_{jt} + \iota_{MStar\times t} + \varepsilon_{jt}, \tag{6}$$

where $p \in \{VW, b^+, b^-, b_0^-\}$ denotes one of the four portfolio rating measures; θ_{jt} denotes fund j's ESG type in month t and the set $\Theta = \{Exc, Imp, Opp, Men\}$ collects all ESG fund types; X_{jt} is a vector of fund-level controls (log age, log assets, expenses, turnover, fund flows, fund flow volatility, and fund betas with respect to Fama–French–Carhart six factors); and $\iota_{MStar\times t}$ represents Morningstar category $\times$ month fixed effects. Non-ESG funds serve as the excluded category, and standard errors are clustered by fund and month. The fixed-effect structure ensures that we compare funds within the same Morningstar category in the same month, effectively controlling for time-varying differences in investment styles. As an initial investigation, we also estimate the specification with a single ESG dummy replacing the category-specific indicators. In our baseline specification, we use MSCI ESG ratings; Internet Appendix IV reports robustness results using a combined measure that aggregates MSCI, Refinitiv, and Sustainalytics ratings.

Table 2 reports the estimates of these panel regressions. When comparing ESG funds as a group against non-ESG funds (top panel), we find that ESG funds tend to hold long positions in, and overweight relative to their benchmarks, stocks with higher ESG ratings. Specifically, the value-weighted and over-weighted portfolio ESG ratings of ESG funds are on average 39% (0.087/0.225) and 46% (0.095/0.207) higher than those of non-ESG funds, respectively, both significant at the

1% level. ESG funds also tend to exclude from their portfolios stocks with lower ESG ratings, although the economic magnitude is smaller: 8.4% (-0.024/0.286) lower, still significant at the 1% level. There is no strong difference in ESG ratings for stocks that funds hold but underweight relative to their benchmarks. These results show that, as an overall group, ESG funds tilt their portfolios toward higher-ESG stocks and, to some extent, away from lower-ESG ones. These average results, however, mask important heterogeneity across ESG fund types. The advantage of our granular classification is that it allows us to identify which strategies drive the documented average behavior and to uncover novel patterns of ESG investing.

Impact funds display by far the strongest tilt toward high-ESG stocks. Their value-weighted and over-weighted portfolio ESG ratings are on average 108% (0.244/0.225) and 127% (0.264/0.207) higher than those of comparable non-ESG funds, respectively—almost three times larger than the coefficients obtained for ESG funds as a whole. Moreover, the benchmark stocks that impact funds choose not to hold have significantly lower ESG ratings than those excluded by non-ESG funds: 28% lower (-0.079/0.286), at the 1% significance level. The ESG rating of stocks held but underweighted is not statistically different from that of non-ESG funds. Taken together, these results indicate that impact funds tilt their portfolios toward stocks with significantly higher ESG ratings than both non-ESG funds and the average ESG fund, consistent with a directional preference for high-ESG holdings.

Exclusionary funds also display a positive tilt, but of smaller magnitude. Their value-weighted and over-weighted portfolio ESG ratings are 28% (0.064/0.225) and 32% (0.066/0.207) higher than non-ESG funds, respectively, both significant at the 1% level. Interestingly, the stocks that they hold but underweight also have a higher ESG rating than those of non-ESG funds (36% higher, at 5% significance). This pattern is consistent with the mechanics of an exclusionary strategy: removing low-ESG stocks from the investment universe raises the average ESG rating of the entire investible set, so that both overweighted and underweighted positions have higher ratings than those of unconstrained non-ESG funds. Crucially, the magnitude of the tilt is significantly smaller than that of impact funds, since portfolio selection within the investible set is not guided by ESG preferences.

Opportunity funds display a different pattern. Their value-weighted and over-weighted portfolio ESG ratings are only slightly higher than those of non-ESG funds—21% (0.047/0.225) and 25% (0.052/0.207), respectively—and their under-held and not-held portfolios show no significant differences. This is consistent with their disclosed strategy: opportunity funds do not have a directional preference over ESG ratings. They use ESG information to identify risk-return opportunities, which may involve overweighting some high-ESG stocks while underweighting others, leading to an ambiguous net effect. The small positive tilt may reflect the favorable price trend of high-ESG stocks during our sample period, which could create speculative opportunities that these funds exploit.

Mention-only funds behave similarly to opportunity funds: their portfolio ESG ratings are slightly higher than those of non-ESG funds, but of comparable magnitude. Despite the lack of clarity in their ESG disclosures, this pattern suggests that the objective of mention only funds with regard to ESG information is more likely pecuniary than non-pecuniary.

4.2 Decomposing ESG Ratings: Externality vs. Materiality

The analysis above relies on overall ESG ratings, which aggregate diverse underlying dimensions into a single score. Some of these dimensions capture the externalities that a firm imposes on the world through its actions—for example, its carbon emissions or labor practices—while others measure the material ESG risks and opportunities that could affect the firm's financial performance—for example, its exposure to regulatory risk or its ability to attract talent. Since our classification is designed to capture fund intentions, a higher fund-level ESG rating may not necessarily indicate a greater propensity toward impact investing, particularly if the portfolio tilt is driven primarily by the materiality component of ESG ratings.

To address this concern, we decompose firms' ESG ratings into two components: an *ESG-externality* component ($ESGx$), capturing metrics related to the impact that a firm has on the outside world; and an *ESG-materiality* component ($ESGm$), measuring how E, S, or G exposures might materially affect the firm's financial performance. This decomposition allows us to revisit the portfolio analysis above and examine which component drives the documented portfolio tilts.

4.2.1 Decomposition Methodology

Our decomposition proceeds in two steps. First, we collect all underlying sub-components of the ESG ratings provided by a given rating agency. Second, we classify each sub-component as belonging to either the ESG-externality group ($ESGx$) or the ESG-materiality group ($ESGm$). To perform this classification, we use ChatGPT 4.1, querying it through separate API calls for each sub-component with a prompt that describes the two categories and asks the model to assign exactly one.[12] Internet Appendix III reports the exact prompt utilized, the resulting categorization of each sub-component, and the model's explanation for each assignment.

We apply this methodology to MSCI ESG ratings, which at their most disaggregated level comprise 38 sub-components ranging from "Carbon Emissions Score" to "Opportunities in Clean

[12]The categorization is robust to using alternative ChatGPT models.

Tech Score." The distribution of these sub-components across the Environmental (E), Social (S), and Governance (G) pillars is as follows:

	E	S	G
ESG externality	6	8	1
ESG materiality	8	7	8

The E and S sub-components are fairly evenly split between externality and materiality, whereas most G sub-components are assigned to materiality. This is expected, since governance metrics such as board structure or executive compensation primarily capture factors material to financial performance rather than externalities imposed on the outside world.

Once the categorization is completed, we construct a *component rating* for each group $g \in \{ESGx, ESGm\}$ by averaging the scores of the sub-components belonging to that group:

$$s_{it}^g = \frac{1}{N_{it}^g} \sum_{u \in \mathcal{U}_{it}^g} s_{it}^g(u), \tag{7}$$

where $\mathcal{U}_{it}^g$ is the set of non-missing sub-components for firm i at time t in group g, and N_{it}^g is the number of elements in that set.[13] Similarly to the overall rating, we then cross-sectionally standardize each composite rating within month:

$$z_{it}^g \equiv \frac{s_{it}^g - E_t^{cs}[s_{it}^g]}{\sqrt{Var_t^{cs}[s_{it}^g]}}. \tag{8}$$

To understand how the two component ratings relate to the overall ESG rating, we estimate the regression

$$z_{it}^{ESG} = \beta_x \cdot z_{it}^{ESGx} + \beta_m \cdot z_{it}^{ESGm} + \varepsilon_{it}. \tag{9}$$

Table 3 reports the results under four specifications: unconstrained OLS; a constrained version imposing $\beta_x + \beta_m = 1$; a specification in which z_{it}^{ESGm} is orthogonalized with respect to z_{it}^{ESGx}; and one combining both the constraint and orthogonalization. Estimates are stable across all specifications: $\hat{\beta}_x$ ranges from 0.403 to 0.507, while $\hat{\beta}_m$ ranges from 0.518 to 0.570. The ESG-materiality component carries slightly more weight than the externality component in determining overall ESG ratings, though both are economically important. The similarity of constrained and unconstrained estimates indicates that overall ESG ratings can be well-approximated as convex

[13]For a meaningful aggregation, all sub-components have been normalized to the same $[0, 100]$ range.

combinations of the two components, and the negligible effect of orthogonalization confirms that our constructed components are largely independent.[14]

4.2.2 Portfolio Analysis by Rating Component

We now repeat the regression analysis in (6), replacing the overall ESG rating with the two component ratings z_{it}^{ESGx} and z_{it}^{ESGm}. Table 4 reports the results.

When ESG funds are considered as a single group (top panel), the average ESG-externality and ESG-materiality ratings of their value-weighted portfolios are, respectively, 8% (0.041/0.516) and 19% (0.059/0.304) higher than those of non-ESG funds, both significant at the 1% level. Patterns are similar for over-weighted portfolios. These aggregate results, however, once again conceal substantial heterogeneity across fund types.

Impact funds are the only type whose portfolios display significantly higher ESG-externality ratings than those of non-ESG funds. Their value-weighted and over-weighted externality ratings are 27% (0.140/0.516) and 31% (0.146/0.466) higher, respectively, both significant at the 1% level. They also display significantly higher ESG-materiality ratings—52% (0.157/0.304) and 56% (0.160/0.284) higher for value-weighted and over-weighted portfolios, respectively. Impact funds are thus the only category that tilts toward stocks scoring highly on both dimensions. This is consistent with their dual mandate: as funds with non-pecuniary ESG preferences, they seek firms that generate positive externalities; but as mutual funds subject to fiduciary duties, they also have pecuniary interests that lead them to favor firms with strong ESG-materiality profiles.

Exclusionary funds present an asymmetric pattern. Their ESG-externality ratings are not statistically different from those of non-ESG funds in either the value-weighted or over-weighted portfolios. In contrast, their ESG-materiality ratings are significantly higher—16% (0.048/0.304) for value-weighted and 13% (0.037/0.284) for over-weighted portfolios. This suggests that the exclusionary screens applied by these funds disproportionately remove stocks with low ESG-materiality scores—plausibly firms in industries associated with material ESG risks (e.g., fossil fuels, weapons)—while having a less pronounced effect on the externality dimension.

Opportunity funds display no significant tilt in ESG-externality ratings, but their ESG-materiality ratings are significantly higher than those of non-ESG funds—11% (0.033/0.304) for value-weighted and 13% (0.038/0.284) for over-weighted portfolios. This is a key finding: it reveals that the slight overall ESG tilt documented for opportunity funds in Table 2 is entirely driven by the materiality component. This is precisely what we would expect from funds that use ESG information to

[14] We have also estimated specification (9) within month and within Fama–French 48 industries, obtaining stable estimates of β_x and β_m across both dimensions.

improve risk-adjusted returns: they select stocks with favorable ESG-materiality profiles because these scores capture risks and opportunities relevant to financial performance, not because they have a preference for firms that generate positive externalities.

Mention-only funds display a pattern similar to opportunity funds: no significant differences in ESG-externality ratings, but higher ESG-materiality ratings in their value-weighted (11%) and over-weighted (19%) portfolios. This reinforces the interpretation that mention-only funds, like opportunity funds, seem to use ESG information for pecuniary purposes.

In sum, the decomposition reveals a sharp distinction: impact funds are the only type that tilts toward stocks with higher ESG-externality ratings. All other ESG fund types that display any tilt relative to non-ESG funds do so exclusively through the materiality component. This finding validates both our classification methodology and our decomposition of ESG ratings, as the revealed portfolio behavior aligns closely with the stated objectives of each fund type.

4.3 Fund Holdings Across ESG Ratings

The regression analysis above compares average portfolio-level ESG ratings across fund types. To complement this analysis, we now study the full distribution of portfolio holdings across ESG-rating quintiles. This allows us to examine not only whether ESG funds tilt toward high-ESG stocks on average, but also how their allocation varies across the entire spectrum of ESG ratings.

We estimate the following saturated specification:

$$w_{jqt}^{sum} = \sum_{\theta \in \Theta} \beta_\theta \, \mathbf{1}_{\{\theta_{jt}=\theta\}} + \sum_{n=1}^{5} \gamma_q \, \mathbf{1}_{\{q=n\}} + \sum_{\theta \in \Theta} \sum_{n=1}^{5} \varphi_{\omega,q} \, \mathbf{1}_{\{\theta_{jt}=\theta\}} \mathbf{1}_{\{q=n\}} + \eta' X_{jt} + \iota_{Mstar \times t} + \varepsilon_{jt}, \quad (10)$$

where w_{jqt}^{sum} denotes fund j's cumulative portfolio weight in stocks with ESG ratings in quintile q in month t; θ_{jt} and Θ are defined as before; X_{jt} includes fund-level controls; and $\iota_{M \times t}$ denotes Morningstar category $\times$ month fixed effects. The omitted categories are non-ESG funds and the middle quintile ($q = 3$), and standard errors are clustered by fund and month.

The estimated relative cumulative weights held by each ESG fund type in each quintile, relative to non-ESG funds, are therefore given by

$$\Delta \hat{w}_{\theta q}^{sum} = \begin{cases} \hat{\beta}_\theta & \text{for } q = 3, \\ \hat{\beta}_\theta + \hat{\varphi}_{\theta,q} & \text{for } q \in \{1, 2, 4, 5\}. \end{cases} \quad (11)$$

Figure 5 plots these estimates for the four ESG fund types across quintiles of overall ESG ratings (left panel), ESG-externality ratings (center panel), and ESG-materiality ratings (right panel). Table 5 reports the full set of estimates and their statistical significance.

The left panel of Figure 5 reveals different allocation patterns across fund types. Impact funds display a steep, monotonically increasing profile: they hold approximately 6.7 percentage points less portfolio weight in the lowest ESG quintile and 9.6 percentage points more in the highest quintile, relative to non-ESG funds (both significant at the 1% level). This translates into a spread of over 16 percentage points between the bottom and top quintiles—a large and economically meaningful tilt toward high-ESG stocks. Exclusionary funds display a qualitatively similar but much flatter pattern, with a spread of roughly 4 percentage points. Opportunity and mention only funds display modest tilts, with their allocations across quintiles remaining closer to those of non-ESG funds.

The center panel sharpens the distinction. Impact funds exhibit a pronounced upward-sloping pattern across quintiles of ESG-externality ratings, with a spread of nearly 14 percentage points between the bottom and top quintiles. In contrast, exclusionary, opportunity, and mention only funds display essentially flat profiles: none of these types allocates significantly more weight to stocks in the highest quintile of ESG-externality ratings than in the lowest. This confirms that only impact funds actively tilt their portfolios toward stocks that generate positive externalities.

The right panel reveals a more nuanced picture. Impact funds again display a steep upward-sloping pattern, with a spread of roughly 11 percentage points between the bottom and top quintiles. Exclusionary funds show a moderate tilt, underweighting the lowest quintile by about 2.4 percentage points (significant at the 5% level) while displaying no significant overweighting of the top quintile—consistent with a strategy that removes low-materiality stocks without actively seeking high-materiality ones.

Opportunity funds present an interesting pattern. Their allocation across quintiles of ESG-materiality ratings is essentially flat. This may appear surprising for two reasons. First, these funds explicitly claim to use ESG information for financial purposes. Second, in Table 4 we documented that opportunity funds exhibit significantly higher *average* ESG-materiality ratings than non-ESG funds. The apparent tension between a higher portfolio-level average and a flat distributional profile can be reconciled by the substantial heterogeneity within the opportunity category. A subset of opportunity funds may tilt strongly toward high-materiality stocks, pulling the average upward, while many others reference ESG considerations in their prospectuses without systematically conditioning their portfolio allocation on ESG-materiality scores. When these funds are pooled together, the concentrated tilts of the former group are diluted by the latter, producing a higher mean but no discernible pattern across quintiles.

To explore this heterogeneity, we isolate a subgroup of opportunity funds that demonstrate a revealed sensitivity to ESG-materiality information, defined as those with particularly high holdings in the top quintile of ESG-materiality ratings. This subgroup, labeled "Opp, high q5" in Figure 5, displays a dramatically different pattern: a steep monotonically increasing profile across materiality quintiles, with approximately 3.8 percentage points less weight in the bottom quintile and 8.0 percentage points more in the top quintile, both highly significant. The spread between the bottom and top quintiles (nearly 12 percentage points) is comparable in magnitude to that of impact funds across overall ESG ratings. This finding confirms that, within the broad opportunity category, there are funds that do systematically exploit ESG-materiality information in their portfolio construction.

Taken together, the distributional analysis reinforces and extends the findings from the portfolio-level regressions. These empirical patterns are strongly aligned with the economic objectives that each fund type discloses in its prospectus, providing key evidence that our classification captures meaningful differences in investment behavior.

4.4 Theoretical Framework

We conclude this section by developing a parsimonious theoretical framework that formalizes how ESG information enters the portfolio optimization problem of each fund type, and verify whether the empirical patterns documented above are consistent with the model's predictions. Given the scope of this analysis, we focus on a static partial equilibrium setting where asset prices are exogenously given.[15]

Assets. Company i is characterized by the triplet (r_i, m_i, x_i), where r_i is the equity excess return realized at the end of the period, m_i is the ESG materiality, and x_i the ESG externalities created by the company over that period.[16] Intuitively, a company can take ESG corporate actions that affect its own performance (generating ESG materiality) or society at large (generating ESG externalities). For example, investing in energy-efficient operations may reduce costs and improve margins (materiality), while reducing carbon emissions benefits the environment (externality) regardless of whether it affects the firm's bottom line.

We assume (r_i, m_i, x_i) are jointly normal with

$$r_i \sim \mathcal{N}(\bar{r}, \sigma_r^2), \qquad m_i \sim \mathcal{N}(0, \sigma_m^2), \qquad x_i \sim \mathcal{N}(0, \sigma_x^2), \tag{12}$$

[15] For studies that focus on the asset pricing implications of ESG investing, see Pástor, Stambaugh, and Taylor (2021), Pedersen, Fitzgibbons, and Pomorski (2021), and Avramov, Lioui, Liu, and Tarelli (2024).

[16] We assume that m_i and x_i are already expressed in a metric that makes them comparable to r_i.

where the unconditional moments are constant and average ESG materiality and externality are normalized to zero. Importantly, while ESG materiality affects financial performance positively, ESG externalities can have either sign. We capture this distinguishing feature by assuming

$$\mathbb{C}\text{ov}(r_i, m_i) = \sigma_r \sigma_m \alpha_i \qquad \text{and} \qquad \mathbb{C}\text{ov}(r_i, x_i) = \sigma_r \sigma_x \beta_i, \tag{13}$$

where $\alpha_i \in [0, 1)$ is the average correlation between company i's excess return and its ESG materiality, and $\beta_i \in (-1, 1)$ the average correlation between the return and ESG externalities. The sign restriction on α_i reflects that material ESG factors have, by definition, a positive relationship with financial performance. For instance, a company that proactively reduces its reliance on fossil fuels lowers its exposure to carbon pricing regulation and stranded asset risk, thereby protecting its long-term financial performance. In contrast, β_i is unrestricted: a company that pays significantly above-market wages to improve worker welfare may see its returns improve if higher compensation attracts talent and boosts productivity, or decline if the additional labor costs are not offset by productivity gains. Both ESG materiality and ESG externalities are orthogonal across companies.

Ratings. A rating agency collects information about companies' ESG materiality and ESG externalities in the form of noisy signals about m_i and x_i,

$$s_{mi} = m_i + \sigma_{\epsilon m} \cdot \epsilon_{mi}, \qquad s_{xi} = x_i + \sigma_{\epsilon x} \cdot \epsilon_{xi}, \tag{14}$$

and aggregates them into an overall ESG rating

$$s_i = \omega_{mi} \cdot s_{mi} + \omega_{xi} \cdot s_{xi}, \tag{15}$$

where $(\omega_{mi}, \omega_{xi})$ capture the relative importance of ESG materiality and ESG externality in the rating. Noise components are orthogonal across assets and signals. Throughout our analysis, we assume that ESG funds observe the individual signals s_{mi} and s_{xi}, rather than only the aggregate rating s_i. This assumption is natural given that all sub-components of ESG ratings are easily available, allowing funds to construct their own ESG-materiality and ESG-externality measures—precisely as we did in our empirical decomposition in Section 4.2.

Funds. In line with our classification in Section 2, we distinguish among three ESG fund types: *exclusionary* (E), *impact* (I), and *opportunity* (O).Each fund type $\theta \in \{E, I, O\}$ starts with capital $W_{\theta 0}$, allocates ϕ_θ to risky assets, and ends the period with $W_{\theta 1} = W_{\theta 0}(1 + r_f) + \phi_\theta' r$. All fund types maximize the expectation of a CARA objective $f(W_{\theta 1}) = -\exp\{-\gamma W_{\theta 1}\}$, conditional on their information set, yielding a mean-variance allocation problem.[17]

- *Non-ESG* funds (N) serve as a benchmark. Their objective is only driven by financial considerations, ignoring ESG information entirely. Their allocation problem is

$$\phi_N^* = \arg\max_{\phi_N} \; \mathbb{E}[\phi_N' r] - \frac{\gamma}{2}\mathrm{Var}[\phi_N' r]. \tag{16}$$

- *Exclusionary* funds maximize the same financial objective as non-ESG funds, but over a restricted investment universe. Without taking a stand on which exclusionary fund excludes which asset, but capturing the idea that low-ESG assets are more likely to be excluded, we model the average asset allocation of exclusionary funds as a scaled-down version of that of non-ESG funds, where the scaling positively depends on the assets' ESG ratings:

$$\phi_{Ei}^* = \phi_{Ni}^* \cdot \Lambda(s_i), \tag{17}$$

  where $\Lambda(s_i) = (1+e^{-s_i})^{-1}$ is a standard logistic function, satisfying $\Lambda(s_i)' > 0$, $\lim_{s_i\to\infty}\Lambda(s_i) = 1$, and $\lim_{s_i\to-\infty}\Lambda(s_i) = 0$. The scaling can be interpreted as the likelihood that an asset is not excluded, or the fraction of exclusionary funds that do not exclude that asset. The higher the ESG rating of an asset, the lower the fraction of funds that exclude it, and consequently the higher the scaling.

- *Impact* funds incorporate both financial and non-pecuniary ESG considerations, conditioning on the ESG signals (s_m, s_x). We model the non-pecuniary component as driven by the ESG externalities a company generates, making the impact funds' objective $f(\varphi W_{1I} + (1-\varphi)\phi_I' x)$, where φ governs the relative importance of financial versus non-pecuniary considerations. Their allocation problem is

$$\begin{aligned}\phi_I^* = \arg\max_{\phi_I} \; & \varphi \left(\mathbb{E}[\phi_I' r | s_m, s_x] - \frac{\gamma}{2}\varphi\,\mathrm{Var}[\phi_I' r | s_m, s_x]\right) \\ & + (1-\varphi)\left(\mathbb{E}[\phi_I' x | s_m, s_x] - \frac{\gamma}{2}(1-\varphi)\,\mathrm{Var}[\phi_I' x | s_m, s_x] - \gamma\varphi\,\mathrm{Cov}[\phi_I' r, \phi_I' x | s_m, s_x]\right). \end{aligned} \tag{18}$$

- *Opportunity* funds use ESG information purely for financial gain, conditioning on (s_m, s_x):

$$\phi_O^* = \arg\max_{\phi_O} \; \mathbb{E}[\phi_O' r | s_m, s_x] - \frac{\gamma}{2}\mathrm{Var}[\phi_O' r | s_m, s_x]. \tag{19}$$

[17] We use the same objective function to emphasize that differences across fund types stem from how they incorporate ESG information, not from different risk attitudes.

Opportunity funds care about companies' ESG-materiality and ESG-externality ratings only to the extent that they can help improve their financial performance. Therefore, their allocation may include investments in both high- and low-ESG companies, depending on how ESG ratings correlate with returns.

All three ESG fund types use ESG information, but in economically distinct ways. Exclusionary funds restrict their investible set without altering the allocation within it. Impact and opportunity funds both condition on ESG signals, but only impact funds have a non-pecuniary motive that directly rewards ESG externalities. In practice, these types need not be mutually exclusive, but we maintain a stark distinction to isolate each type's defining feature.

Optimal allocations. Assuming returns are independent across assets—it provides a more transparent analysis—the next proposition characterizes the optimal allocation of each fund type.

Proposition 1. *The optimal capital allocation to asset i by non-ESG funds is $\phi^*_{Ni} = \bar{r}/(\gamma\sigma_r^2)$, whereas the corresponding amounts invested by exclusionary, impact, and opportunity funds are*

$$\phi^*_{Ei} = \frac{\bar{r}}{\gamma\sigma_r^2} \cdot \Lambda(s_i), \qquad \phi^*_{Ii} = \frac{\varphi\hat{r}_i + (1-\varphi)\hat{x}_i}{\gamma(\varphi^2\hat{\sigma}^2_{ri} + (1-\varphi)^2\hat{\sigma}^2_{xi} + 2\varphi(1-\varphi)\hat{\sigma}_{rxi})}, \qquad \phi^*_{Oi} = \frac{\hat{r}_i}{\gamma\hat{\sigma}^2_{ri}}, \tag{20}$$

respectively, where the conditional moments $\hat{r}_i$, $\hat{x}_i$, $\hat{\sigma}^2_{ri}$, $\hat{\sigma}^2_{xi}$, and $\hat{\sigma}_{rxi}$ are given by

$$\hat{r}_i = \mathbb{E}[r_i|s_m, s_x] = \bar{r} + \alpha_i\frac{\sigma_r\sigma_m}{\sigma_m^2 + \sigma_{\epsilon m}^2}s_{mi} + \beta_i\frac{\sigma_r\sigma_x}{\sigma_x^2 + \sigma_{\epsilon x}^2}s_{xi}, \tag{21}$$

$$\hat{x}_i = \mathbb{E}[x_i|s_m, s_x] = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{\epsilon x}^2}s_{xi}, \tag{22}$$

$$\hat{\sigma}^2_{ri} = \mathbb{V}\text{ar}[r_i|s_m, s_x] = \sigma_r^2\left(1 - \alpha_i^2\frac{\sigma_m^2}{\sigma_m^2 + \sigma_{\epsilon m}^2} - \beta_i^2\frac{\sigma_x^2}{\sigma_x^2 + \sigma_{\epsilon x}^2}\right), \tag{23}$$

$$\hat{\sigma}^2_{xi} = \mathbb{V}\text{ar}[x_i|s_m, s_x] = \sigma_x^2\left(1 - \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{\epsilon x}^2}\right), \tag{24}$$

$$\hat{\sigma}_{rxi} = \mathbb{C}\text{ov}[r_i, x_i|s_m, s_x] = \sigma_r\sigma_x\beta_i\left(1 - \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{\epsilon x}^2}\right). \tag{25}$$

Since all assets share the same unconditional return distribution, the non-ESG fund allocates equally across all assets. Differences in the optimal allocations of ESG funds are therefore driven entirely by heterogeneity in ESG ratings (s_{mi}, s_{xi}) and in the asset-specific correlations (α_i, β_i).

The allocation of opportunity funds depends on ESG ratings only through the conditional expected return $\hat{r}_i$. The materiality signal s_{mi} enters with a coefficient proportional to $\alpha_i > 0$, generating an unambiguous tilt toward high-materiality stocks. The externality signal s_{xi} enters with a coefficient proportional to β_i, which can be positive or negative, so the net effect on portfolio allocation is ambiguous. In contrast, the allocation of impact funds is shaped by two forces: $\hat{r}_i$ (weighted by φ) and the conditional expected externality $\hat{x}_i$ (weighted by $1-\varphi$). The externality signal enters directly through $\hat{x}_i$, creating a tilt toward high-externality stocks that is entirely absent from the opportunity funds' problem. This is the key theoretical distinction: impact funds are the only type for which ESG-externality information has a direct, unambiguous effect on portfolio allocation, over and above any indirect effect through expected returns. Finally, the exclusionary fund's allocation depends on the overall rating s_i through $\Lambda(\cdot)$, which mechanically increases (decreases) the weight of assets with high (low) ESG scores. It is worth pointing out that the larger weights of high-ESG assets are not driven by an explicit intention to increase the exposure to those assets (*positive screening*), but rather by the constraint to remove the exposure to low-ESG assets (*negative screening*).

Numerical predictions. We now verify that these allocations reproduce the empirical patterns documented in Tables 2–4 and Figure 5. We generate asset heterogeneity by varying α_i, β_i, s_{mi}, and s_{xi} independently: 5 equally spaced values of $\alpha_i \in [0, 0.6]$ and $\beta_i \in [-0.5, 0.5]$, and 15 representative realizations for each signal. Other parameter values are: $\bar{r} = 0.1$, $\sigma_r = \sigma_x = \sigma_m = 0.15$, $\sigma_{\epsilon x} = \sigma_{\epsilon m} = 0.4$, $\omega_{mi} = \omega_{xi} = 0.5$, $\gamma = 10$, $\varphi = 0.5$.

Using the optimal allocations from Proposition 1, we first compute value-weighted portfolio ratings $S_\theta^{ESG} = \sum_i w_{\theta i} s_i$, $S_\theta^{ESGx} = \sum_i w_{\theta i} s_{xi}$, and $S_\theta^{ESGm} = \sum_i w_{\theta i} s_{mi}$, where $w_{\theta i} = \phi^*_{\theta i} / \sum_\ell \phi^*_{\theta \ell}$. The table below reports these portfolio ratings relative to non-ESG funds (which are zero by construction):

	ESG	*ESGx*	*ESGm*
Exc	0.03	0.03	0.03
Imp	0.08	0.12	0.04
Opp	0.02	0.00	0.04

Despite the parsimony of the setting, the model reproduces the main empirical patterns. Impact funds have the highest portfolio ratings across all three measures, with the largest difference on the externality dimension—consistent with their non-pecuniary preference for firms that generate positive externalities. Opportunity funds show no difference in their portfolio ESG-externality rating relative to non-ESG funds, but do tilt their holdings toward stocks with high ESG-materiality ratings—consistent with their purely pecuniary use of ESG information. Finally, the negative screening of exclusionary funds leads to a modest and symmetric increase in portfolio ESG ratings across all three measures, reflecting the mechanical effect of removing low-rated stocks from the investible set.

We then construct quintiles of each rating measure and compute cumulative portfolio weights within each quintile, $w^{sum}_{\theta q} = \sum_{i \in Q_q(\zeta)} w_{\theta i}$, for $\theta \in \{E, I, O\}$, $q \in \{1, \ldots, 5\}$, and $\zeta \in \{s_i, s_{mi}, s_{xi}\}$. Figure 6 reports the results as deviations from non-ESG allocations, and the patterns closely mirror the empirical evidence in Figure 5. Across quintiles of overall ESG ratings (left panel), impact funds display a steep monotonically increasing profile, exclusionary funds a moderate upward slope, and opportunity funds a flatter pattern. Across externality quintiles (center panel), the distinction sharpens: impact funds exhibit a pronounced tilt while exclusionary and opportunity funds are essentially flat—opportunity funds because the ambiguous sign of β_i washes out any systematic pattern, and exclusionary funds because their screens depend on the overall ESG rating. Across materiality quintiles (right panel), both impact and opportunity funds display upward-sloping profiles, since both benefit from the positive correlation between materiality and returns ($\alpha_i > 0$), while exclusionary funds show a moderate mechanical tilt.

Overall, the theoretical predictions align closely with the empirical evidence across all three dimensions. This confirms that the portfolio allocation patterns documented in this section are not merely statistical artifacts but rather the natural outcome of optimizing portfolios under the different objectives that our classification captures.

5 Novel Insights from Granular ESG Classification

5.1 Greenwashing

A central concern in the ESG investment debate is whether funds that market themselves as sustainable actually deliver on their promises. Funds that claim to pursue ESG objectives but hold portfolios indistinguishable from, or even worse than, conventional funds are commonly labeled "greenwashers." Our classification allows us to bring greater precision to this question, both in defining which funds should be evaluated for greenwashing and in measuring the extent to which it occurs.

An important conceptual point is that greenwashing, properly defined, requires a gap between what a fund promises and what it delivers. In our framework, this concern applies most directly to *impact* funds, since they are the only type that explicitly promises to invest on the basis of non-pecuniary ESG preferences—to channel capital toward firms that generate positive externalities. If an impact fund's portfolio exhibits ESG ratings no better than those of a typical non-ESG fund, there is a clear disconnect between its stated objective and its revealed behavior.

For the other ESG fund types, the greenwashing label is less appropriate. Opportunity funds explicitly disclose that they use ESG information for financial purposes, not to generate positive externalities. A low ESG rating in their portfolio is not a broken promise—it is consistent with their stated mandate. Similarly, exclusionary funds promise only to avoid certain categories of stocks, not to tilt toward high-ESG ones, and mention-only funds make no specific ESG commitments at all. Nonetheless, to provide a complete picture, we report greenwashing statistics for all ESG fund types, while emphasizing that the economic interpretation differs across categories.

Identifying greenwashers. We flag a fund as a potential greenwasher if its portfolio-level ESG rating falls in the *bottom* of the distribution of non-ESG fund ratings for a *sufficiently large fraction* of its lifetime. Formally, a fund receives the greenwasher label if its value-weighted portfolio ESG rating (or one of the two component ratings) falls below a given percentile of the cross-sectional distribution of non-ESG fund ratings for more than a specified share of the months in which the fund appears in our sample.

We vary both thresholds to assess robustness. For the definition of "bottom," we consider two cutoffs: below the 10th percentile and below the 25th percentile of the non-ESG distribution. For the definition of "sufficiently large fraction," we consider two persistence thresholds: at least 50% and at least 75% of the fund's lifetime. The combination of these two dimensions produces four increasingly stringent criteria, ranging from a lenient standard (below the 25th percentile for at least 50% of the fund's lifetime) to a strict one (below the 10th percentile for at least 75%). We apply this procedure separately to three rating measures: the overall ESG rating, the ESG-externality component ($ESGx$), and the ESG-materiality component ($ESGm$). For impact funds, the externality component is the most economically relevant metric, since their defining feature is a non-pecuniary preference for firms that generate positive externalities.

Findings. Figure 7 reports the percentage of funds flagged as potential greenwashers in each ESG category, for each combination of rating measure, percentile threshold, and persistence threshold.

The most important finding is that *impact funds have the lowest greenwashing rates across nearly all criteria and rating measures*. Under the strictest definition (below the 10th percentile for at least 75% of lifetime), no impact fund is flagged on the overall ESG rating or the materiality component, and only 2% are flagged on the externality component. Even under the most lenient definition (below the 25th percentile for at least 50% of lifetime), only 5% of impact funds are flagged on the overall ESG rating, 11% on the externality component, and 5% on the materiality component. These rates are the lowest or among the lowest across all ESG fund types. The low incidence of greenwashing among impact funds is reassuring: it indicates that the vast majority of funds claiming non-pecuniary ESG objectives do in fact hold portfolios with ESG ratings well above the bottom of the non-ESG distribution. Combined with the evidence from Sections 4.1–4.3

that impact funds display the largest portfolio tilts toward high-ESG stocks, this suggests that concerns about widespread greenwashing among ESG funds are overstated.

Turning to the other fund types, mention-only and opportunity funds display higher flagging rates than impact funds across all criteria and rating measures. Under the lenient definition, 17% of mention only funds and 24% of opportunity funds are flagged on the overall ESG rating; on the externality component, the figures are 27% and 23%, respectively. These rates are comparable to, and in some cases exceed, those of non-ESG funds, consistent with the finding in Section 4.1 that these fund types hold portfolios whose ESG profiles are only marginally different from non-ESG funds. However, as noted above, interpreting these figures as greenwashing requires caution: neither type promises to hold high-ESG portfolios, so a low ESG rating does not constitute a broken promise in the same way it would for an impact fund.

Exclusionary funds occupy an intermediate position. Their flagging rates are generally lower than those of opportunity and mention only funds, but higher than those of impact funds. This is consistent with the mechanics of an exclusionary strategy: removing certain low-ESG stocks from the investible set raises the average ESG rating of the portfolio, but the absence of positive screening means that some exclusionary funds may still hold portfolios with overall ESG ratings in the lower tail.

In sum, our analysis reveals that greenwashing is rare among the funds for which it matters most: impact funds display the lowest greenwashing rates across nearly all criteria and rating measures, indicating that the vast majority of funds claiming non-pecuniary ESG objectives hold portfolios consistent with those claims. An important caveat to both this analysis and the portfolio-level results in Section 4, however, is that they rely on portfolio holdings disclosed at the end of each quarter. If funds strategically adjust their portfolios around disclosure dates, the holdings observed by investors and researchers may not be representative of the positions held during the remainder of the quarter. This concern is particularly relevant for ESG funds: a fund could maintain a low-ESG portfolio for most of the quarter, tilt toward high-ESG stocks shortly before the disclosure date, and revert afterward—appearing to deliver on its ESG mandate while in practice doing so only intermittently. In the next subsection, we directly address this concern by studying the dynamics of ESG factor exposure around quarterly disclosure events, which allows us to assess whether the portfolio patterns documented so far reflect persistent investment behavior or strategic pre-disclosure repositioning.

5.2 Green Window Dressing

A fund engaging in window dressing temporarily inflates a portfolio characteristic before a mandatory disclosure date, only to reverse the position once the snapshot has been taken. Parise and Rubin (2025) provide evidence of such strategic behavior among ESG funds: they document that the average ESG factor loading of these funds rises significantly in the days before quarterly disclosures and reverts immediately afterward—evidence of "green" window dressing. Our granular classification of ESG funds allows us to ask a sharper question: which fund types drive this aggregate finding?

Hypotheses. Two opposing views generate distinct predictions. Under the first view, funds adhere to their mandates. Impact and exclusionary funds hold high-ESG assets as a persistent feature of their strategy—baked into the objective function for impact funds, and a structural consequence of the constrained investible universe for exclusionary ones. Neither type would need to temporarily inflate an exposure they already maintain, and transaction costs would deter the round-trip repositioning that window dressing requires. By contrast, opportunity and mention-only funds operate without binding ESG constraints yet are often classified as sustainable by third-party rating providers (Figure 1), creating an incentive to inflate ESG exposure at the disclosure snapshot to secure a ratings boost and attract ESG-motivated flows. This view predicts window dressing concentrated in opportunity and mention only funds, and absent among impact and exclusionary funds.

Under the second view, funds deviate from their mandates. Persistent adherence to a strict ESG mandate involves genuine return trade-offs, so the funds under the most demanding mandates—impact and exclusionary—face the greatest temptation to relax ESG exposure mid-quarter and restore it before the snapshot. Under this view, window dressing should be most pronounced precisely among impact and exclusionary funds.

Empirical design. We follow the event-study design of Parise and Rubin (2025). For each quarterly portfolio disclosure event d (the last trading day of the quarter), we estimate the ESG factor loading $\hat{\beta}^{\mathrm{ESG}}_{j,d,\mathrm{w}}$ of fund j by OLS within three non-overlapping windows w: a *control* (ctrl) window spanning the second month of the quarter (excluding the first and last trading days), a *pre-disclosure* (pre) window covering the 10 trading days immediately before the disclosure date, and a *post-disclosure* (post) window covering the 10 trading days beginning two days after. Our baseline specification augments the Fama-French-Carhart six factors with the returns on the Morningstar ESG index (FFC6+ESG); robustness results using alternative factor models are reported in Internet Appendix V.

For each fund type $\theta \in \{E, I, O, ESG, N\}$, we compute three statistics:

$$\Delta\hat{\beta}^{\text{ESG}}_{\theta,\text{pre}} = \frac{1}{N_\theta} \sum_{(j,d):\, j\in\theta_d} \hat{\beta}^{\text{ESG}}_{j,d,\text{pre}} - \hat{\beta}^{\text{ESG}}_{j,d,\text{ctrl}}, \tag{26}$$

$$\Delta\hat{\beta}^{\text{ESG}}_{\theta,\text{rev}} = \frac{1}{N_\theta} \sum_{(j,d):\, j\in\theta_d} \hat{\beta}^{\text{ESG}}_{j,d,\text{post}} - \hat{\beta}^{\text{ESG}}_{j,d,\text{pre}}, \tag{27}$$

$$\Delta\hat{\beta}^{\text{ESG}}_{\theta,\text{net}} = \frac{1}{N_\theta} \sum_{(j,d):\, j\in\theta_d} \hat{\beta}^{\text{ESG}}_{j,d,\text{post}} - \hat{\beta}^{\text{ESG}}_{j,d,\text{ctrl}}, \tag{28}$$

where $N_\theta = \sum_d |\theta_d|$ is the total number of fund-event pairs in category θ. $\Delta\hat{\beta}^{\text{ESG}}_{\theta,\text{pre}}$ measures the pre-disclosure increase in ESG loading relative to the within-quarter baseline, with a positive value being necessary condition for window dressing. $\Delta\hat{\beta}^{\text{ESG}}_{\theta,\text{rev}}$ captures the post-disclosure reversal, with a negative value confirming that the increase is unwound. $\Delta\hat{\beta}^{\text{ESG}}_{\theta,\text{net}}$ measures the net change from baseline to post-disclosure, with a value indistinguishable from zero indicating a fully temporary repositioning. Inference relies on a wild bootstrap procedure that accounts for cross-sectional correlation of fund residuals across calendar dates, following Parise and Rubin (2025).[18]

Findings. Table 6 reports level estimates and changes in ESG exposure by fund type (Panel A), as well as cross-type differences relative to non-ESG funds (Panel B). Figure 8 presents these results graphically.

We begin by reproducing the main finding of Parise and Rubin (2025) within our sample.[19] Pooling all ESG fund types, we find a significant pre-disclosure increase in ESG exposure ($\Delta\hat{\beta}^{\text{ESG}}_{\text{ESG,pre}} = 0.221$, 1% significance) accompanied by a full post-disclosure reversal ($\Delta\hat{\beta}^{\text{ESG}}_{\text{ESG,rev}} = -0.185$, 1% significance), with the net post-control change statistically indistinguishable from zero. Non-ESG funds also display a positive pre-control delta (0.151, 1% significance), but critically their reversal is insignificant and the net change remains positive (0.104, 5% significance), indicating a persistent, rather than strategic, increase in ESG loading. Notably, the cross-type difference between the aggregate ESG and non-ESG pre-control deltas is itself not statistically significant. As the decomposition by fund type will make clear, this is because the aggregate ESG effect averages over types with sharply divergent behaviors. Indeed, decomposing the ESG universe by fund type reveals a striking pattern of heterogeneity, strongly consistent with the first hypothesis above.

[18] We use $B = 9{,}999$ bootstrap iterations. Cross-type tests use the same global draw to form the bootstrap distribution of the difference $\Delta\hat{\beta}^{\text{ESG}}_A - \Delta\hat{\beta}^{\text{ESG}}_B$; see their Internet Appendix for a complete description.

[19] Our sample of ESG funds is larger than that used by Parise and Rubin (2025), who include only the subset classified as sustainable by Morningstar or containing ESG keywords in the fund name (see Figure 1 for the overlap).

Impact funds exhibit no window dressing. In the control window, they display the highest ESG loading across all fund types (0.329, 1% significance), exceeding non-ESG funds by 0.163 (1% significance). This elevated baseline, also shown in the left panel of Figure 8, is consistent with their non-pecuniary ESG objective generating a structurally high loading even in the period furthest from any disclosure incentive. Their ESG exposure is essentially flat across all three windows, as revealed by the green line in the center panel of Figure 8. The pre-control delta (0.018), the reversal (0.010), and the net change (0.048) are all statistically insignificant. Notably, impact funds are the only type whose control-period ESG loading is significantly higher than that of non-ESG funds, yet they show no pre-disclosure repositioning whatsoever. Exclusionary funds display a nearly identical profile: the pre-control delta (0.001), the reversal (0.076), and the net change (0.099) are all insignificant. This is consistent with ESG exposure being determined by structural investment restrictions rather than by active repositioning within the investible universe.

Opportunity funds exhibit the classic window dressing pattern. In the control window, their ESG loading (0.217) is statistically indistinguishable from that of non-ESG funds, consistent with an investment approach that does not require persistent ESG tilts. In the pre-disclosure window, however, their loading nearly triples to 0.553—a pre-control delta of 0.336 (1% significance)—followed by an almost exact reversal of −0.338 (1% significance). The post-disclosure loading returns to 0.217, numerically identical to the control level, with a net change of precisely −0.000. The center panel of Figure 8 illustrates this pattern vividly: the ESG beta trajectory of opportunity funds traces an inverted V, spiking at the pre-disclosure window and reverting to baseline immediately afterward.

Mention-only funds display the most extreme pre-disclosure increase. Their ESG loading rises from 0.276 in the control window to 0.737 pre-disclosure, yielding the largest pre-control delta among all fund types (0.462, 1% significance). Unlike opportunity funds, however, the reversal is only partial: −0.245 (1% significance), offsetting roughly half of the pre-disclosure increase. The post-disclosure loading settles at 0.493, significantly above the control level, with a net change of 0.220 (5% significance)—the only category for which this measure is reliably different from zero. This partial reversal suggests that mention only funds engage in strategic pre-disclosure repositioning, but that some of their pre-disclosure increase may also reflect a genuine (if transient) shift in portfolio composition that is not fully unwound.

These findings demonstrate that green window dressing is not a general property of ESG funds. It is concentrated entirely in the two types with the weakest ESG mandates (opportunity and mention-only funds) and absent among those with more binding commitments (impact and exclusionary funds). The aggregate ESG window dressing documented by Parise and Rubin (2025) is, in our sample, an average of a null effect for roughly one quarter of ESG fund observations and a pronounced effect for the remainder, which accounts for the vast majority of ESG AUM, as documented in Section 3.

This finding has direct implications for investors and regulators. An investor relying on a fund's pre-disclosure ESG factor loading to assess its ESG orientation would be systematically misled: the funds displaying the highest ESG loading at the disclosure snapshot are not the impact funds with genuine non-pecuniary preferences, but the opportunity and mention only-funds with the weakest ESG mandates. The persistent ESG commitment of impact funds is most visible precisely when no disclosure incentive is present—in the mid-quarter control window—not at the moment investors observe the portfolio. More broadly, these results reinforce the paper's central message: aggregate ESG indicators obscure economically important heterogeneity, and a classification based on investment objectives is necessary to accurately assess which funds deliver on their ESG promises.

5.3 ESG Violations

The analyses in Section 4 established that different ESG fund types hold portfolios with systematically different ESG *ratings*. Ratings are assessments produced by third-party agencies based on firms' policies, disclosures, and management practices. They primarily measure how firms manage ESG issues (i.e., the systems and commitments they have in place), rather than providing a comprehensive measure of realized ESG outcomes. As a result, a firm can score well by adopting extensive environmental policies and producing detailed sustainability disclosures, while still generating negative externalities or committing ESG violations. The well-documented disagreement across ESG rating providers further suggests that these ratings reflect subjective assessments of ESG management and disclosure quality, rather than objective measures of firms' ESG outcomes.

This raises a pointed question about the portfolio tilts documented in Section 4: are impact funds genuinely avoiding firms that harm the world, or are they simply selecting firms that *look good* on ESG? In other words, are they simply chasing high ratings? To address this, we construct a novel dataset of firm-level ESG violations—legally adjudicated instances of environmental contamination, labor abuses, consumer protection failures, fraud, and other regulatory breaches—and study how fund holdings covary with realized ESG misconduct. Unlike ratings, violations cannot be managed through better disclosure or policy adoption. They represent hard evidence of actual harm or regulatory failure, providing a fundamentally different lens through which to evaluate whether our fund classification captures meaningful differences in real-world ESG outcomes.

Data on ESG violations. We collect firm-level data on corporate misconduct from Violation Tracker, covering the period 2015–2023. The dataset records the firm involved, the date, the nature of the offense, and the associated penalty. ESG violations span a wide range of offenses: environmental violations include breaches of the Clean Air Act and Clean Water Act, improper disposal of hazardous waste, oil spills, and illegal release of toxic substances such as PFAS or PCBs; social violations include employment discrimination, workplace safety failures, violations of the Americans with Disabilities Act, off-label promotion of medical products, and consumer protection abuses; governance violations include accounting fraud, securities violations, anti-competitive practices, and breaches of the False Claims Act.

We classify each violation along two dimensions. First, following the standard ESG taxonomy, we assign each violation to the Environmental (E), Social (S), or Governance (G) pillar. Second, mirroring our decomposition of ESG ratings in Section 4.2, we classify each violation as either ESG-externality—capturing offenses that directly harm society or the environment—or ESG-materiality—capturing offenses that primarily increase a firm's financial and operational risks. We perform this classification using ChatGPT 4.1, querying it with a description of each violation type.The following table reports the resulting cross-tabulation:

	E	S	G
ESG externality	26	40	6
ESG materiality	0	2	26

The structure is striking. All 26 Environmental violations are classified as externality-related, reflecting their direct consequences for ecosystems and public health (e.g., water pollution, toxic gas leaks, hazardous waste mismanagement). The vast majority of Social violations (40 out of 42) are also classified as externality-related, capturing harms to individuals and communities such as employment discrimination, unsafe workplaces, and violations of consumer rights. Governance violations, by contrast, are overwhelmingly materiality-related (26 out of 32), reflecting failures in corporate accountability—fraud, false claims, anti-competitive practices—that primarily expose firms to financial and legal risk.

Figure 9 provides additional context on the violation data. The top panels display the average number of violations per firm-month, broken down by ESG pillar (top-left) and by our externality-materiality classification (top-right), over the sample period. Social violations are the most frequent, followed by Environmental and Governance violations; correspondingly, externality-related violations outnumber materiality-related ones. The bottom panels display the average penalty amount as a percentage of market capitalization. While Environmental and Social penalties are typically small relative to firm size, Governance penalties—which include fraud settlements and false claims penalties—are substantially larger, consistent with the view that materiality-related violations carry greater direct financial consequences for firms.

Measuring fund exposure to violations. For each fund-month, we compute the portfolio-weighted exposure to violator stocks:

$$\text{ViolatorShare}_{jt} = \sum_i w_{jit} \cdot V_{it}, \tag{29}$$

where w_{jit} is the weight of stock i in fund j's portfolio in month t and $V_{it} \in \{0, 1\}$ equals one if stock i has at least one ESG violation in month t. We then estimate:

$$\text{ViolatorShare}_{jt} = \sum_{\theta \in \Theta} \beta_\theta \, \mathbf{1}_{\{\theta_{jt}=\theta\}} + \eta' X_{jt} + \iota_{MStar \times t} + \varepsilon_{jt}, \tag{30}$$

where θ_{jt} denotes fund j's ESG type in month t and the set $\Theta = \{Exc, Imp, Opp, Men\}$ collects all ESG fund types; X_{jt} is a vector of fund-level controls (log age, log assets, expenses, turnover, fund flows, fund flow volatility, and fund betas with respect to Fama–French–Carhart six factors); and $\iota_{MStar \times t}$ represents Morningstar category $\times$ month fixed effects. Non-ESG funds serve as the excluded category, and standard errors are clustered by fund and month. Each β_θ measures the difference in violation exposure between fund type θ and non-ESG funds within the same investment category and month. We repeat this specification separately for externality-related violations (V^x_{it}) and materiality-related violations (V^m_{it}).

Findings. Column (1) of Table 7 reports the results for overall ESG violations. The average violation exposure among non-ESG funds is 13.8% of portfolio weight, providing a natural benchmark. Impact funds display the largest reduction in violation exposure: 2.21 percentage points less than non-ESG funds (1% significance), corresponding to a 16.0% lower exposure relative to the non-ESG mean. Opportunity funds also hold significantly less in violator stocks: 1.59 percentage points less (1% significance), an 11.5% reduction. Exclusionary funds show a smaller but significant reduction of 1.09 percentage points (5% significance), while mention-only funds display a statistically insignificant difference. These results establish that the portfolio differences across fund types are not limited to ESG ratings: they extend to measurable differences in exposure to firms that have committed actual ESG violations.

Columns (2) and (3) of Table 7 decompose these results by violation type, and here a revealing asymmetry emerges. For externality-related violations—environmental contamination, labor abuses, consumer harm—impact funds display the largest and most significant reduction: 1.94

percentage points less than non-ESG funds (1% significance), a 16.2% decrease. Opportunity funds follow with a 1.36 percentage point reduction (1% significance). Exclusionary and mention only funds show smaller, statistically insignificant reductions on this dimension. For materiality-related violations—fraud, false claims, governance failures—all ESG fund types except mention-only display significant reductions relative to non-ESG funds. Crucially, on this dimension the gap between impact and opportunity funds narrows considerably: impact funds reduce exposure by 20.7% relative to the non-ESG mean, while opportunity funds reduce it by 14.2%.

The pattern that emerges from the violation-type decomposition is economically informative and distinct from what ESG ratings alone can reveal. Impact funds' avoidance of violator stocks is disproportionately concentrated in externality-related offenses, i.e., the violations that directly harm people and the environment, and that align most closely with their stated non-pecuniary preferences. This indicates that impact funds are not simply selecting firms with polished ESG profiles; they are systematically less exposed to firms that cause verifiable real-world harm.

Opportunity funds also avoid violator stocks, but the composition of their avoidance is different. Their reduction in exposure is more evenly distributed across externality-related and materiality-related violations. This pattern is consistent with a strategy that treats violation history as a signal of financial risk (i.e., firms with governance failures, fraud, or regulatory breaches face future penalties, litigation costs, and reputational damage that threaten returns), rather than as a screen motivated by concerns about social or environmental harm.

The distinction between avoiding firms that *harm the world* and avoiding firms that *face financial risk from ESG failures* echoes the externality-materiality decomposition of ESG ratings from Section 4.2, but now emerges from a completely independent data source. The convergence of evidence across ESG ratings and ESG violations reinforces the paper's central finding: our classification captures economically meaningful differences in fund objectives that manifest consistently across multiple dimensions of ESG investment behavior.

6 Conclusion

The label "ESG fund" encompasses strategies with fundamentally different objectives, yet the academic literature, industry practice, and regulatory frameworks have largely treated these strategies as interchangeable. This paper shows that they are not.

By classifying US active equity mutual funds based on the objective for which they use ESG information, we uncover a landscape that defies the prevailing narrative. The explosive growth of ESG investing has been driven not by funds with non-pecuniary preferences but by opportunity funds that integrate ESG factors into their return-maximizing mandates. Impact funds, the only type that genuinely channels capital toward firms generating positive externalities, represent a small and shrinking share of the ESG universe. Yet they are also the type that most consistently delivers on its promises: they hold the highest-ESG portfolios, display no evidence of greenwashing or window dressing, and are the least exposed to firms that commit ESG violations.

The externality-materiality decomposition introduced in this paper proves essential for understanding these differences. Impact funds are the only type that tilts toward stocks with higher ESG-externality ratings. All other fund types that tilt at all do so exclusively through the materiality channel. This pattern emerges consistently whether we examine ESG ratings, the distribution of holdings across rating quintiles, the predictions of our theoretical framework, or the entirely independent domain of corporate violations. The robustness of this finding across such diverse sources of evidence underscores that our classification captures a real and economically meaningful distinction in how funds use ESG information.

These findings carry direct implications. For investors, our results provide a clear framework for assessing whether a fund's portfolio behavior aligns with its stated objectives. For regulators, they suggest that disclosure mandates organized around the *importance* of ESG factors miss the more consequential distinction: the *objective* for which those factors are used. And for researchers, they demonstrate that pooling ESG funds into a single category can produce misleading conclusions, and that the investor types featured in theoretical models of ESG investing have distinct and identifiable empirical counterparts in the mutual fund industry. Recognizing the different shades of ESG funds is a necessary first step toward accurately evaluating the promises, the behavior, and ultimately the real-world impact of sustainable finance.

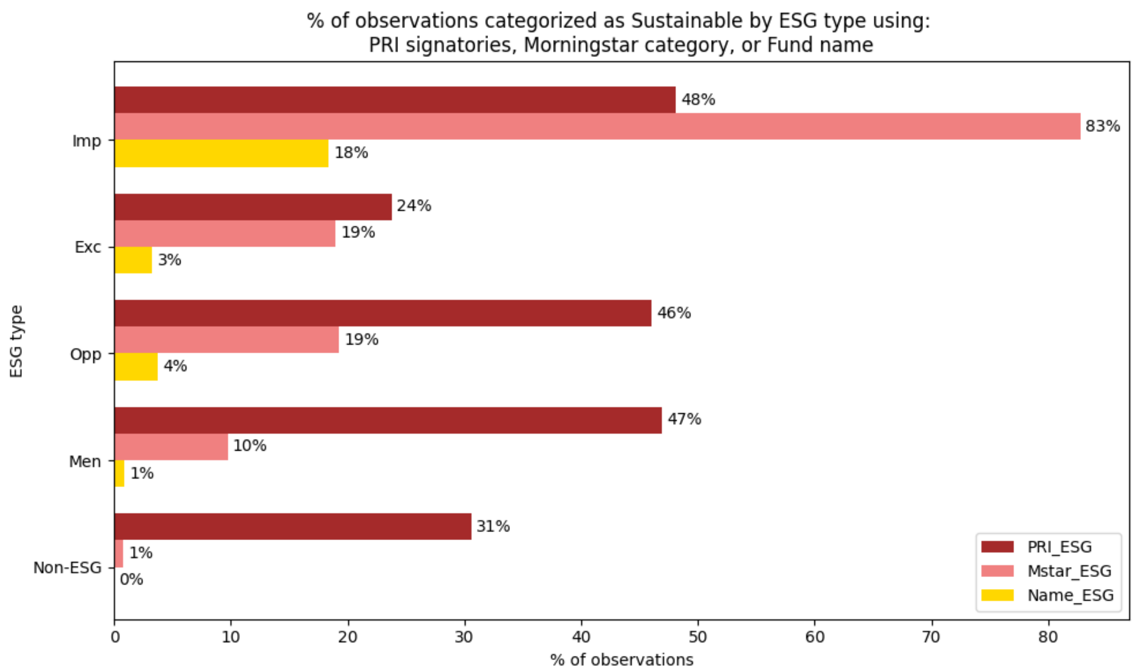

Figure 1: Classification Comparisons

This figure displays, for each ESG type identified by our classification, the percentage of fund-month observations that would be classified as sustainable under three alternative methodologies: Morningstar's Sustainability categorization (variable: *Sustainable Investment Overall*, available from 2020); the presence of sustainability-related terms in the fund name (see Internet Appendix I for the full list of terms); and whether the fund's family is a signatory to the Principles for Responsible Investment (PRI). Each bar represents the overlap between a given ESG type in our classification and a given alternative method.

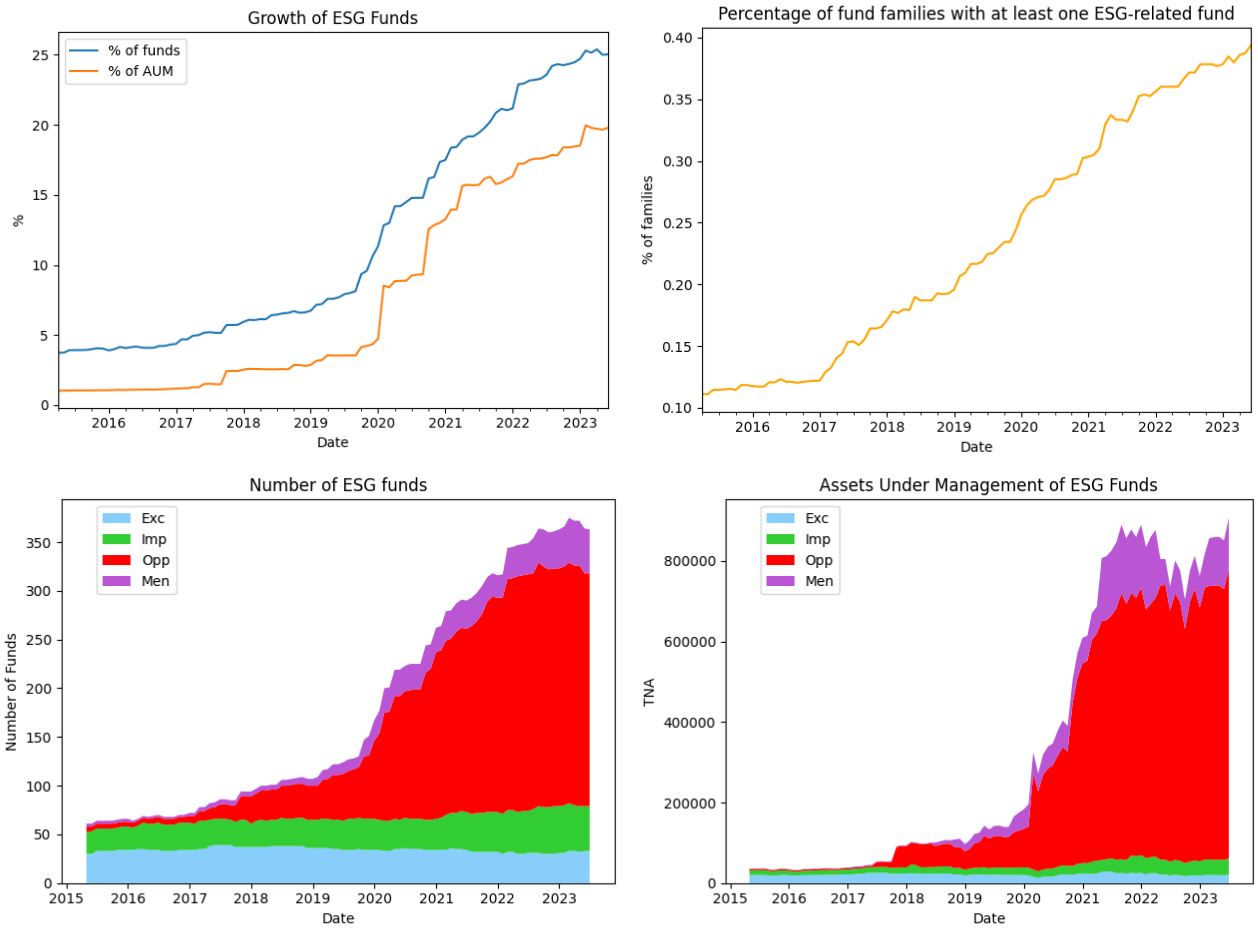

Figure 2: Growth of ESG Funds

This figure displays the time-series evolution of ESG investing among active equity US mutual funds over the period 2015–2023. The top-left panel plots the percentage of ESG funds in terms of number and AUM. The top-right panel plots the percentage of fund families with at least one ESG fund. The bottom panels decompose the number of ESG funds (left) and their total net assets in \$1000s (right) by ESG category.

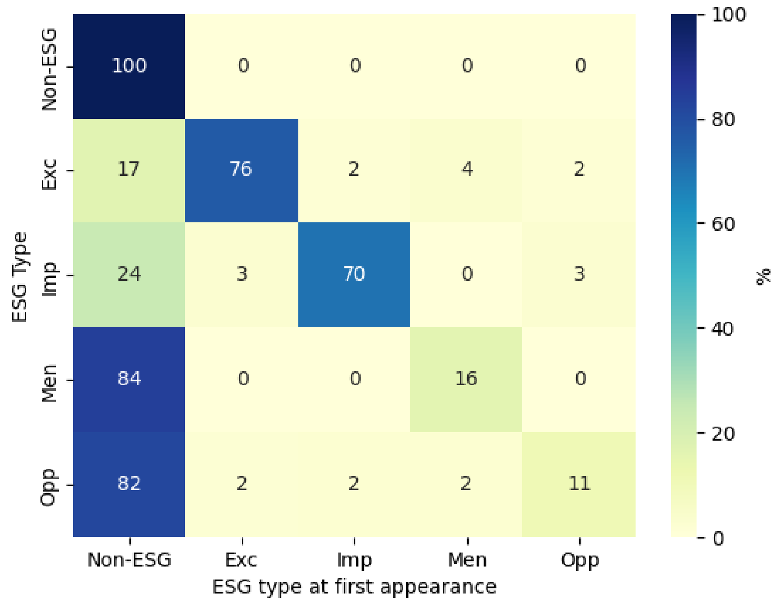

Figure 3: New versus Existing Funds

This figure displays the percentage of funds in each ESG type (rows) broken down by their ESG type at first appearance in our dataset (columns). Diagonal entries indicate the share of funds that were classified in their current type from inception.

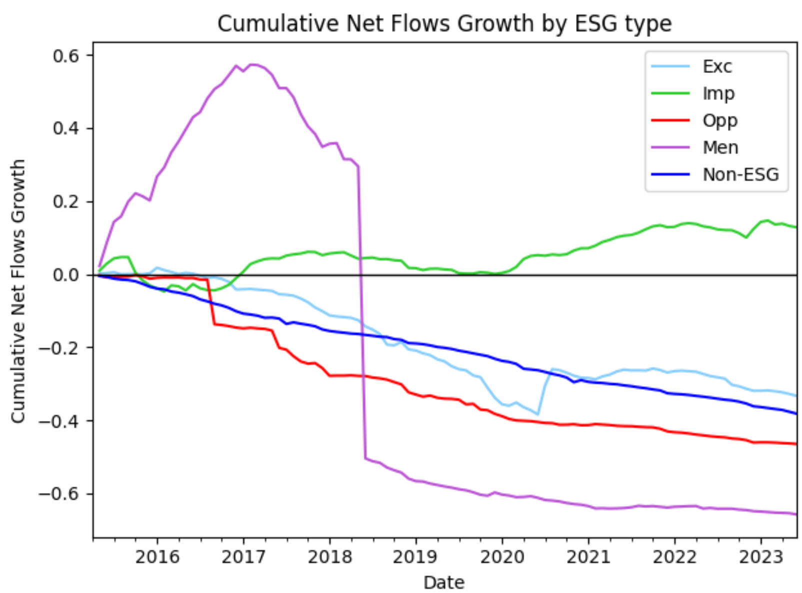

Figure 4: Capital Flows to ESG Funds

This figure displays the cumulative percentage change in AUM due solely to net flows, by ESG type. The measure compounds monthly net-flow growth rates starting from the first month of our sample.

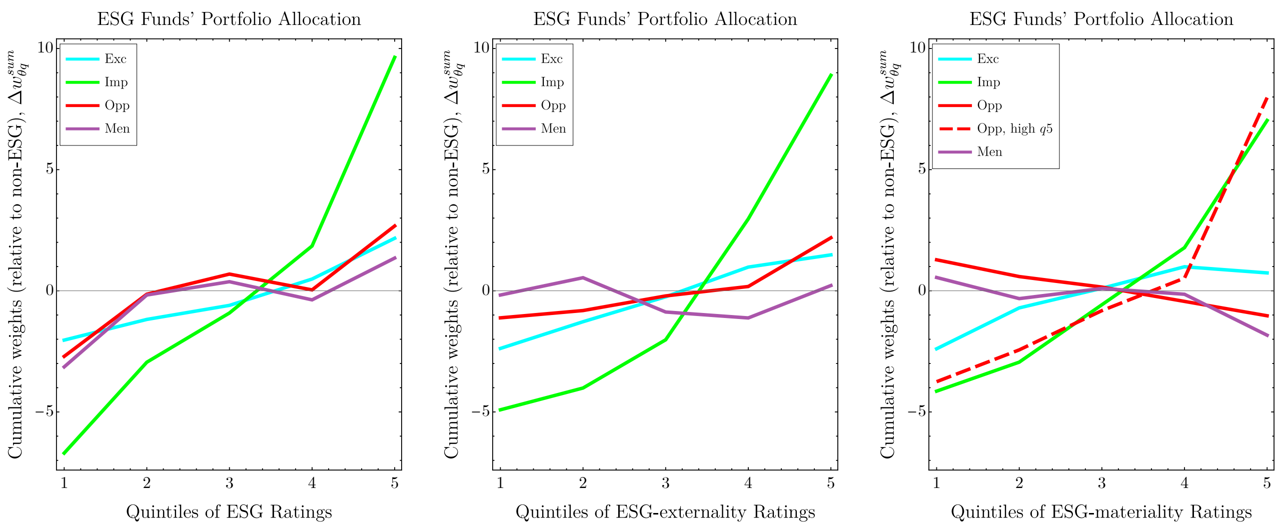

Figure 5: Fund Holdings Across ESG-rating Quintiles

This figure plots the estimated relative cumulative portfolio weights $\Delta\hat{w}_{\theta q}^{sum}$ in (11) for each ESG fund type across quintiles of overall ESG ratings (left panel), ESG-externality ratings (center panel), and ESG-materiality ratings (right panel). Values represent the difference in cumulative portfolio weight held in each quintile relative to non-ESG funds in the same Morningstar category and month, after controlling for fund-level characteristics. In the right panel, "Opp, high $q5$" denotes the subgroup of opportunity funds with the highest holdings in the top quintile of ESG-materiality ratings. The underlying estimates and their statistical significance are reported in Table 5.

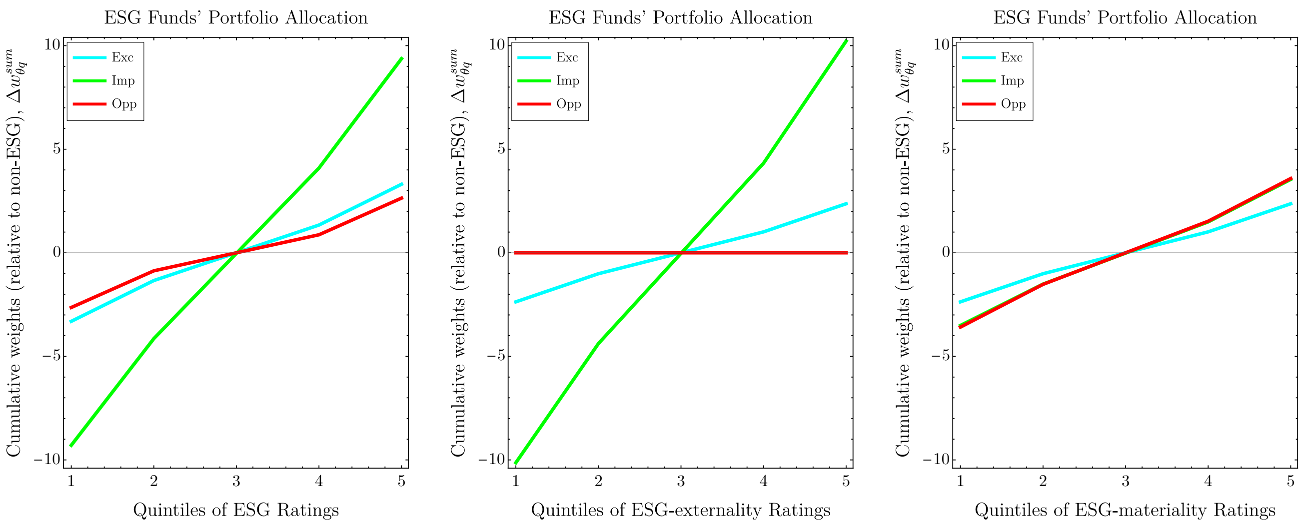

Figure 6: Fund Holdings Across ESG-Rating Quintiles: Theoretical Predictions

This figure plots the theoretical portfolio weights within each quintile of overall ESG ratings (left panel), ESG-externality ratings (center panel), and ESG-materiality ratings (right panel), expressed as deviations from the allocations of non-ESG funds.

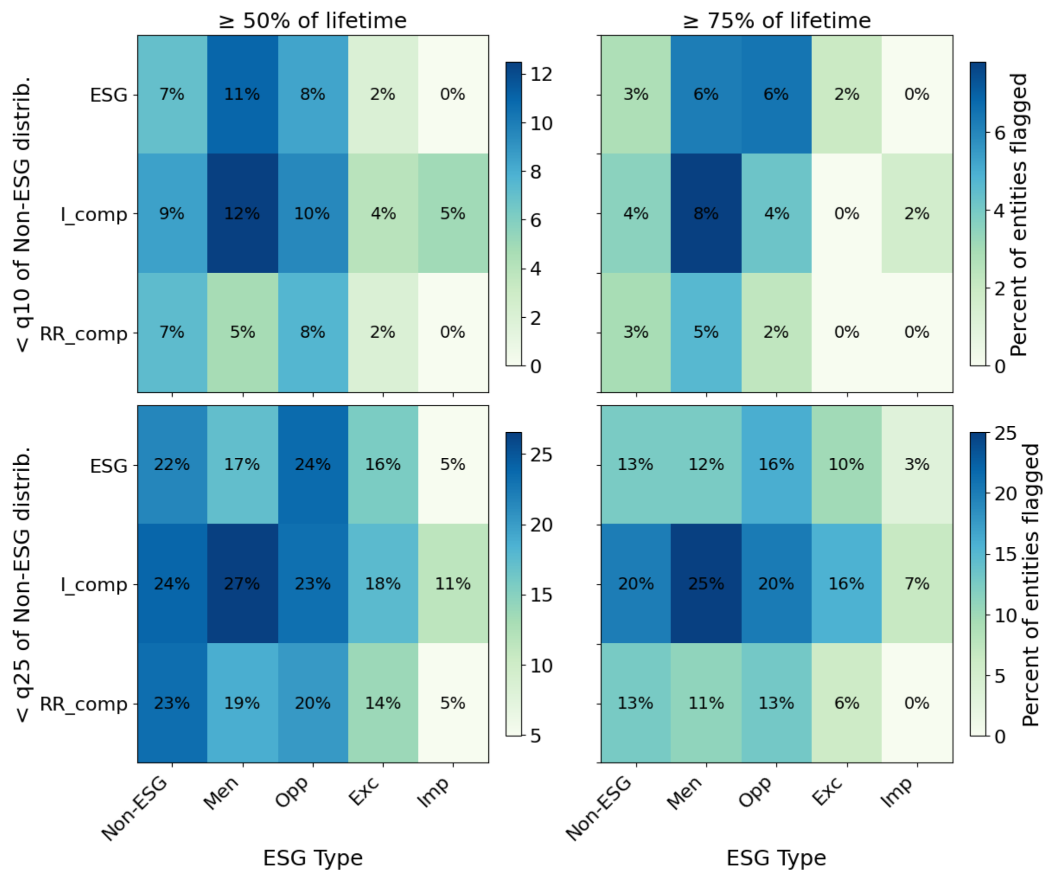

Figure 7: Greenwashing

This figure reports the percentage of funds in each ESG type flagged as potential greenwashers, defined as funds whose portfolio-level ESG rating falls below a given percentile of the non-ESG fund distribution for more than a specified fraction of the fund's lifetime. Rows correspond to three rating measures: overall ESG, ESG-externality ($ESGx$), and ESG-materiality ($ESGm$). Columns correspond to ESG fund types. The top panels use the 10th percentile threshold; the bottom panels use the 25th percentile. The left panels require the condition to hold for at least 50% of the fund's lifetime; the right panels require at least 75%.

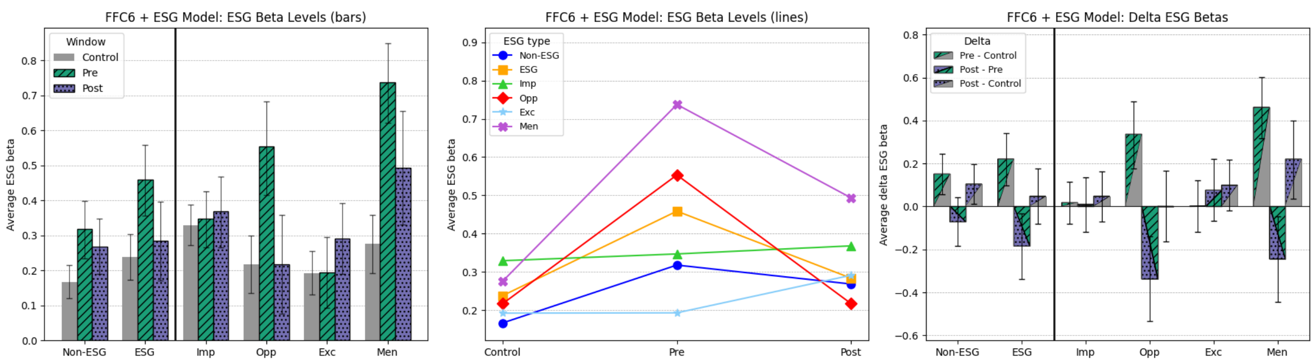

Figure 8: ESG Window Dressing

This figure displays average ESG factor loadings estimated from the FFC6+ESG model with 10 trading-day pre- and post-disclosure windows. The left panel plots ESG beta levels by window (control, pre-disclosure, post-disclosure) for each fund type. The center panel connects the three windows for each fund type to illustrate the dynamics of ESG exposure around disclosure dates. The right panel plots the delta ESG betas (pre minus control, post minus pre, post minus control) with bootstrap confidence intervals. Table 6 reports the underlying estimates.

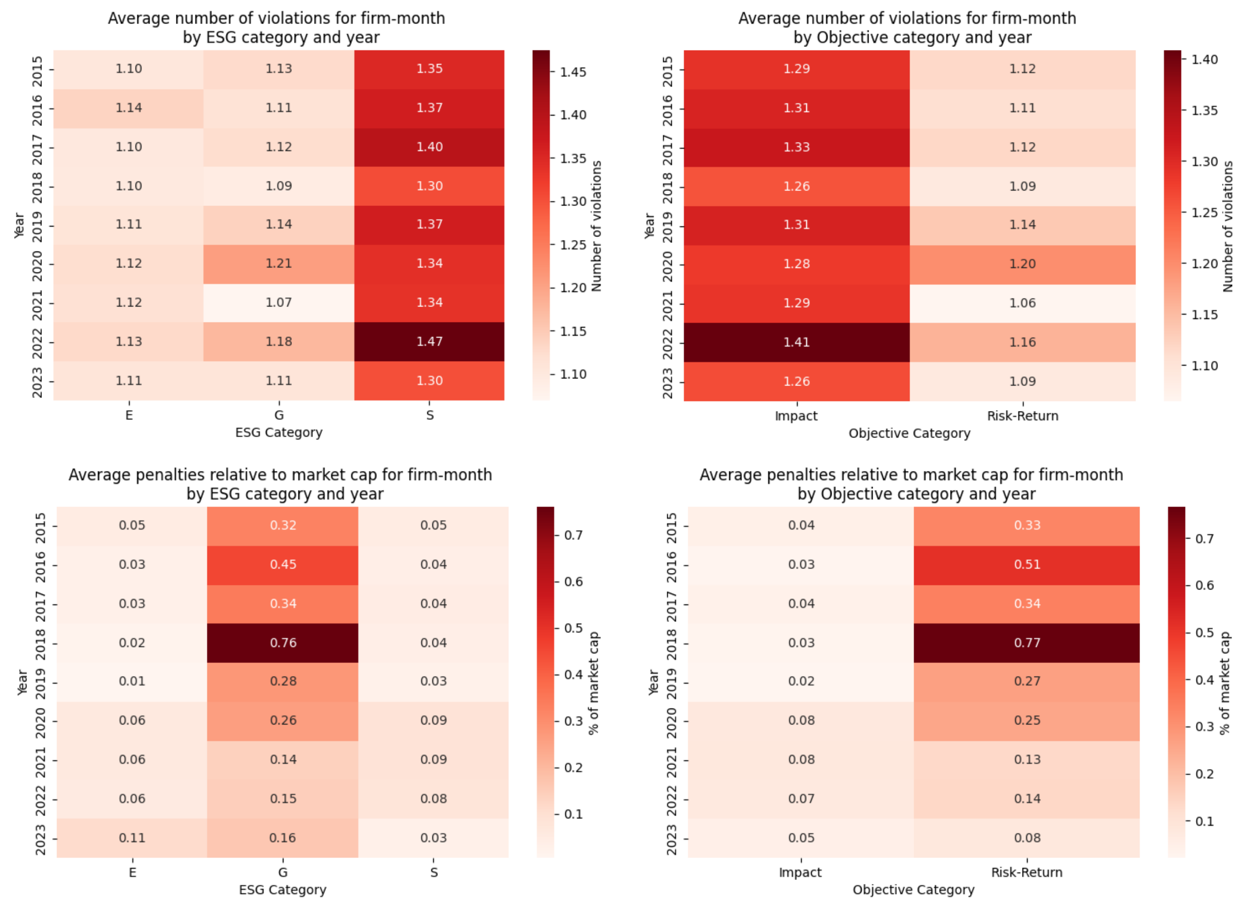

Figure 9: ESG Violations

This figure summarizes the ESG violation data over the period 2015–2023. The top panels display the average number of violations per firm-month, broken down by ESG pillar (left) and by our externality-materiality classification (right). The bottom panels display the average penalty amount as a percentage of market capitalization, with the same breakdowns. Violations are classified into E, S, and G pillars and into externality and materiality categories using ChatGPT 4.1, as described in Section 5.3.

Table 1: Classification Comparison: False Positives and False Negatives

This table compares our classification of ESG funds with three alternative methodologies commonly used in the literature: Morningstar's Sustainability categorization (variable: *Sustainable Investment Overall*); a classification based on the presence of sustainability-related terms in the fund name (see Internet Appendix I for the full list of terms); and whether the fund's family is a signatory to the Principles for Responsible Investment (PRI). We report misclassification rates using our classification as the benchmark. In Panel A, we define as correctly classified any fund that our method categorizes as impact. In Panel B, we expand this definition to also include exclusionary funds. False positive rates are the percentage of observations classified as sustainable by the alternative method that do not belong to the relevant categories under our classification. False negative rates are the percentage of observations belonging to the relevant categories under our classification that are not identified as sustainable by the alternative method. Statistics for the Fund Name and PRI methodologies cover the full sample period 2015–2023, while those for Morningstar cover the sample period 2020–2023.

Panel A: Impact Funds

	Morningstar	Fund Name	PRI Signatory
False Positive (Type 1 Error)	50%	46%	97%
False Negative (Type 2 Error)	20%	82%	52%

Panel B: Impact or Exclusionary Funds

	Morningstar	Fund Name	PRI Signatory
False Positive (Type 1 Error)	41%	36%	95%
False Negative (Type 2 Error)	48%	89%	64%

Table 2: Fund Level ESG Ratings

This table reports estimates of the regression specification in (6), where the dependent variable is one of four measures of fund-level portfolio ESG ratings: value-weighted, over-weighted, under-held, or not-held, as defined in equations (2)–(5). ESG ratings are cross-sectionally standardized within month as in equation (1). In the top panel, $\mathbf{1}_{jt}^{ESG}$ is a single dummy for ESG funds. In the bottom panel, it is replaced by separate dummies for each ESG category (Exc, Imp, Opp, Men). Non-ESG funds are the excluded category. Controls include log age, log assets, expenses, turnover, flows, flow volatility, and fund betas with respect to Fama–French–Carhart six factors. All specifications include Morningstar category × month fixed effects. Standard errors are clustered by fund and month; t-statistics are reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Portfolio ESG Ratings			
	value-weighted (1)	over-weighted (2)	under-held (3)	not-held (4)
ESG	0.087***	0.095***	0.059*	-0.024***
	(6.74)	(6.84)	(1.85)	(-3.16)
Adj. R^2	0.541	0.469	0.146	0.643
Exc	0.064***	0.066***	0.157**	-0.011
	(2.93)	(2.82)	(2.17)	(-1.09)
Imp	0.244***	0.264***	0.100	-0.079***
	(8.58)	(8.04)	(1.27)	(-3.89)
Opp	0.047***	0.052***	0.017	-0.014
	(3.11)	(3.29)	(0.46)	(-1.45)
Men	0.048**	0.065**	0.012	-0.000
	(2.23)	(2.58)	(0.18)	(-0.02)
Adj. R^2	0.550	0.478	0.147	0.644
Controls	Y	Y	Y	Y
MStar×Month FE	Y	Y	Y	Y
Obs.	143,635	143,634	92,030	143,637
$\bar{y}_{\text{Non-ESG}}$	0.225	0.207	0.437	0.286

Table 3: ESG-Externality and ESG-Materiality Ratings

This table reports estimates of the regression specification in (9). Models (1) and (3) are unconstrained; Models (2) and (4) impose the constraint $\beta_x + \beta_m = 1$, where $\hat{\beta}_m$ is obtained as $1 - \hat{\beta}_x$ and its standard error is not separately defined. Models (3) and (4) orthogonalize $z_{it}^{ESG^m}$ with respect to $z_{it}^{ESG^x}$ before estimation. t-statistics are reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	z_{it}^{ESG}			
	(1)	(2)	(3)	(4)
z_{it}^{ESGx}	0.403***	0.430***	0.507***	0.482***
	(0.012)	(0.001)	(0.012)	(0.001)
z_{it}^{ESGm}	0.543***	0.570	0.543***	0.518
	(0.010)	(.)	(0.010)	(.)
Constant	N	N	N	N
Constrained	N	Y	N	Y
Orthogonalized	N	N	Y	Y
AdjR2	0.54	0.39	0.54	0.50
Obs	241,864	241,864	241,864	241,864

Table 4: Fund-Level ESG-Externality and ESG-Materiality Ratings

This table reports estimates of the regression specification in (6), where the dependent variable is one of four measures of fund-level portfolio ratings—value-weighted, over-weighted, under-held, or not-held—constructed using the ESG-externality rating $z_{it}^{ESG^x}$ (columns 1–4) or the ESG-materiality rating $z_{it}^{ESG^m}$ (columns 5–8), as defined in Section 4.2. In the top panel, $\mathbf{1}_{jt}^{ESG}$ is a single dummy for ESG funds. In the bottom panel, it is replaced by separate dummies for each ESG category (Exc, Imp, Opp, Men). Non-ESG funds are the excluded category. Controls and fixed effects are as in Table 2. Standard errors are clustered by fund and month; t-statistics are reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Portfolio ESG-externality Rating				Portfolio ESG-materiality Rating			
	value-weighted	over-weighted	under-held	not-held	value-weighted	over-weighted	under-held	not-held
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ESG	0.041***	0.049***	0.044**	0.014	0.059***	0.062***	0.037	-0.014**
	(4.01)	(4.13)	(2.10)	(1.46)	(6.87)	(6.39)	(1.51)	(-2.20)
Adj. R^2	0.870	0.818	0.403	0.821	0.807	0.741	0.538	0.823
Exc	0.035	0.035	0.021	0.067***	0.048***	0.037**	0.114**	0.001
	(1.50)	(1.32)	(0.47)	(3.26)	(3.78)	(2.40)	(2.30)	(0.06)
Imp	0.140***	0.146***	0.047	-0.004	0.157***	0.160***	0.089	-0.041**
	(5.98)	(4.97)	(0.92)	(-0.13)	(7.39)	(6.14)	(1.57)	(-2.25)
Opp	0.015	0.024*	0.051*	-0.001	0.033***	0.038***	-0.007	-0.012*
	(1.23)	(1.77)	(1.91)	(-0.15)	(3.11)	(3.31)	(-0.25)	(-1.68)
Men	0.008	0.020	0.048	0.019	0.034**	0.055***	0.022	0.000
	(0.38)	(0.81)	(0.85)	(1.53)	(2.47)	(3.40)	(0.38)	(0.05)
Adj. R^2	0.871	0.819	0.403	0.822	0.809	0.742	0.539	0.824
Controls	Y	Y	Y	Y	Y	Y	Y	Y
MStar×Month FE	Y	Y	Y	Y	Y	Y	Y	Y
Obs.	143,635	143,634	92,017	143,637	143,635	143,634	92,030	143,637
$\bar{y}_{\text{Non-ESG}}$	0.516	0.466	0.816	0.594	0.304	0.284	0.521	0.380

Table 5: Distribution of Fund Holdings Across ESG Ratings

This table reports estimates of the regression specification in (10), where the dependent variable is the cumulative portfolio weight held by fund j in stocks belonging to quintile q of ESG ratings in month t. The omitted categories are non-ESG funds and the middle quintile ($q = 3$). Reported coefficients represent the relative cumulative weights $\Delta\hat{w}^{sum}_{\omega q}$ as defined in equation (11). Panel A uses quintiles of overall ESG ratings; Panel B uses quintiles of ESG-externality ratings; Panel C uses quintiles of ESG-materiality ratings. In Panel C, "Opp (high q5)" denotes the subgroup of opportunity funds with the highest holdings in the top quintile of ESG-materiality ratings. Controls and fixed effects are as in Table 2. Standard errors are clustered by fund and month; t-statistics are reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	q1 (1)	q2 (2)	q3 (3)	q4 (4)	q5 (5)
Panel A: Quintiles of ESG Ratings					
Exc	-2.035*** (0.757)	-1.182** (0.463)	-0.602 (0.377)	0.485 (0.455)	2.168** (0.985)
Imp	-6.700*** (0.616)	-2.948*** (0.539)	-0.916** (0.419)	1.849*** (0.478)	9.635*** (1.381)
Opp	-2.707*** (0.541)	-0.141 (0.328)	0.690*** (0.235)	0.037 (0.305)	2.676*** (0.663)
Men	-3.141*** (0.848)	-0.174 (0.597)	0.375 (0.453)	-0.374 (0.676)	1.351 (1.163)
Panel B: Quintiles of ESG-externality Ratings					
Exc	-2.380* (1.235)	-1.277 (1.049)	-0.248 (0.638)	0.978 (0.931)	1.483 (1.894)
Imp	-4.912*** (1.038)	-4.013*** (1.176)	-2.023*** (0.705)	2.967*** (0.762)	8.893*** (2.090)
Opp	-1.117* (0.650)	-0.816 (0.581)	-0.212 (0.371)	0.177 (0.460)	2.189** (1.047)
Men	-0.175 (1.566)	0.540 (1.376)	-0.876 (0.748)	-1.121 (1.079)	0.217 (2.560)
Panel C: Quintiles of ESG-materiality Ratings					
Exc	-2.394** (1.146)	-0.706 (0.582)	0.104 (0.468)	0.992 (0.645)	0.738 (1.131)
Imp	-4.143*** (0.882)	-2.945*** (0.705)	-0.576 (0.467)	1.780*** (0.531)	7.031*** (1.571)
Opp	1.279** (0.517)	0.587 (0.385)	0.154 (0.275)	-0.438 (0.321)	-1.031 (0.695)
Men	0.549 (0.9930)	-0.323 (0.8024)	0.096 (0.5348)	-0.143 (0.5367)	-1.834 (1.3611)

Table 6: ESG Window Dressing

This table reports estimates of ESG factor loadings around quarterly portfolio disclosure dates, by fund type. ESG betas are estimated from the FFC6+ESG model within three non-overlapping windows: a control window spanning the second month of the quarter, a pre-disclosure window covering the 10 trading days before the disclosure date, and a post-disclosure window covering the 10 trading days beginning two days after. Panel A reports level ESG betas for each window (columns 1-3) and their differences (columns 4-6): pre minus control in (26), post minus pre in (27), and post minus control in (28). Panel B reports the corresponding differences relative to non-ESG funds. Inference is based on a wild bootstrap procedure with 9,999 iterations; bootstrap standard errors are reported in parentheses. For pre minus control and post minus pre, significance stars reflect one-sided p-values; for levels and post minus control, they reflect two-sided p-values. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Levels of ESG Betas			Changes in ESG Betas		
	Control (1)	Pre (2)	Post (3)	Pre − Control (4)	Post − Pre (5)	Post − Control (6)
Panel A: Estimates						
Non-ESG	0.167***	0.318***	0.269***	0.151***	-0.072	0.104**
	(0.024)	(0.042)	(0.040)	(0.048)	(0.059)	(0.047)
ESG	0.238***	0.459***	0.284***	0.221***	-0.185***	0.048
	(0.033)	(0.052)	(0.057)	(0.061)	(0.078)	(0.066)
Exc	0.193***	0.193***	0.292***	0.001	0.076	0.099
	(0.032)	(0.052)	(0.052)	(0.061)	(0.074)	(0.060)
Imp	0.329***	0.347***	0.368***	0.018	0.010	0.048
	(0.030)	(0.041)	(0.051)	(0.050)	(0.065)	(0.059)
Opp	0.217***	0.553***	0.217***	0.336***	-0.338***	-0.000
	(0.042)	(0.069)	(0.072)	(0.080)	(0.100)	(0.083)
Men	0.276***	0.737***	0.493***	0.462***	-0.245***	0.220**
	(0.043)	(0.059)	(0.083)	(0.073)	(0.102)	(0.093)
Panel B: Estimate Differences Relative to Non-ESG						
ESG	0.071***	0.141***	0.015	0.070	-0.113*	-0.057
	(0.024)	(0.041)	(0.041)	(0.048)	(0.059)	(0.048)
Exc	0.026	-0.125***	0.023	-0.150***	0.148***	-0.005
	(0.020)	(0.045)	(0.032)	(0.049)	(0.056)	(0.038)
Imp	0.163***	0.029	0.100***	-0.134***	0.082	-0.057
	(0.022)	(0.035)	(0.037)	(0.042)	(0.052)	(0.043)
Opp	0.050	0.235***	-0.051	0.185***	-0.266***	-0.105
	(0.034)	(0.057)	(0.059)	(0.067)	(0.083)	(0.068)
Men	0.109***	0.419***	0.225***	0.310***	-0.173*	0.116
	(0.036)	(0.054)	(0.074)	(0.065)	(0.092)	(0.082)

Table 7: Fund-Level Exposure to ESG Violations

This table reports estimates of the regression specification in (30), where the dependent variable is the portfolio-weighted share of stocks with at least one ESG violation in month t, as defined in equation (29). Column (1) uses all ESG violations (V_{it}); column (2) uses ESG-externality violations (V_{it}^{x}); column (3) uses ESG-materiality violations (V_{it}^{m}). The classification of violations into externality and materiality categories is described in Section 5.3. The omitted category is non-ESG funds. Controls include log age, log assets, expenses, turnover, flows, flow volatility, and fund betas with respect to Fama–French–Carhart six factors. All specifications include Morningstar category $\times$ month fixed effects. Standard errors are clustered by fund; t-statistics are reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	All violations (1)	ESG-externality violations (2)	ESG-materiality violations (3)
Exc	-0.011**	-0.009	-0.004**
	(-1.98)	(-1.76)	(-2.50)
Imp	-0.022***	-0.019***	-0.005***
	(-4.22)	(-4.41)	(-2.70)
Opp	-0.016***	-0.014***	-0.004***
	(-5.34)	(-4.53)	(-3.70)
Men	-0.005	-0.004	-0.001
	(-1.51)	(-1.38)	(-0.63)
Adj. R^2	0.727	0.693	0.592
Controls	Y	Y	Y
MStar$\times$Month FE	Y	Y	Y
Obs.	146,595	146,595	146,595
$\bar{y}_{\text{Non-ESG}}$	0.138	0.120	0.026

References

Andrikogiannopoulou, Angie, Philipp Krueger, Shema Frédéric Mitali, and Filippos Papakonstantinou, 2022, Discretionary Information in ESG Investing: A Text Analysis of Mutual Fund Prospectuses, SSRN Working Paper 4082263.

Avramov, Doron, Abraham Lioui, Yang Liu, and Andrea Tarelli, 2024, Dynamic ESG Equilibrium, *Management Science* forthcoming.

Baker, Malcolm, Mark L. Egan, and Suproteem K. Sarkar, 2022, How Do Investors Value ESG?, NBER Working Paper 30708.

Berk, Jonathan, and Jules H. van Binsbergen, 2024, The Impact of Impact Investing, *Journal of Financial Economics* forthcoming.

Broccardo, Eleonora, Oliver Hart, and Luigi Zingales, 2022, Exit versus Voice, *Journal of Political Economy* 130, 3101–3145.

Ceccarelli, Marco, Simon Glossner, and Mikael Homanen, 2022, Catering through Transparency: Voluntary ESG Disclosure by Asset Managers and Fund Flows, *SSRN Electronic Journal*.

Cremers, K. J. Martijn, and Antti Petajisto, 2009, How Active is Your Fund Manager? A New Measure That Predicts Performance, *Review of Financial Studies* 22, 3329–3365.

Cremers, Martijn, Timothy B. Riley, and Rafael Zambrana, 2024, The Complex Materiality of ESG Ratings: Evidence from Actively Managed ESG Funds, SSRN Working Paper 4335688.

Dasgupta, Amil, Dirk Jenter, Richmond Mathews, and Paul Voss, 2026, Are socially responsible funds viable?, .

Edmans, Alex, Tom Gosling, and Dirk Jenter, 2026, Sustainable investing in practice: Objectives, beliefs, and limits to impact, .

Evans, Richard B., 2010, Mutual Fund Incubation, *Journal of Finance* 65, 1581–1611.

Farroukh, Abed El Karim, Jarrad Harford, and David (Dongheon) Shin, 2023, What Does ESG Investing Mean and Does It Matter Yet?, SSRN Working Paper 4484626.

Frankel, Richard M., Xiumin Martin, Hanmeng Ivy Wang, and Mengting Ashley Yu, 2024, Mutual Fund ESG Disclosures, SSRN Working Paper 4688461.

Ghoul, Sadok El, and Aymen Karoui, 2022, Fund Performance and Social Responsibility: New Evidence using Social Active Share and Social Tracking Error, *Journal of Banking & Finance* 143, 106598.

Gibson Brandon, Rajna, Simon Glossner, Philipp Krueger, Pedro Matos, and Tom Steffen, 2022, Do Responsible Investors Invest Responsibly?, *Review of Finance* 26, 1389–1432.

Goldstein, Itay, Alexandr Kopytov, Lin Shen, and Haotian Xiang, 2022, On ESG Investing: Heterogeneous Preferences, Information, and Asset Prices, NBER Working Paper 29839.

Heinkel, Robert, Alan Kraus, and Josef Zechner, 2001, The Effect of Green Investment on Corporate Behavior, *Journal of Financial and Quantitative Analysis* 36, 431–449.

Hong, Harrison, and Marcin Kacperczyk, 2009, The Price of Sin: The Effects of Social Norms on Markets, *Journal of Financial Economics* 93, 15–36.

Inderst, Roman, and Marcus M Opp, 2025, Sustainable finance versus environmental policy: Does greenwashing justify a taxonomy for sustainable investments?, *Journal of financial economics* 163, 103954.

Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng, 2008, Unobserved Actions of Mutual Funds, *Review of Financial Studies* 21, 2379–2416.

Kim, Soohun, and Aaron Yoon, 2023, Analyzing Active Fund Managers' Commitment to ESG: Evidence from the United Nations Principles for Responsible Investment, *Management Science* 69, 741–758.

Lowry, Michelle B., Pingle Wang, and Kelsey D. Wei, 2025, Are All ESG Funds Created Equal? Only Some Funds Are Committed, *Forthcoming, Review of Financial Studies*.

Morningstar, 2022, Morningstar Sustainable Attributes, Report.

Oehmke, Martin, and Marcus M. Opp, 2024, A Theory of Socially Responsible Investment, .

Parise, Gianpaolo, and Mirco Rubin, 2025, Green window dressing, *The Journal of Finance* 80, 3555–3588.

Pastor, Lubos, Robert F. Stambaugh, and Lucian A. Taylor, 2024, Green Tilts, NBER Working Paper 31320.

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

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2015, Scale and Skill in Active Management, *Journal of Financial Economics* 116, 23–45.

——— , 2021, Sustainable investing in equilibrium, *Journal of Financial Economics* 142, 550–571.

Raghunandan, Aneesh, and Shivaram Rajgopal, 2022, Do ESG Funds Make Stakeholder-Friendly Investments?, *Review of Accounting Studies* 27, 822–863.

Shive, Sophie, and Hayong Yun, 2013, Are mutual funds sitting ducks?, *Journal of Financial Economics* 107, 220–237.

Starks, Laura T., 2023, Presidential Address: Sustainable Finance and ESG Issues—Value versus Values, *Journal of Finance* 78, 1837–1872.

van der Beck, Philippe, 2023, Flow-Driven ESG Returns, .

Yang, Keer, and Ayako Yasuda, 2025, Decoding sustainable investment strategies: Bridging intentions and outcomes, .

Zerbib, Olivier David, 2022, A Sustainable Capital Asset Pricing Model (S-CAPM): Evidence from Environmental Integration and Sin Stock Exclusion, *Review of Finance* 26, 1345–1388.

Zhu, Qifei, 2020, The Missing New Funds, *Management Science* 66, 1193–1204.

INTERNET APPENDIX
"Different Shades of ESG Funds"

by Simona Abis, Andrea M. Buffa and Meha Sadasivam

I ESG Dictionary

ESG-related terminology. ESG, environmental, carbon, SRI, human rights, green, climate change, renewable energy, social responsibility, pollution, sustainable business practice, sustainable development goals, clean energy, SDG, greenhouse emissions, greenhouse gas emissions, cluster munitions, child labor, alternative energy, solar, adopt more responsible, adopt responsible, implement responsible, responsible supply, biblically responsible, catholic church, islam, toxic, irresponsible marketing, gambling, tobacco, palm oil, pesticides, small arms, thermal coal, nuclear, alcohol, animal testing, adult entertainment, fossil fuel, weapons, specialty leather, genetically modified organism, military contracting, abortion, stem cells, pornography, firearm, handgun, liquor, casino, anti-family, terrorism, social screens, environmental screens, governance screens.

ESG-terms searched in fund names. ESG, SRI, SDG, CSR, environment, social, governance, impact, climate, equality, ethic, responsible, green, clean, sustain, renew, diversity, women, low carbon, progressive energy, fossil fuel free, alternative energy, energy solution, eco leader, ecological strategy.

II Examples of Prospectus Classifications

This section reports example extracts from PIS sections that were classified in each of the ESG fund types defined in Section 2.

II.A Exclusionary

Fund name: Alger Responsible Investing Fund; Date: 2018-03-31

> under normal circumstances, the fund invests at least % of its net assets, plus any borrowings for investment purposes, in equity securities of companies of any size with an environmental, social and governance esg rating of bb or above by msci or an equivalent rating by another esg rating agency that also demonstrate, in the view of fred alger management, inc. , promising growth potential. (...) fred alger management, inc. employs fundamental analysis to identify innovative and dynamic companies and uses msci's esg ratings to consider how such stocks rank within an industry or sector based on a company's conduct in offering products or services that promote positive

environmental, social and/or governance policies, or have a positive impact in these areas, addressing concerns such as climate change, resource depletion, health and safety, employee relations and diversity, bribery and corruption, and fostering board diversity and structure. the fund does not bar companies in any industries.

Fund name: Northern US Quality ESG Fund; Date: 2019-08-31

in seeking long-term capital appreciation, the fund will invest, under normal circumstances, at least % of its net assets plus borrowings for investment purposes in equity securities of large and mid-capitalization United States companies that nti believes have favorable environmental, social and governance esg characteristics under a third-party vendors rating methodology. (...) Using a quantitative, factor based approach, the fund intends to invest in companies that: i meet certain criteria for esg factors as provided by a third-party research vendor; ii exhibit strong business fundamentals, solid management and reliable cash flows; and iii are located, headquartered in, incorporated in or otherwise organized in the united states. (...) the fund is managed according to a quantitative model developed by nti to define an investable universe, nti excludes securities of companies involved in esg controversies or those that violate global norms like the united nations global compact. nti also removes companies that do a poor job of managing their esg risks and opportunities relative to their peers as well as those with material involvement in controversial business practices, including, but not limited to, tobacco and civilian firearms. nti engages a third-party research vendor to provide esg data for United States companies. the third-party vendor identifies esg areas of risk and opportunity, evaluates exposure and management, and ranks and rates companies against their industry peers. after defining the investable universe, nti evaluates the quality of the remaining securities and removes those securities that do not meet the proprietary quality methodology. (...) the fund is constructed based on an optimization methodology designed to take active exposure by overweighting and underweighting securities based on their esg and relative financial quality rankings. (...) further, the carbon footprint of the portfolio is reduced relative to the companies in the benchmark russell index.

Fund name: Wells Fargo Large Cap Core Fund; Date: 2020-11-30

(...) generally, we avoid investments in issuers we deem to have significant alcohol, gaming or tobacco business.

II.B Impact

Fund name: Eventide Gilead Fund, date: 2017-04-30

The funds investment adviser, eventide asset management, llc eventide or the adviser analyzes the performance of potential investments not only for financial strengths and outlook, but also for the company's ability to operate with integrity and create value for customers, employees, and other stakeholders. While few companies may reach these ideals in every area of their business, these principles articulate the advisers highest expectations for corporate behavior. (...) The adviser seeks to invest in companies that reflect the following values: respecting the value and freedom of all people; this includes the right to life at all stages and freedom from addictive behaviors caused by gambling, pornography, tobacco and alcohol. Demonstrating a concern for justice and peace through fair and ethical relationships with customers, suppliers and business partners and through avoidance of products and services that promote weapons production and proliferation. Promoting family and community; this includes protecting children from violent forms of entertainment and also includes serving low income communities. Exhibiting responsible management practices, including fair-dealing with employees, communities, competitors, suppliers and customers as demonstrated by a company's record regarding litigation, regulatory actions against the company and its record of promoting products and services that improve the lives of people. Practicing environmental stewardship; this includes practices considered more sustainable than those of industry peers, reduction in environmental impact when compared to previous periods, and/or the use of more efficient and cleaner energy sources. Consistent with the advisers values, the fund may also invest in community development institutions that serve the financial needs of low-to-moderate income families and communities.

Fund name: Domini Impact Equity Fund; Date: 2018-12-31

The fund will invest in companies that Domini Impact Investments llc Domini believes have strong environmental and social profiles. The fund may also invest in companies that Domini believes help create products and services that provide sustainability solutions and are evaluated using fundamental analysis. The fund may sell a security if the issuer fails to meet Domini's social and environmental standards or sustainability themes. (...) While pursuing their financial objectives, impact investors seek to use their investments to create a more fair and sustainable world. Domini believes that by factoring social and environmental sustainability standards into their investment decisions, investors can encourage greater corporate accountability. Domini evaluates the funds potential investments against its social and environmental standards based on the businesses in which an issuer engages, as well as on the quality of the issuers relations with key stakeholders, including communities, customers, ecosystems, employees, investors, and suppliers. Domini's interpretation and application of its social and environmental standards are subjective and may evolve over time.

Fund name: Neuberger Berman Sustainable Equity Fund; Date: 2020-12-31

The portfolio managers employ a research driven and valuation sensitive approach to stock selection, with a focus on long term sustainability. This sustainable investment approach seeks to identify high quality, well-positioned companies with leadership that is focused on ESG as defined by best in class operating practices. (...) Among companies that meet these criteria, the portfolio managers look for those that show leadership in environmental, social and governance considerations, including progressive workplace practices and community relations. In addition, the portfolio managers typically look at a company's record in public health and the nature of its products. The portfolio managers judge firms on their corporate citizenship overall, considering their accomplishments as well as their goals. While these judgments are inevitably subjective, the fund endeavors to avoid companies that derive revenue from gambling or the production of alcohol, tobacco, weapons, or nuclear power. The fund also does not invest in any company that derives its total revenue primarily from non-consumer sales to the military.

II.C Opportunity

Fund name: UBS US Small Cap Growth Fund; Date: 2020-10-31

> The fund is classified by UBS am Americas as an ESG-integrated fund. The fund's investment process integrates material sustainability and/or environmental, social and governance ESG considerations into the research process. ESG integration is driven by taking into account material ESG risks which could impact investment returns, rather than being driven by specific ethical principles or norms. The analysis of material sustainability/ESG considerations can include many different aspects, including, for example, the carbon footprint, employee health and well-being, supply chain management, fair customer treatment and governance processes of a company. The fund's portfolio managers may still invest in securities with a higher ESG risk profile where the portfolio managers believe the potential compensation outweighs the risks identified.

Fund name: Artisan Value Fund; Date: 2022-01-31

> As part of the teams analysis of a company's business prospects, among other factors, the team considers certain environmental, social and governance ESG factors relating to the company. These ESG factors may include the impact of environmental regulatory change, the use of human, natural and physical resources and corporate governance structures and practices. When the team deems a factor material to the value of a company, the team incorporates it into its decision-making process.

Fund name: Hartford Equity Income Fund; Date: 2022-02-28

> As part of this analysis, Wellington Management evaluates financial and competitive conditions, management quality, potential earnings, free cash flow, dividends, and other related measures or indicators of value, including financially material environmental, social, and/or governance ESG characteristics based on Wellington Managements proprietary ESG research. Wellington Management believes the integration of financially material ESG characteristics into its investment process allows it to better assess strategic business issues that may impact the long-term quality, dividend sustainability, and ultimately the performance of a company. The factors that Wellington Management considers as part of its fundamental analysis, including the assessment of financially material ESG characteristics, contribute to its overall evaluation of a company's risk and return potential.

II.D Mention-only

Fund name: Hartford Dividend & Growth Fund; Date: 2019-03-31

> (...) Wellington management also evaluates a company's business environment, management quality, balance sheet, income statement, anticipated earnings, revenues and dividends, and other related measures or indicators of value, including certain environmental, social and/or governance ESG factors.

Fund name: GMO Quality Fund; Date: 2021-06-30

> In addition, gmo may consider esg environmental, social, and governance criteria as well as trading patterns, such as price movement or volatility of a security or groups of securities.

Fund name: Franklin Mutual US Value Fund; Date: 2022-02-28

> (...) environmental, social and governance ESG related assessments of companies are also considered (...)

III Decomposition of ESG Ratings

Classification prompt. To classify each rating sub-component into ESG-externality and ESG-materiality categories, we query ChatGPT 4.1 through independent API calls, one for each sub-component. The prompt provides the model with the name and description of the sub-component and asks it to assign exactly one of two categories: Category 1 (ESG externality), corresponding to metrics of primary importance for funds that evaluate companies' ESG issues with the objective of generating a positive impact on the world; and Category 2 (ESG materiality), corresponding to metrics of primary importance for funds that evaluate material risks and opportunities associated with companies' ESG issues with the objective of enhancing financial performance. The full prompt is reproduced below.

> *ESG is a set of principles that considers environmental, social, and corporate governance issues. Mutual Funds concerned with ESG issues can be classified into two categories, depending on their investment motives/mandates:*
>
> - *Category 1: Funds that, in their portfolio selection process, evaluate companies' ESG issues with the objective of generating a positive impact on the world.*
> - *Category 2: Funds that, in their portfolio selection process, evaluate material risks and opportunities associated with companies' ESG issues with the objective of enhancing financial performance.*
>
> *Instructions: You will be given a metric and its description, as used by rating agencies to produce ESG-based ratings of companies. After carefully evaluating the metric and its description, select exactly one of the following options:*

- *"Category 1" if the metric is of primary importance for funds in Category 1.*
- *"Category 2" if the metric is of primary importance for funds in Category 2.*
- *"None" only if you cannot reasonably decide between the two categories, or if neither applies.*

You must select only one option. You are not allowed to select both categories. If your answer marks both categories, it will be considered invalid.

Please provide your answer in the following format:

- *Category: [1 or 2 or None]*
- *Explanation: [A brief explanation of why you chose this category, including why you did *not* choose the other category or why you selected "None."]*

Thank you for your time and consideration.

MSCI sub-component categorization. We apply this methodology to the 38 sub-components of MSCI ESG ratings. The following table reports the resulting classification. The model assigns 15 sub-components to the ESG-externality category and 22 to the ESG-materiality category, with one sub-component assigned to "None."

ESG externality	**ESG materiality**
Access to Communications Score	Business Ethics and Fraud Score
Access to Finance Score	Chemical Safety Score
Access to Healthcare Score	Business Ethics Score
Biodiversity & Land Use Score	Corporate Behavior Score
Carbon Emissions Score	Corporate Governance Score
Community Relations Score	Corruption & Instability Score
Controversial Sourcing Score	Energy Efficiency Score
Tax Transparency Score	Financial System Instability Score
Electronic Waste Score	Financing Environmental Impact Score
Financial Product Safety Score	Health & Safety Score
Human Capital Development Score	Climate Change Vulnerability Score
Opportunities in Clean Tech Score	Insuring Health & Demographic Risk Score
Opportunities in Green Building Score	Opportunities in Renewable Energy Score
Opportunities in Nutrition and Health Score	Packaging Material & Waste Score
Raw Material Sourcing Score	Privacy & Data Security Score
	Product Carbon Footprint Score
	Product Safety and Quality Score
	Responsible Investment Score
	Supply Chain Labor Standards Score
	Tax Gap Score
	Toxic Emissions and Waste Score
	Water Stress Score

Extension to other rating providers. We apply the same prompt to classify the sub-components of Refinitiv and Sustainalytics ratings. Since each of these providers has hundreds of sub-components, we do not report the full categorization for brevity. Instead, the following table summarizes the cross-tabulation of the number of sub-components belonging to the Environmental (E), Social (S), and Governance (G) pillars that are assigned to the ESG-externality and ESG-materiality categories, for each of the three providers.

	MSCI			Refinitiv			Sustainalytics		
	E	S	G	E	S	G	E	S	G
ESG externality	6	8	1	97	129	38	64	63	29
ESG materiality	8	7	8	16	24	90	6	16	47

For MSCI, the sub-components are split roughly evenly between the externality and materiality categories. Refinitiv and Sustainalytics, by contrast, have a substantially larger number of externality-related sub-components, reflecting a greater emphasis on metrics that capture firms' impact on society and the environment relative to metrics that capture financially material ESG risks and opportunities.

IV Combined ESG Ratings

In our baseline analysis, we use MSCI ESG ratings because they offer the broadest coverage for our sample period. In this appendix, we verify the robustness of our main results by replacing the MSCI rating with a combined measure that aggregates ESG ratings from MSCI, Refinitiv, and Sustainalytics. For each firm-month, we average across all available provider ratings (after cross-sectional standardization within each provider) to obtain a single combined overall ESG rating and, analogously, combined ESG-externality and ESG-materiality component ratings. This approach mitigates the concern that our findings might be driven by idiosyncrasies of a single rating provider.

Overall ESG ratings. Table IV.1 replicates the analysis of Table 2 using the combined rating. The key patterns are robust. Impact funds continue to display by far the strongest tilt toward high-ESG stocks: their value-weighted and over-weighted portfolio ratings are 45% (0.193/0.430) and 51% (0.203/0.396) higher than those of non-ESG funds, respectively, both significant at the 1% level. As with MSCI ratings, these coefficients are roughly three times larger than those obtained by pooling all ESG funds into a single category. Exclusionary funds show a smaller but significant positive tilt, while opportunity and mention-only funds display modest differences relative to non-ESG funds. The ordering of fund types and the overall pattern of significance are fully consistent with the baseline, confirming that the results are not driven by the choice of rating provider.

ESG-externality and ESG-materiality ratings. Table IV.2 replicates the analysis of Table 4 using the combined rating. The central finding from the baseline analysis is preserved: impact funds are the only type whose portfolios display significantly higher ESG-externality ratings than those of non-ESG funds. Their value-weighted and over-weighted externality ratings are 18% (0.098/0.541) and 22% (0.099/0.459) higher, respectively, both significant at the 1% level. No other fund type shows a significant tilt on the externality dimension in the value-weighted or over-weighted portfolios, confirming that the distinctive behavior of impact funds on the externality dimension is not an artifact of MSCI's particular rating methodology.

On the materiality dimension, impact funds again display the largest and most significant tilt: 22% (0.093/0.423) for value-weighted and 24% (0.093/0.391) for over-weighted portfolios. A notable difference relative to the baseline emerges for opportunity funds: their ESG-materiality tilt, which is significant in the MSCI-only specification (Table 4, columns 5–6), becomes statistically insignificant with the combined rating. This attenuation likely reflects the composition of the combined materiality component. Because Refinitiv and Sustainalytics assign a much larger share of their sub-components to the externality category (as shown in the Internet Appendix III), their materiality components are constructed from a narrower and potentially less overlapping set of metrics. When averaged with the MSCI materiality component, the resulting combined materiality measure may introduce noise that dilutes the relatively modest tilt of opportunity funds, while leaving the larger tilt of impact funds intact. Overall, these results confirm that the portfolio allocation patterns documented in Section 4 reflect genuine differences in fund behavior rather than artifacts of a particular rating methodology.

Table IV.1: Fund-Level ESG Ratings: Combined Rating

This table replicates the analysis in Table 2 using a combined ESG rating that aggregates MSCI, Refinitiv, and Sustainalytics ratings. The dependent variable is one of four measures of fund-level portfolio ESG ratings: value-weighted, over-weighted, under-held, or not-held, as defined in equations (2)–(5). In the top panel, $\mathbf{1}_{jt}^{ESG}$ is a single dummy for ESG funds. In the bottom panel, it is replaced by separate dummies for each ESG category (Exc, Imp, Opp, Men). Non-ESG funds are the excluded category. Controls and fixed effects are as in Table 2. Standard errors are clustered by fund and month; t-statistics are reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Portfolio ESG Ratings (Combined)			
	value-weighted (1)	over-weighted (2)	under-held (3)	not-held (4)
ESG	0.071***	0.078***	0.039*	-0.002
	(6.23)	(6.09)	(1.94)	(-0.27)
Adj. R^2	0.794	0.807	0.391	0.794
Exc	0.057***	0.047**	0.117**	0.030**
	(2.88)	(2.11)	(2.44)	(2.14)
Imp	0.193***	0.203***	0.094**	-0.027
	(7.31)	(6.13)	(2.24)	(-1.14)
Opp	0.036**	0.046***	-0.003	-0.009
	(2.61)	(3.04)	(-0.11)	(-0.88)
Men	0.054**	0.074***	0.005	0.014
	(2.48)	(2.93)	(0.13)	(1.18)
Adj. R^2	0.796	0.809	0.392	0.794
Controls	Y	Y	Y	Y
MStar×Month FE	Y	Y	Y	Y
Obs.	143,637	143,636	92,185	143,637
$\bar{y}_{\text{Non-ESG}}$	0.430	0.396	0.702	0.531

Table IV.2: Fund-Level ESG-Externality and ESG-Materiality Ratings: Combined Rating

This table replicates the analysis in Table 4 using a combined ESG rating that aggregates MSCI, Refinitiv, and Sustainalytics ratings. The dependent variable is one of four measures of fund-level portfolio ratings—value-weighted, over-weighted, under-held, or not-held—constructed using the ESG-externality rating $z_{it}^{ESG^x}$ (columns 1–4) or the ESG-materiality rating $z_{it}^{ESG^m}$ (columns 5–8). In the top panel, $\mathbf{1}_{jt}^{ESG}$ is a single dummy for ESG funds. In the bottom panel, it is replaced by separate dummies for each ESG category (Exc, Imp, Opp, Men). Non-ESG funds are the excluded category. Controls and fixed effects are as in Table 2. Standard errors are clustered by fund and month; t-statistics are reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Portfolio ESG-externality Rating (Combined)				Portfolio ESG-materiality Rating (Combined)			
	value-weighted	over-weighted	under-held	not-held	value-weighted	over-weighted	under-held	not-held
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ESG	0.031***	0.036***	0.033*	0.024**	0.028***	0.029***	0.015	0.004
	(3.29)	(3.30)	(1.87)	(2.29)	(4.01)	(3.56)	(0.97)	(0.68)
Adj. R^2	0.911	0.876	0.683	0.823	0.858	0.845	0.487	0.800
Exc	0.026	0.017	0.032	0.075***	0.027*	0.015	0.054	0.027**
	(1.40)	(0.81)	(0.83)	(3.06)	(1.97)	(0.94)	(1.59)	(2.59)
Imp	0.098***	0.099***	0.053	0.032	0.093***	0.093***	0.052	0.000
	(4.46)	(3.33)	(1.04)	(1.08)	(6.08)	(4.50)	(1.47)	(0.02)
Opp	0.009	0.020	0.022	0.002	0.008	0.012	-0.012	-0.005
	(0.80)	(1.55)	(1.15)	(0.23)	(0.93)	(1.21)	(-0.58)	(-0.72)
Men	0.029	0.041	0.050	0.022	0.011	0.025	0.011	0.015*
	(1.31)	(1.61)	(1.39)	(1.55)	(0.84)	(1.64)	(0.33)	(1.72)
Adj. R^2	0.912	0.876	0.683	0.823	0.858	0.846	0.488	0.800
Controls	Y	Y	Y	Y	Y	Y	Y	Y
MStar×Month FE	Y	Y	Y	Y	Y	Y	Y	Y
Obs.	143,637	143,636	92,194	143,637	143,647	143,646	92,197	143,637
$\bar{y}_{\text{Non-ESG}}$	0.541	0.459	1.074	0.615	0.423	0.391	0.711	0.528

V Alternative Factor Models for Green Window Dressing

In Section 5.2, we report green window dressing results using the FFC6+ESG specification as our baseline. This appendix examines the sensitivity of these results to the choice of factor model. Figure V.1 displays ESG beta levels and delta betas for four specifications of increasing richness: a bivariate model including only the market return and the ESG index (Mkt+ESG), the Fama–French three-factor model augmented with the ESG index (FF3+ESG), the Fama–French–Carhart four-factor model with the ESG index (FFC4+ESG), and the full six-factor model with the ESG index (FFC6+ESG).

The key patterns documented in Section 5.2 are robust across specifications that include style factor controls, but the bivariate Mkt+ESG model produces materially different results. In this specification, ESG beta estimates are negative for several fund types in the control and post-disclosure windows. This is driven by the correlation structure between the ESG index and the omitted style factors, reported in the following table:

	Mkt	SMB	HML	RMW	CMA	MOM
ESG	0.991***	0.144***	-0.035	-0.114***	-0.267***	-0.182***
	(0.000)	(0.025)	(0.026)	(0.025)	(0.024)	(0.025)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ (two-sided Pearson test).
SE in parentheses, delta method: $(1 - \rho^2)/\sqrt{N-2}$.

The ESG index is negatively correlated with the investment factor CMA, the momentum factor MOM, and the profitability factor RMW, all significant at the 1% level. When these factors are omitted from the regression, any fund-level variation in these style exposures across windows is absorbed by the ESG loading, contaminating the window dressing estimates. Combined with the near-perfect correlation between the ESG index and the market return, the inclusion of style factors is necessary to obtain interpretable ESG loading estimates.

Once style factors are included, the results stabilize. Starting from the FF3+ESG model, ESG beta levels are positive for all fund types across all windows, and the characteristic inverted-V profile for opportunity and mention-only funds begins to emerge. The FFC4+ESG model, which adds momentum, and the FFC6+ESG model, which further adds profitability and investment, yield progressively more stable estimates. In the FFC6+ESG specification, the green window dressing pattern for opportunity and mention-only funds is sharp and statistically significant, while the flat profiles for impact and exclusionary funds are cleanly identified.

A potential concern with the FFC6+ESG specification is the limited number of degrees of freedom: with eight parameters estimated over $n = 10$ observations per window, individual fund-level betas are necessarily imprecise. Two considerations mitigate this concern. First, our wild bootstrap procedure does not rely on the precision of individual estimates. It resamples the OLS residuals of each fund-event-window cell and propagates estimation noise directly into the bootstrap distribution of the cross-sectional mean, without requiring asymptotic approximations at the fund level. Second, the quantity of interest is the cross-sectional average across all fund-event pairs within a category, not any individual fund's beta. With approximately 1,556 fund-event pairs per quarter on average, individual imprecision is diversified away, and the precision of the cross-sectional mean improves with the square root of the number of funds.

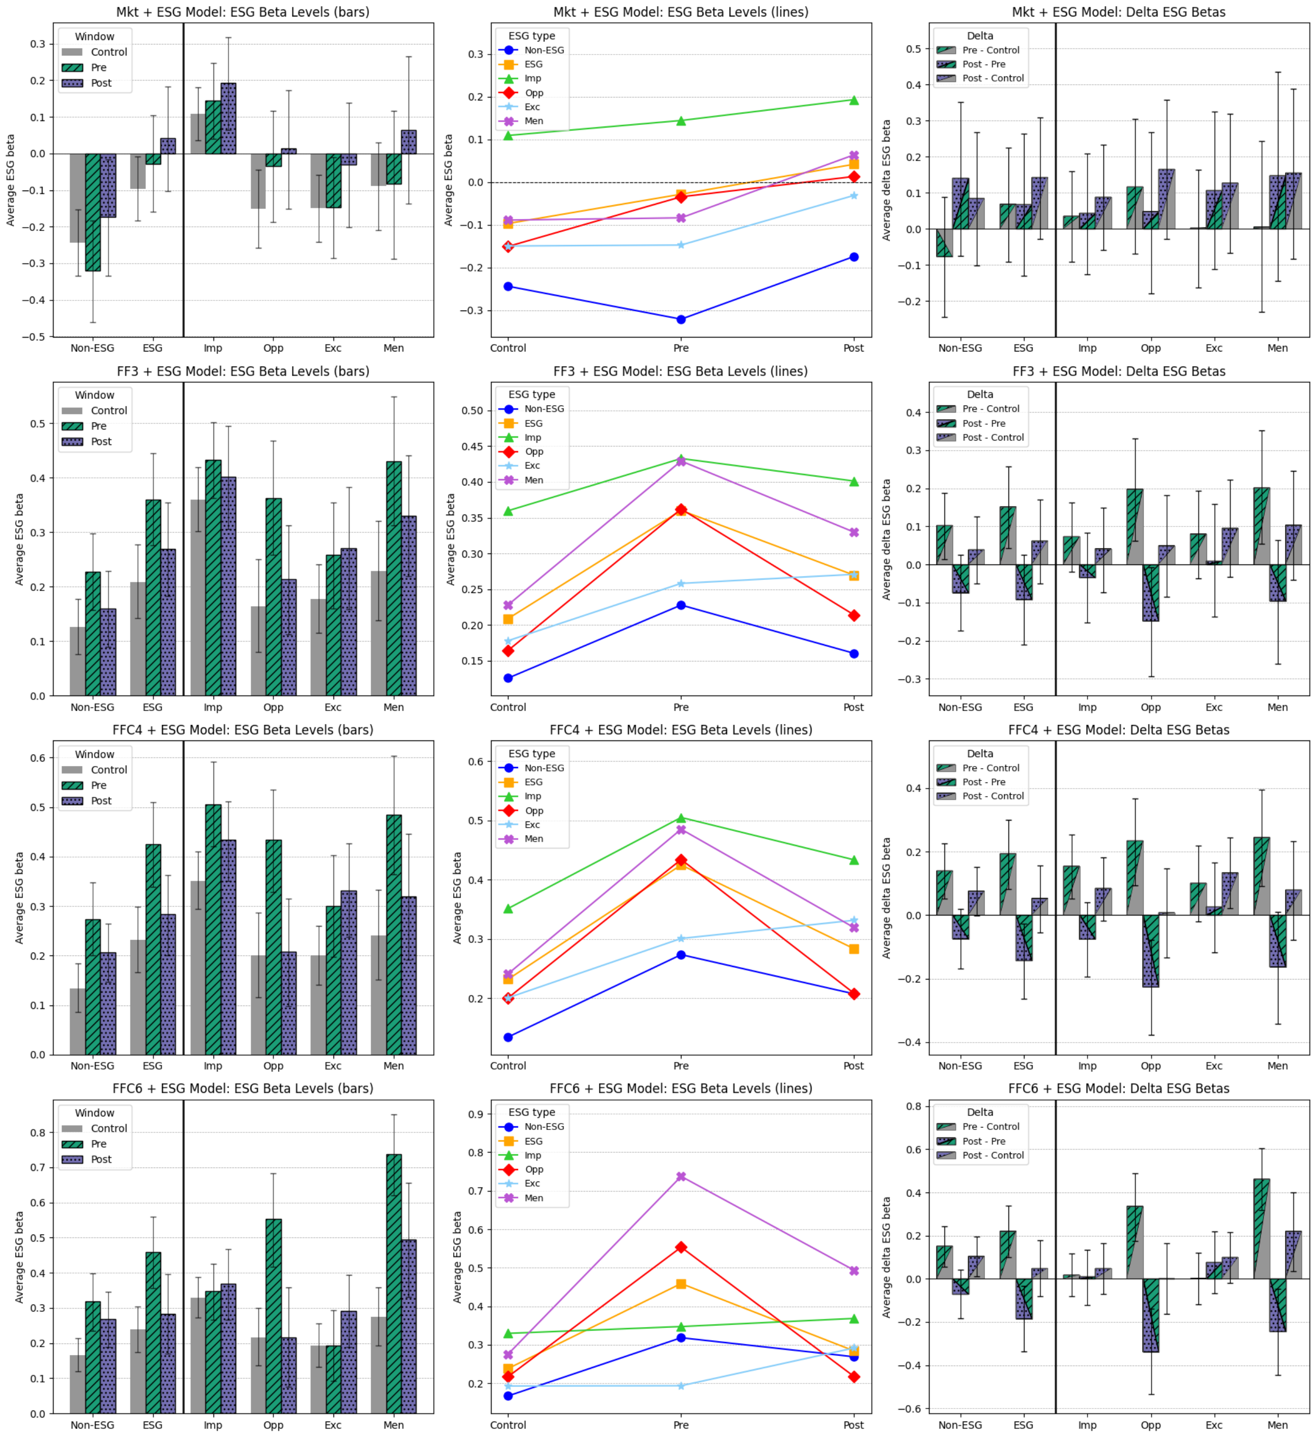

Figure V.1: Green Window Dressing: Alternative Factor Models

This figure displays average ESG factor loadings and their changes around quarterly disclosure dates under four factor model specifications. Each row corresponds to a different specification: Mkt+ESG (top), FF3+ESG (second), FFC4+ESG (third), and FFC6+ESG (bottom). Within each row, the left panel plots ESG beta levels by window (control, pre-disclosure, post-disclosure) for each fund type, the center panel connects the three windows to illustrate the dynamics of ESG exposure, and the right panel plots the delta ESG betas with bootstrap confidence intervals.

VI LLM for ESG Fund Classification

The manual classification described in Section 2.2 produces a richly labeled corpus that can be leveraged to train an automated classifier. In this appendix, we describe a hierarchical deep learning model trained on these labels, designed to extend our classification to newly filed prospectuses at scale.

VI.A Dataset

We extract 12,664 passages from the PIS sections of US mutual fund prospectus filings. Each passage is assigned one of seven mutually exclusive labels: the six ESG categories from Section 2.2 (Exc, Imp, Imp Act, Opp, Opp Act, Men) and a residual non-ESG category (None). Passage lengths vary considerably, ranging from 42 to over 17,000 characters, with a median around 2,100 characters.

The label distribution is heavily skewed. Non-ESG passages account for 88.4% of the corpus, and the smallest ESG class (Impact Activist) contains only 56 observations. During the manual classification, the annotators recorded the specific text spans within each passage that justified their labeling decision. As we describe below, these spans play a central role in both the augmentation strategy and the model's input representation.

VI.B Preprocessing

We transform the raw labeled data into model-ready splits through six steps: semantic clustering, deduplication, class balancing, train-validation-test splitting, data augmentation, and hierarchical label construction.

Semantic clustering. We embed all passages using a pre-trained sentence transformer (all-MiniLM-L6-v2) and apply agglomerative clustering to the resulting 384-dimensional vectors, using average linkage with a cosine distance threshold of 0.7. The clusters capture groups of textually similar prospectuses and serve two purposes: guiding the removal of near-duplicates and ensuring that related documents do not straddle the training and evaluation sets.

Deduplication. We remove near-duplicates in two passes. In the first pass, we compare all pairs within each cluster and discard any passage whose cosine similarity to a retained passage exceeds 0.97. In the second pass, we repeat this threshold globally to catch duplicates that fall in different clusters. This process eliminates 53.8% of the original corpus, leaving 5,849 unique passages.

Class balancing. Even after deduplication, non-ESG passages outnumber ESG ones by a ratio of roughly 7:1. We reduce this gap by randomly undersampling the None class to 90% of the total ESG count, producing a working dataset of 1,445 passages.

Cluster-based splitting. We partition the data into training (60%), validation (20%), and test (20%) sets using a stratified procedure that operates at the cluster level. Within each class, clusters are shuffled and their constituent samples allocated across splits in proportion. Because entire clusters are assigned to the same split, semantically similar passages cannot appear in both the training set and either evaluation set, preventing data leakage.

Data augmentation. To further address class imbalance, we apply two augmentation strategies to the training set only (the evaluation sets remain untouched).

The first strategy swaps annotator-identified relevant text spans across same-class passages. For each ESG training sample, we randomly replace one of its relevant spans with a span drawn from a different passage carrying the same label, generating two augmented variants per sample. An additional round targets the Impact Activist class to compensate for its extreme scarcity. Per-class caps (5× the original count, relaxed to 10× for Impact Activist) prevent any single class from becoming over-represented.

The second strategy applies back-translation. Each original ESG passage is translated from English to French and back using pre-trained models (Helsinki-NLP opus-mt-en-fr and opus-mt-fr-en), producing a paraphrased variant. To preserve the marker tokens described below, the translation is applied only to unmarked portions of the text, and long passages are split into chunks of 350 words before translation. One back-translated variant is generated per original ESG sample, subject to the same per-class caps.

Throughout both strategies, annotator-identified relevant spans are wrapped with class-specific marker tokens (e.g., `[Opp]`...`[/Opp]`). Twelve such tokens are defined, one opening and one closing marker for each ESG category. This tagged text constitutes the model's input. After augmentation, the final training set contains 2,615 passages.

Hierarchical label construction. From the flat seven-class label, we derive four auxiliary labels that correspond to the stages of our hierarchical classifier: (i) a binary indicator separating ESG from non-ESG passages; (ii) a four-way family label grouping ESG passages into Exclusionary, Impact family, Opportunity family, and Mention-only; (iii) a binary indicator distinguishing Impact from Impact Activist; and (iv) a binary indicator distinguishing Opportunity from Opportunity Activist.

VI.C Model Architecture

Our classifier follows a hierarchical design in which a shared document encoder feeds into four specialized classification heads, each responsible for one stage of the labeling hierarchy.

Document encoder. We build on the Longformer (Beltagy et al., 2020), a transformer architecture that combines local sliding-window attention with selective global attention, enabling efficient processing of sequences up to 4,096 tokens. The encoder contains approximately 149 million parameters. We designate the `[CLS]` token for global attention so that its final hidden state (768 dimensions) aggregates information from the entire passage. This vector serves as the shared document representation for all four classification heads. To accommodate the class-specific marker tokens, we expand the tokenizer vocabulary by 12 entries and resize the embedding layer accordingly.

Classification heads. Each head is a two-layer feedforward network with the structure: Dropout(0.1) → Linear → LayerNorm → GELU → Dropout(0.1) → Linear. The four heads differ in their output dimensionality and hidden-layer width:

- *Binary head*: separates ESG from non-ESG passages ($768 \rightarrow 384 \rightarrow 2$).
- *Family head*: assigns ESG passages to one of four families ($768 \rightarrow 384 \rightarrow 4$).
- *Impact Activist head*: distinguishes Impact from Impact Activist ($768 \rightarrow 192 \rightarrow 2$).
- *Opportunity Activist head*: distinguishes Opportunity from Opportunity Activist ($768 \rightarrow 192 \rightarrow 2$).

The activist heads use a narrower hidden layer (192 rather than 384), reflecting the simpler nature of the within-family binary distinction.

Hierarchical inference. At prediction time, the heads are queried sequentially. If the binary head classifies a passage as non-ESG, no further heads are consulted. Otherwise, the family head determines the ESG family. Exclusionary and Mention-only passages receive their final labels directly; Impact family and Opportunity family passages are routed to the corresponding activist head for the final distinction. Seven-class probabilities are obtained by multiplying softmax outputs along the path through the hierarchy—for instance, $P(\text{Imp}) = P(\text{ESG}) \times P(\text{family} = \text{Impact}) \times P(\text{activist} = \text{Imp})$.

VI.D Training

Loss function. We train each head with focal loss (Lin et al., 2017), a variant of cross-entropy that down-weights confidently classified examples and concentrates the gradient signal on difficult cases. For predicted probability p_t on the correct class, the loss is $\text{FL}(p_t) = -\alpha_t(1 - p_t)^\gamma \log(p_t)$. The class weights α_t are set by inverse frequency, with additional multiplicative adjustments for underrepresented classes (Impact family: 1.2×, Mention-only: 1.5×, Impact Activist: 2.5×, Opportunity Activist: 1.3×). The focusing parameter γ is set to 1.5 for all heads except the Impact Activist head, where $\gamma = 2.0$ provides stronger emphasis on hard examples in this extremely sparse class.

The overall objective is a weighted combination of the four head-specific losses:

$$\mathcal{L} = \mathcal{L}_{\text{binary}} + \mathcal{L}_{\text{family}} + 0.5\,\mathcal{L}_{\text{imp\_act}} + 0.5\,\mathcal{L}_{\text{opp\_act}}, \tag{1}$$

where each component is computed only over applicable samples (e.g., the family loss excludes non-ESG passages, and the activist losses are restricted to their respective families).

Optimization. We train for up to 20 epochs with an effective batch size of 16 (8 per device, 2 gradient accumulation steps), a peak learning rate of 1.5×10^{-5} decayed via cosine annealing with a 10% warmup phase, weight decay of 0.01, gradient clipping at norm 1.0, label smoothing of 0.1, and mixed-precision (fp16) computation. All passages are tokenized to a maximum of 2,048 tokens. Early stopping halts training after 5 epochs without improvement in validation macro F1, and the best checkpoint is retained.

VI.E Performance

We evaluate the trained model on both the validation and test sets. All reported metrics are computed on held-out data that contains no augmented samples and no passages from the same semantic clusters as the training set.

Training dynamics. Figure VI.2 tracks the training and validation loss (left panel) alongside the validation macro F1 score (right panel) across epochs. Training loss declines steadily, while validation loss levels off after approximately 10 epochs with no sign of overfitting. The macro F1 score rises sharply in the early epochs and plateaus near its peak, at which point early stopping selects the best checkpoint.

Validation set. Figure VI.3 presents the confusion matrix on the validation set. The classifier correctly identifies the large majority of passages in each class, with most predictions concentrated along the diagonal. Residual errors occur primarily between closely related categories within the same family—Impact versus Impact Activist, and Opportunity versus Opportunity Activist—where the textual signals distinguishing the two are often limited to a few phrases describing the fund's engagement approach. Some Mention-only passages are misclassified as None or Opportunity, consistent with the inherent ambiguity of a category defined by the absence of a clearly articulated ESG strategy.

Test set. Figure VI.4 reports the corresponding confusion matrix on the held-out test set. The error patterns are consistent with those observed on the validation set: strong diagonal performance across all seven classes, with the within-family activist distinction accounting for most of the residual misclassification. This consistency confirms that the model generalizes reliably and that the cluster-based splitting strategy effectively prevents leakage of training signal into the evaluation data.

Summary. The hierarchical Longformer classifier achieves strong performance across all seven classes despite severe class imbalance. Its ability to reliably distinguish among ESG fund types whose prospectus language often differs by only a few words confirms that the distinctions captured by our manual classification are systematic and learnable. The model provides a scalable tool for extending the classification to future prospectus filings as they become available.

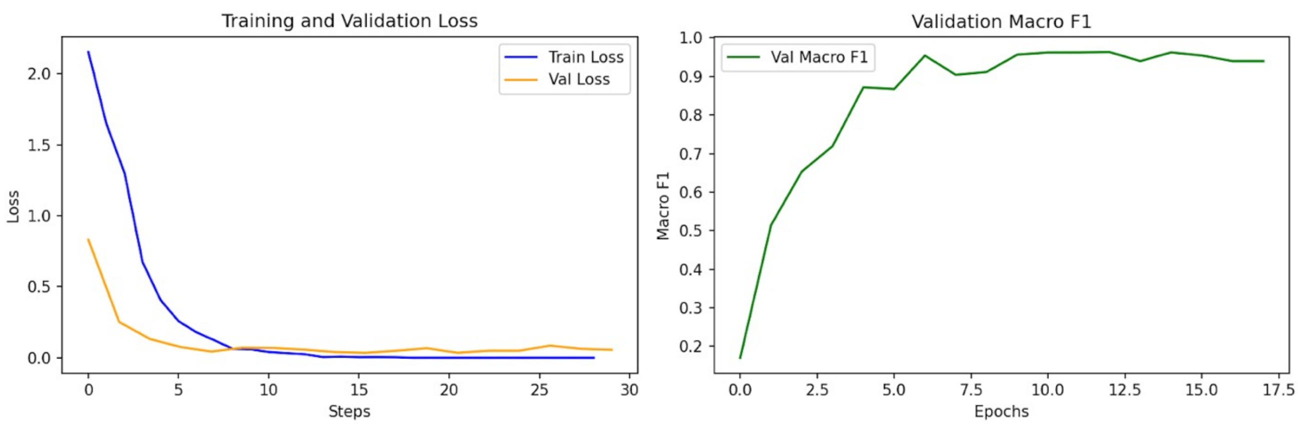

Figure VI.2: Training and Validation Curves

The left panel plots training and validation loss by epoch. The right panel plots the validation macro F1 score. The vertical dashed line marks the epoch at which early stopping selects the best checkpoint.

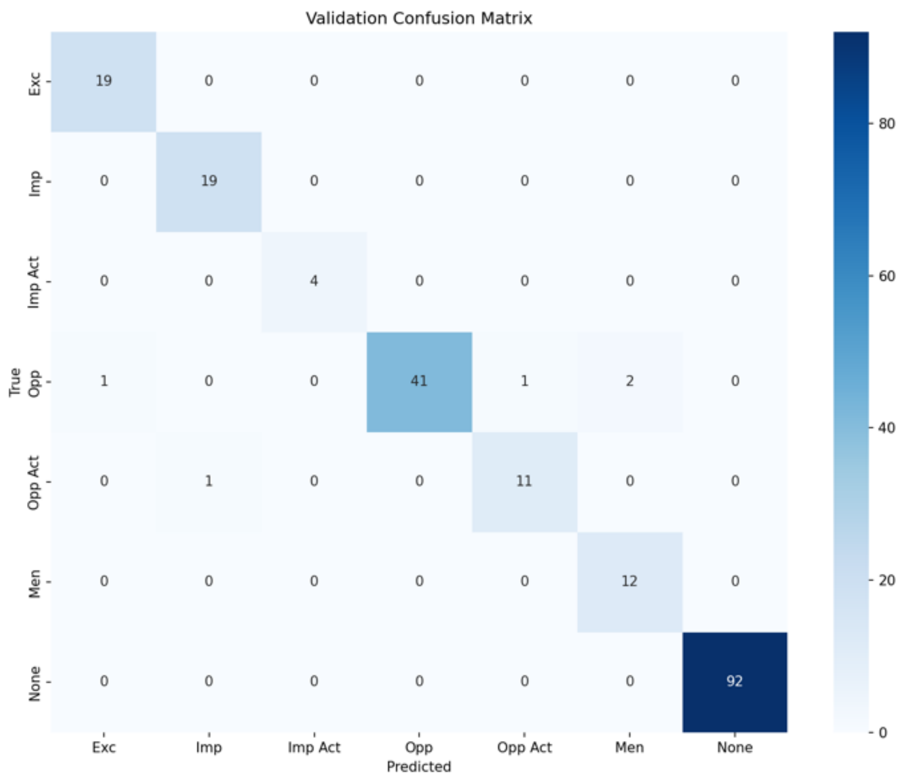

Figure VI.3: Validation Set: Confusion Matrix

Confusion matrix for the seven-class hierarchical classifier on the validation set. Rows correspond to true labels and columns to predicted labels.

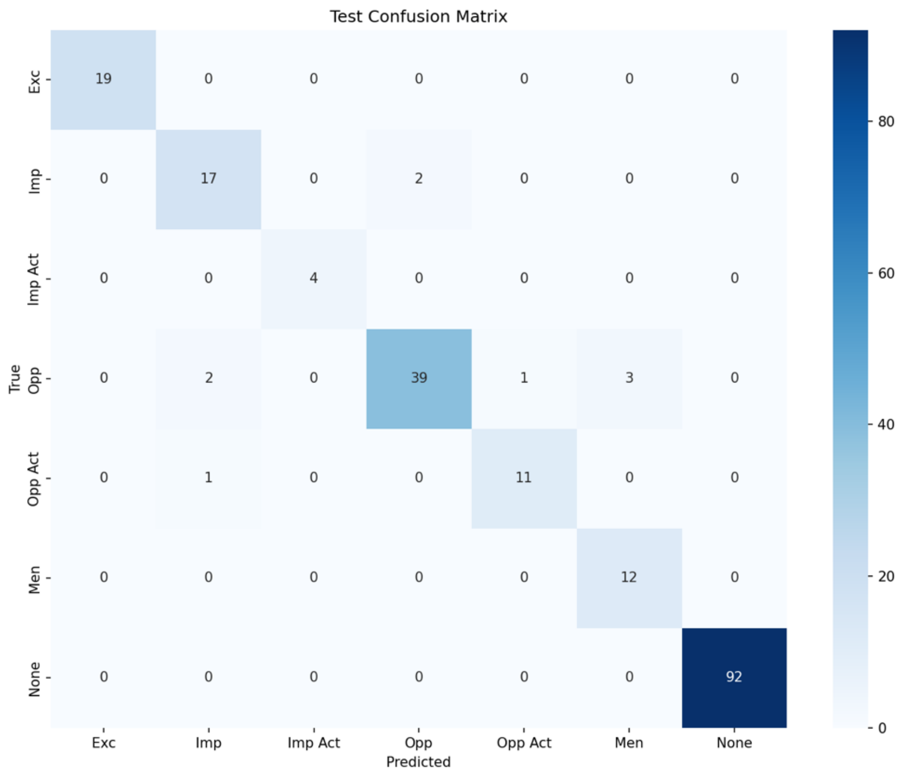

Figure VI.4: Test Set: Confusion Matrix

Confusion matrix for the seven-class hierarchical classifier on the held-out test set. Rows correspond to true labels and columns to predicted labels.

References

Beltagy, Iz, Matthew E. Peters, and Arman Cohan, 2020, Longformer: The Long-Document Transformer, *arXiv preprint arXiv:2004.05150*.

Lin, Tsung-Yi, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár, 2017, Focal Loss for Dense Object Detection, *Proceedings of the IEEE International Conference on Computer Vision*, 2980–2988.