GREEN WINDOW DRESSING IN EUROPEAN ESG FUNDS

ADAM TEJME

VILHELM HÖLJÖ

Bachelor Thesis

Stockholm School of Economics

2026

Green Window Dressing in European ESG Funds

Abstract:

This study examines the prevalence of green window dressing among EU-domiciled SFDR Article 9 equity mutual funds. Green window dressing refers to the practice of temporarily adjusting fund holdings around disclosure dates to present a more sustainable portfolio to investors. Building on the methodology of Parise and Rubin (2025), this study uses return gap analysis and ESG beta comparisons to identify abnormal trading around disclosure dates and whether it serves to overstate ESG exposure. We document a significant return gap, consistent with disclosure-related trading. However, this trading does not reflect a significant change in ESG exposure. Contrary to U.S. evidence, we find no support for green window dressing among EU-domiciled funds. This study provides the first empirical test of green window dressing in the highly regulated EU ESG climate. The results suggest that stricter ESG regulations and higher levels of ESG maturity in the EU may limit the prevalence of green window dressing.

Keywords:

Sustainable finance, ESG, Greenwashing, Window dressing, Green window dressing, Mutual funds

Authors:

Adam Tejme (25888)
Vilhelm Höljö (26136)

Tutor:

Alvin Chen, Assistant Professor, Department of Finance

Examiner:

Ramin Baghai, Professor, Department of Finance

Bachelor Thesis
Bachelor Program in Business and Economics
Stockholm School of Economics

1 Introduction

Over the last decade, ESG and sustainable investing have become central to global financial markets, attracting substantial capital inflows and growing regulatory attention. Academic studies document that a growing share of investors incorporate ESG considerations into their investment decisions (Kräussl, Oladiran & Stefanova, 2024). As investors increasingly demand sustainable investment products, concerns regarding greenwashing among firms and funds have also increased. While greenwashing can take many forms, a particularly subtle and difficult-to-detect form is green window dressing, where fund managers temporarily buy "green" stocks and sell "brown" stocks before mandatory reporting dates to present a more sustainable portfolio than they actually hold (Parise & Rubin, 2025). This, in turn, increases the fund's reported ESG exposure and presents a misleadingly sustainable portfolio to investors. If investors rely on reported holdings to evaluate sustainability, this can distort capital allocation and undermine the credibility of sustainable investment markets.

Although green window dressing is a relatively recent concern, it builds on a much older and well-documented practice in the mutual fund literature. Traditional window dressing has been prevalent for decades, with early studies documenting such activity in the 1990s (Lakonishok, Shleifer, Thaler & Vishny, 1991; Musto, 1999). The practice involves adjusting portfolio holdings around disclosure dates to appear more attractive to investors, for example, by selling stocks prior to reporting if they are perceived as excessively risky or inconsistent with the fund strategy.

Building on these traditional incentives, the ESG setting adds a sustainability-specific motive that gives rise to green window dressing. As funds with stronger perceived ESG characteristics often attract greater capital inflows (Hartzmark & Sussman, 2019), managers have incentives to adjust portfolio holdings to present a more sustainable portfolio to investors. However, funds may benefit from not holding exclusively green portfolios to preserve diversification benefits and retain exposure to potentially higher-yielding brown assets (Parise & Rubin, 2025). This tension between holding a sustainable portfolio and potentially forgoing higher returns is what incentivizes fund managers to engage in green window dressing.

Recent evidence confirms that these incentives translate into observable behavior. Parise and Rubin (2025) provide a landmark empirical study identifying green window dressing among U.S. mutual funds. They use return gap analysis and ESG exposure tests to compare the returns of reported holdings with the returns funds generate between disclosures. These tests show that funds temporarily increase their ESG exposure around mandatory reporting dates. They also find that this behavior results in higher sustainability ratings and capital inflows for funds engaging in the practice. Their findings establish green window dressing as an empirically detectable phenomenon, but leave open whether it is a general feature of ESG funds or specific to the U.S. institutional environment.

Extending current research, our study aims to investigate green window dressing among funds domiciled in the EU, where the institutional context and investor preferences greatly differ from those in the U.S. ESG regulations in the EU, especially regarding ESG

reporting, are considerably stricter (Allcott, Egan, Smeets & Yang, 2026). This can partly be attributed to the recently implemented SFDR (Sustainable Finance Disclosure Regulation) framework. It requires managers of EU-domiciled funds to classify their funds into one of three categories: Article 6 funds, which have no specific sustainability focus; Article 8 funds, which promote environmental or social characteristics; and Article 9 funds, which have sustainability as their primary investment objective (European Commission, 2025). Our analysis focuses on SFDR Article 9 funds, since these funds face the strongest sustainability obligations; they are required to have sustainable investment as their core objective.

Beyond regulatory differences, studies also note a difference in both ESG performance and engagement from investors between the EU and the U.S. EU investors have greater ESG concerns, and firms have higher ESG performance (Gibson, Glossner, Krueger, Matos & Steffen, 2022). This raises the question of whether stricter regulations and a more mature ESG climate reduce greenwashing behavior or if they could instead incentivize funds to engage in green window dressing to appear to comply. Our study, therefore, extends current research on green window dressing to the EU context by addressing the following question:

Do EU-domiciled SFDR Article 9 mutual funds rebalance their portfolios around mandatory disclosure dates, and if so, is this rebalancing directed at overstating ESG exposure?

This study follows the approach of Parise and Rubin (2025) to ensure comparability, while adjusting certain parameters to suit the EU context. We select a primary sample that runs from March 2021 to November 2025, using stock return data from 230 SFDR Article 9 EU-domiciled equity mutual funds. We choose March 2021 as the start date for our main sample to align with the implementation of the SFDR framework and capture the current regulatory environment. For each fund and report date, we construct a counterfactual portfolio using the fund's reported holdings and calculate the returns for this portfolio. It serves as a proxy for how the fund would have performed if it continued to hold its reported portfolio. First, we conduct a return gap analysis, where we compare the returns of the counterfactual portfolio with the reported realized returns of each fund. The aim of this test is to identify if the fund engages in trading around disclosure and whether the realized portfolio differs from the counterfactual portfolio. Second, we conduct an ESG beta comparison. We regress the realized fund returns and the counterfactual returns against an ESG index while controlling for the market. This serves to identify whether the realized and counterfactual portfolios systematically differ in ESG exposure. Thus, we can determine whether funds engage in green window dressing.

The results of our study indicate the presence of traditional window dressing, but we find no evidence of green window dressing among SFDR Article 9 EU-domiciled funds. We identify a statistically significant return gap of -0.036% per day in the ten days leading up to the mandatory disclosure date, consistent with trading fees and other effects related to portfolio rebalancing. This suggests that fund holdings are being altered to a greater extent around disclosure dates, likely to present a more attractive portfolio to investors. In contrast to the U.S.-focused study by Parise and Rubin (2025), we do not find a statistically significant ESG beta difference. In the 15 days following disclosure, we

estimate significant positive ESG betas, indicating that our funds hold greener portfolios than the market as a whole. However, the realized and counterfactual ESG betas are not significantly different, meaning that we do not find evidence of funds systematically altering their ESG exposure when engaging in window dressing. We analyze our results further by conducting a cross-sectional beta difference analysis split by pre-disclosure return gap size. We find no significant ESG beta difference for the funds with the largest pre-disclosure return gap, strengthening the interpretation that this portfolio rebalancing is not tied to changes in ESG exposure. Taken together, we find evidence indicating traditional window dressing, but no evidence of green window dressing. We argue that the lack of evidence of green window dressing in our EU fund sample could be the result of several factors. First, the high ESG maturity of EU firms reduces the incentive for funds to artificially increase ESG exposure around disclosure dates. Second, EU investors appear to place greater emphasis on substantive ESG performance rather than ESG labels alone. Third, the EU's stricter ESG regulatory framework increases the expected costs associated with misleading sustainability claims and greenwashing behavior.

Our study extends current research on green window dressing to the EU context. We contribute to the literature by suggesting that green window dressing is not a universal feature of ESG funds, but appears to depend on the institutional context, including the maturity of the sustainable investment market, investors' ESG preferences, and the strictness of ESG regulation.

2 Literature Review

This paper contributes to the growing literature on sustainable investing and window dressing by investigating the prevalence of green window dressing in EU-based mutual funds. Traditional window dressing is a widely studied phenomenon in the mutual fund literature, with early studies such as Lakonishok et al. (1991) showing that pension fund managers sell poorly performing stocks just before reporting the holdings, to make the portfolio appear more attractive to investors. Similarly, Musto (1999) compares disclosed and undisclosed mutual fund holdings and finds that managers adjust their portfolios shortly before disclosure dates to present less risky portfolios to investors. These studies highlight that disclosure requirements may create incentives for temporary portfolio manipulation.

Beyond these traditional motives, the growth of sustainable investing has created new incentives for portfolio rebalancing around disclosure. ESG and sustainable investing have rapidly become integral to financial markets, increasing the importance of sustainability characteristics in investment decisions. With investors increasingly allocating capital towards sustainable investment products, concerns about greenwashing among funds have increased and numerous studies have documented its prevalence. Several studies show that fund managers have incentives to portray their investments as more sustainable than they actually are. Heeb, Kölbel, Paetzold and Zeisberger (2023) show that investors prioritize investments that are branded as sustainable but show indifference to the actual ESG impact of the funds, thereby creating incentives for fund managers to signal sustainability without altering investment behavior. Similarly, Lowry, Wang and Wei (2026) show that many ESG funds lack proper incentives to create ESG-

positive impact through their investments. Together, these studies suggest that investor demand can encourage superficial ESG signaling, creating conditions where green window dressing can emerge.

Recent studies investigate whether these ESG incentives translate to observable disclosure-related trading. The primary study on which our paper builds is Parise and Rubin (2025), which investigates the prevalence of green window dressing in U.S. mutual funds. Previous research supports strong incentives for funds to engage in greenwashing and, consequently, green window dressing, highlighting the importance of empirically assessing its prevalence. Using several empirical tests, including ESG exposure changes around reporting dates and return gap analysis, Parise and Rubin (2025) find strong evidence that funds temporarily increase their ESG exposure prior to disclosure dates. Additionally, they also examine the incentives for green window dressing and find that green window dressing leads to an increase in fund sustainability ratings and performance. A study by Andrikogiannopoulou, Papakonstantinou, Song and Zu (2025) finds that 15% of U.S. mutual funds actively engage in green window dressing, most commonly in funds with poor past performance.

A gap in current research is that no study has investigated green window dressing among EU-domiciled funds. There have been studies comparing greenwashing between the U.S. and the EU, highlighting differences in ESG regulation as well as behavioral differences (Gibson et al., 2022). These institutional differences make the EU context particularly relevant for further investigation, and it is important to assess whether green window dressing is still prevalent, given the stricter regulatory environment with higher ESG maturity. Our study therefore focuses on the EU and contributes to the literature by providing the first empirical investigation of green window dressing in EU-domiciled funds.

3 Data and Methodology

3.1 Data Sources

Morningstar Direct serves as the primary source of fund-level data, including daily fund returns and reported holdings. Morningstar is regarded as a leading database for fund research and is commonly utilized in financial research regarding mutual funds (Hartzmark & Sussman, 2019). This is complemented with fund expense ratios from LSEG's Lipper Fund Research Database accessed through Refinitiv Workspace. Stock return data is retrieved from the Compustat database through WRDS. The market factor used is taken from Kenneth French's data library and complemented with the Morningstar Europe Sustainability Total Return Index (MSEGEUEG) from Morningstar Direct. Historic daily exchange rate data from LSEG IBES, accessed through WRDS, is used for currency conversions. The data sources used are established research databases, recommended by institutions such as the Stanford Graduate School of Business Library (2025).

Table I
Descriptive statistics of the sample

This table reports descriptive statistics for the fund sample. The sample consists of 230 EU-domiciled, euro-denominated, actively managed equity mutual funds classified as Article 9 under SFDR. The sample period runs from March 2021 to November 2025. Fund size is in millions of USD. Expense ratio is the annual expense ratio reported in percent. Return is the annualized mean daily return of the funds.

Variable	N	Mean	Std. Dev.	P25	Median	P75
Fund Size (USD, mm)	230	665.03	1,207.71	65.37	232.07	687.39
Expense Ratio (%)	230	1.95	0.56	1.61	2.00	2.34
Return, annualized (%)	230	4.36	4.16	1.51	4.41	6.97

Table I presents descriptive statistics for the fund sample. Our sample has an average fund size of USD 665 million, which is comparable to the average TNA reported by Parise and Rubin (2025) of USD 669 million. Our median fund size is substantially lower than the mean, at USD 232 million, reflecting that our sample includes a few large funds that increase the mean. Our mean expense ratio of 1.95% per year is significantly higher than the average expense ratio of 0.92% in Parise and Rubin. This is likely due to our data sample using the highest expense ratio share class for each fund. The mean annual return of our fund sample is 4.36%, which is considered low for equity mutual funds and may be attributed to the turbulent market conditions during the studied period, combined with the high expense ratios of our fund share classes (Ascari, Bonam & Smadu, 2024; McKibbin, Noland & Shuetrim, 2025).

3.2 Sample Selection

Our original fund sample consists of all 247 euro-denominated SFDR Article 9 equity mutual funds in the Morningstar Direct database available during our sample period, with 230 funds remaining after data cleaning. Article 9 funds, defined as funds with a sustainable investment objective, are selected for our sample, as these funds face the strictest sustainability requirements (European Commission, 2025). They are consequently required to maintain high ESG exposure, and fund managers are obligated to report sustainable holdings at the mandatory disclosure dates. Furthermore, Article 9 funds are the most important segment to study in this context, as evidence of green window dressing among these funds would be especially detrimental to the credibility of the SFDR framework. The sample is restricted to equity mutual funds to ensure comparability with Parise and Rubin (2025). Consistent with their methodology, ETFs and index funds are excluded to focus our sample on actively managed funds. The scope is limited to Euro-denominated funds to reduce measurement error related to exchange rate conversions. Maintaining homogeneity in reporting currencies limits possible distortions in the return gap that could arise from exchange rate fluctuations. This is especially important when constructing a counterfactual portfolio for return gap analysis, as differences in fund NAV calculation timing could introduce distortions arising from misaligned exchange rate conversions.

As the funds in our sample have multiple share classes, only one share class per fund is selected, as share classes represent the same underlying portfolio and primarily only differ in fee structures or distribution channels (Morningstar, n.d.). The selected share class is based on the Morningstar Direct primary share class, which is the share class with the highest expense ratio. This is selected as it is the most likely for retail investors to encounter (Morningstar, 2021).

Kenneth French Data Library's European Market Factors are chosen for the market factor, as they are constructed to cover the European market. Our fund sample includes funds with global stock portfolios. However, when analyzing our sample's correlation with market factors shown in Table A.I, it becomes apparent that the sample has a higher correlation with the European factor compared to a global market factor. Thus, the European market factor is more appropriate, as a Global market factor becomes too heavily U.S.-weighted.

In the Parise and Rubin (2025) study, the Morningstar U.S. Sustainability Total Return Index is used to measure funds' green exposure. To ensure comparability with their results and match our European market factor, the Morningstar Europe Sustainability Total Return Index (MSEGEUEG) is used.

All fund and stock returns are retrieved in their domestic currency and converted to USD using the same LSEG IBES exchange rates to ensure consistency and minimize distortions from currency conversion. All analysis is conducted in USD, as this is the base currency for our market factor.

The study looks at data from the period March 2021 until November 2025 to ensure consistent ESG legislation during the period. The sample period begins in March 2021, as this was when the SFDR regulation first came into effect. Prior to the implementation of the SFDR regulation, EU funds were not formally classified into explicit sustainability categories, such as Article 6, 8, and 9. Implementation of the SFDR framework thereby increased incentives for funds to maintain a strong reported sustainability profile in their holdings, especially for Article 9 funds. The sample period ends in November 2025, as this is the last month with full data availability across all data sources.

Figure 1
Cumulative return of the ESG index and market factor during sample period

This figure shows the cumulative return of the ESG index used in our analysis, the Morningstar Europe Sustainability Index, and Kenneth R. French's European market factor including the risk-free rate. The returns are in USD over our sample period from March 2021 to November 2025.

Figure 1 shows the performance of the ESG index we use and the market factor, including the risk-free rate, during the period of our study. The graph shows that the sample period is characterized by turbulent market conditions, with high inflation, wars in Ukraine and the Middle East, and substantial shifts in international trade policy, impacting returns (Ascari et al., 2024; McKibbin et al., 2025). We can also see that the ESG index tracks the market returns closely during the period.

3.3 Data cleaning and adjustments

Due to data limitations, some cleaning and adjustments are needed to ensure accurate analysis of our sample. Funds with missing expense ratios or missing holding data are dropped, as counterfactual portfolio construction for these funds would not be possible. In total, this reduces our sample from 247 funds to 230, which nevertheless remains a substantial sample size and is comparable to Parise and Rubin's (2025) sample of 223 ESG funds.

For these funds, we export the ISINs and CUSIPs of all unique securities they hold. This allows us to use CUSIP to match individual security returns from the North American Compustat database to the securities held by each fund. ISIN is used to match the Global Compustat database returns for securities outside of North America. In total, we are able to match a mean of 87% of fund holdings for our sample. This is in line with Parise and Rubin's (2025) coverage of 89% and should therefore provide sufficient coverage for accurate counterfactual portfolio construction.

As the Compustat database lists all security returns in their local currency, these returns are converted to minimize distortions arising from currency fluctuations. USD is chosen as the base currency, and all currencies are converted to USD, following standard practice in the international mutual fund literature (Cremers, Ferreira, Matos & Starks, 2016).

Some funds in our sample are affected by market holidays in their respective markets, while the underlying holdings are still traded, which results in a substantial return gap for these dates. This is identified by analyzing which funds have a null return while the underlying securities report returns. To address this issue, the affected dates are handled as non-trading days for the counterfactual portfolio as well, shifting the returns to the day after. This means that if markets move significantly on a day when a given fund does not trade, the returns would be shifted and added to the day after, in line with how the fund calculates returns.

Furthermore, some funds in our sample follow a non-calendar fiscal year. This issue is addressed by extracting each fund's end-of-fiscal-year month from the Morningstar Direct database and using the last trading day of this month as the fund's report date. Each fund's mid-year report date is then calculated as the last trading day of the month, 6 months after the end of the fiscal year.

3.4 Method

EU mutual funds are required to report holdings semi-annually under the UCITS framework, once at the end of the financial year and once after half the financial year. This means that EU funds have long periods in between reporting dates, where their holdings are unobservable to outside investors. Thus, if fund managers were to window dress, they would maintain their preferred portfolio during most of the year but rebalance toward a more attractive looking portfolio immediately before the two disclosure dates. They would then revert to their preferred portfolio shortly after reporting. The semi-annual reporting structure also significantly reduces the effective sample size relative to U.S. studies with quarterly disclosures, as it effectively halves the number of observable disclosure events. This also increases reliance on fiscal year-end events, as these make up half our sample, and may differ from interim reports. A version of our return gap analysis excluding year-end reports is provided as additional analysis.

Our method for identifying trading around disclosure and whether green window dressing is prevalent is primarily based on Parise and Rubin's (2025) adaptation of Kacperczyk, Sialm, and Zheng's (2008) definition of the return gap. This is an established method for identifying traditional window dressing in funds, which is based on the premise that even if funds only report holdings twice a year, they report returns daily. Thus, the reported return of a fund can be compared to the counterfactual return of all the holdings reported at disclosure.

Our study consists of two main tests. First, we conduct a return gap analysis where we measure the difference between each fund's realized returns and the returns of their counterfactual portfolio around their disclosure dates. In the absence of any form of window dressing, the two portfolios should exhibit similar returns and thus a return gap of zero. Second, we conduct an ESG beta analysis, where we regress the returns of the realized and counterfactual portfolio against an ESG index to identify if they differ in ESG beta. In the absence of green window dressing, we would expect the beta difference to be zero, as both portfolios should have a similar exposure to ESG stocks.

3.4.1 Return gap analysis

For our main return gap analysis, we analyze a window of ten days before and ten days after each fund's mandatory report date. The report date (t = 0) and the day after (t = 1) are excluded to avoid price fluctuation around the reporting date, in line with the methodology by Parise and Rubin (2025). A 10-day window length is used for the main analysis, as Parise and Rubin (2025) identify this as the optimal window length, thus strengthening our results while maintaining comparability. The 10-day window is long enough to capture statistically significant results, while also being close enough to the report date to isolate disclosure-related rebalancing from funds' regular trading practices. However, results for a longer window are included as a robustness test in the results section. In line with Parise and Rubin's method, only mandatory fund report dates are analyzed, i.e., semi-annually for EU funds, instead of quarterly for U.S. funds. Some of the funds make voluntary portfolio disclosures quarterly or monthly, but these reports are excluded in order to minimize adverse selection bias. Parise and Rubin (2025) argue that voluntary disclosure is likely endogenous to funds' disclosure behavior, which would skew the results.

To construct the counterfactual portfolios from reported holdings, we calculate counterfactual returns using the following formula:

$$R_{i,t}^{H} = \sum_{j=1}^{N_i} w_{j,i,d} \times R_{j,t} + w_{i,d}^{U} \times R_{f,t} \quad (1)$$

The counterfactual returns are calculated daily per fund during each 20-day window around reporting (ten days before to ten days after) for the main return gap analysis, in line with Parise and Rubin's (2025) method, using equation 1. The counterfactual returns are calculated over a 15-day window following reporting for the beta analysis using equation 1. $R_{i,t}^{H}$ is the hypothetical return of each counterfactual portfolio *i* at time *t*. N_i is the number of securities held in the disclosed portfolio. $w_{j,i,d}$ is the weight of each security *j* held by the fund at the report date *d*, $R_{j,t}$ is the daily return of each security, $w_{i,d}^{U}$ is the portion of the portfolio weight that cannot be allocated to a security and $R_{f,t}$ is the risk-free rate. In line with Parise and Rubin, we assume all security IDs that are missing in the data or cannot be mapped to a return to return the risk-free rate. This assumption reflects the possibility that unmatched securities consist primarily of bonds or cash holdings, consistent with evidence from Kacperczyk et al. (2008). This is a conservative choice, as equity returns generally exceed the risk-free rate, thus assigning Rf to unmatched holdings depresses the counterfactual return compared to a fully matched portfolio. As a result, this adjustment mechanically reduces the likelihood of observing a negative return gap.

We calculate the return gap between the realized and counterfactual portfolios in line with Parise and Rubin (2025) as follows:

$$\mathrm{GAP_{i,t}} = \mathrm{R_{i,t}} - \left(\mathrm{R^{H}_{i,t}} - \mathrm{EXP_{i,t}}\right) \quad (2)$$

In equation 2, $\mathrm{GAP_{i,t}}$ is the return gap between the counterfactual and the realized portfolio returns for fund *i* on date *t*. $\mathrm{R_{i,t}}$ is the realized daily net-of-fees returns of the fund, $\mathrm{R^{H}_{i,t}}$ is the counterfactual holdings returns that we calculated in equation 1. $\mathrm{EXP_{i,t}}$ is the daily fund expense that is subtracted from our counterfactual returns to ensure comparability with the realized net fund returns. This is calculated by taking the fund's total annual expense ratio in percent and dividing by 252 trading days to get the daily expenses, in line with Parise and Rubin's (2025) method.

The return gap for all funds is then regressed in two separate regressions, one for the pre-disclosure window, $t \in [t - 10, t - 1]$, and one for the post-disclosure window, $t \in [t + 2, t + 11]$. The mean return gap in each period is estimated using constant-only OLS regression, which is in line with the method from Parise and Rubin (2025). To account for the short estimation windows, bootstrapping is used, which is described further in the robustness section.

3.4.2 ESG beta comparison

Next, we investigate whether the return gap is the result of green window dressing. We regress realized and counterfactual portfolio returns against an ESG factor and a market factor over a 15-day period following reporting. 15 days is chosen over the ten days used in Parise and Rubin (2025) to improve beta estimation precision for our smaller sample while staying close to the report date. Alternative window lengths are also examined in robustness checks. Fund returns are regressed on an ESG factor as follows:

$$R^{k}_{i,t} = \alpha^{k}_{i} + \beta^{MKT,k}_{i} \times MKT_t + \beta^{ESG,k}_{i} \times ESG_t + \varepsilon^{k}_{i,t}, k \in \{R, C\} \quad (3)$$

$$\Delta\beta^{ESG}_{i} = \beta^{ESG,R}_{i} - \beta^{ESG,C}_{i} \quad (4)$$

In Equation 3, k = R denotes the regression on realized net fund returns, and k = C denotes the regression on counterfactual net returns, for fund *i*. α^{k}_{i} is the intercept, $\beta^{MKT,k}_{i}$ is the estimated beta against the market, MKT_t is the daily market return, $\beta^{ESG,k}_{i}$ is the estimated beta against the ESG index, ESG_t is the daily return of the ESG index, and $\varepsilon^{k}_{i,t}$ is the residual error term. The intercept α^{k}_{i} is included to absorb the average return not explained by the two factors over the estimation window but is not central to the empirical analysis; our inference focuses on the ESG loading. Using Equation 3, we run an OLS regression to estimate market beta and ESG beta for both the counterfactual portfolio and the realized portfolio. One ESG beta is estimated per fund over the 15 days following each report period, $t \in [t + 2, t + 16]$. A second constant-only OLS regression is then estimated on all betas to get the mean realized and counterfactual ESG beta of our sample,

in line with the method used by Parise and Rubin (2025). The ESG beta difference $\Delta\beta_i^{ESG}$ is calculated with the resulting realized and counterfactual ESG betas using Equation 4 in a similar constant-only OLS regression. Standard errors are bootstrapped to account for the short beta estimation window, described further in the robustness section.

4 Results

4.1 Main results

Table II
Return gap analysis (10-day window)

This table reports the mean daily return gap in the ten days leading up to disclosure of holdings, $t \in [t - 10, t - 1]$, and the ten days after disclosure excluding the disclosure date and the day after, $t \in [t + 2, t + 11]$. The return gap is calculated as the difference between each fund's realized and counterfactual returns.

	Coefficient	SE	n
Pre-disclosure return gap	−0.0359**	(0.0149)	17,010
Post-disclosure return gap	−0.0104	(0.0158)	16,800

Notes: Coefficients reported in percent (%) per day. Standard errors in parentheses, double-clustered by fund and date with bootstrap (B = 1,000). *** p < 0.01, ** p < 0.05, * p < 0.10.

Table II reports a negative pre-disclosure return gap of -0.0359% per day, significant at the 5% level. This means that in the ten days leading up to disclosure, funds, on average, underperform their counterfactual portfolio corresponding to a cumulative gap of approximately -0.36% over the 10-day window.

The estimated post-disclosure return gap is also negative at -0.0104% per day. This result is statistically insignificant, indicating no statistically significant difference between realized and counterfactual portfolio returns in the ten days following disclosure.

Table III
Return gaps across alternative event windows

This table reports the mean daily return gap in the days before and after disclosure. The return gap is calculated as the difference between each fund's realized and counterfactual returns. Panel A shows the return gap in the days before disclosure and Panel B shows the return gap in the days after disclosure. The return gap is calculated over varying window lengths of 10, 15, 20, and 25 days, while excluding the report date and the day after.

	(1)	(2)	(3)	(4)
Window length (trading days):	10	15	20	25
Panel A: Pre-disclosure return gap				
Coefficient	−0.0359**	−0.0367***	−0.0326***	−0.0286***
	(0.0149)	(0.0123)	(0.0110)	(0.0096)
Observations	17,010	25,515	34,019	41,752
Panel B: Post-disclosure return gap				
Coefficient	−0.0104	0.0043	0.0063	0.0022
	(0.0158)	(0.0139)	(0.0116)	(0.0107)
Observations	16,800	25,200	33,594	41,964

Notes: Coefficients reported in percent (%) per day. Standard errors in parentheses, double-clustered by fund and date with bootstrap (B = 1,000). *** p < 0.01, ** p < 0.05, * p < 0.10.

Table III demonstrates that our results are robust across alternative event windows, ranging from 10 to 25 days. Panel A shows a statistically significant pre-disclosure daily return gap that decreases in magnitude from -0.0359% per day to -0.0286% per day when increasing the window length from 10 to 25 days. The significance level is 5% for the 10-day window and 1% for the longer windows. Panel B shows a post-disclosure daily return gap that switches sign from -0.0104% at the 10-day window to 0.0022% at the 25-day window, with intermediate values of 0.0043% and 0.0063% at the 15- and 20-day windows respectively. However, none of these post-disclosure results are statistically significant.

Table IV
ESG beta analysis (15-day window)

This table reports the estimated mean exposure of the realized and counterfactual portfolios to an ESG factor. This is calculated in the 15 days after disclosure, excluding the report date and the day after, $t \in [t + 2, t + 16]$. Fund returns are regressed against the Morningstar Europe Sustainability Index while controlling for the market factor, using OLS regression. The mean ESG beta difference is calculated as the difference between the realized and counterfactual ESG betas.

	Coefficient	**SE**	**n**
Realized ESG beta	0.4259**	(0.1974)	1,680
Counterfactual ESG beta	0.3747**	(0.1481)	1,680
Beta difference (realized − counterfactual)	0.0512	(0.0931)	1,680

Notes: Standard errors in parentheses, double-clustered by fund and period with bootstrap (B = 1,000). *** p < 0.01, ** p < 0.05, * p < 0.10.

Table IV indicates that the sample exhibits a realized ESG beta of 0.4259 and a counterfactual ESG beta of 0.3747, both statistically significant at the 5% level. However, the estimated beta difference of 0.0512 is statistically insignificant.

4.2 Robustness

Robustness is assessed using bootstrapped double-clustered standard errors and alternative event windows. Bootstrapping is used in line with Parise and Rubin (2025) to adjust for the short windows when calculating the return gap and estimating beta. The bootstrap specification uses 1,000 draws. Standard errors are clustered per fund and per day for the return gap and per fund and window in the beta regression. Double clustering follows Cameron, Gelbach and Miller's (2011) method, to correct for within-fund correlation over time while also accounting for cross-sectional dependence across funds within a given period.

The first robustness test we conduct is a return gap analysis over varying window lengths, 10, 15, 20 and 25 days pre- and post-disclosure, shown in Table III. The negative pre-disclosure return gap remains statistically significant across all event windows. It also shows a decreasing daily gap with a longer window, suggesting that the measured gap is concentrated around disclosure dates.

Furthermore, ESG betas are calculated at varying window lengths as a robustness test shown in Table A.II in the appendix. These results show that our beta difference is insignificant at all window lengths, reinforcing the finding that no statistically significant deviation in ESG exposure can be detected from the reported holdings of the funds in our

sample. The test also shows that ESG betas are always positive, with increasing significance over longer windows. This suggests that the sample funds maintain positive ESG exposure relative to the market over all window lengths. This is consistent with expectations if the ESG factor adequately captures ESG exposure, further strengthening our method.

An additional robustness check we use is the calculation of return gaps while excluding the year-end reports, as shown in Table A.III in the appendix. Note that, due to the smaller sample, the coefficients in this test are not statistically significant but still provide useful insights. Excluding year-end reports reduces the estimated return gap by 25%, indicating that not all of the return gap is driven by window dressing ahead of disclosure. This suggests that part of the estimated effect could be attributed to other year-end effects, which is discussed in 5.2. Nevertheless, the persistence of a negative return gap suggests that the main findings are not solely driven by year-end effects.

5 Discussion

5.1 Overview of findings

Our results do not find evidence of green window dressing among European SFDR Article 9 mutual funds. We find evidence consistent with portfolio rebalancing around disclosure dates, as reflected in the significant negative pre-disclosure return gap. However, the evidence does not suggest that this rebalancing involves systematic substitution from brown to green holdings, as we find no significant ESG beta difference between the counterfactual portfolios the funds report and the portfolios they actually hold following reporting. This contrasts with the evidence from Parise and Rubin's (2025) study regarding U.S. mutual funds, where green window dressing is identified. This suggests that green window dressing is not a universal characteristic of ESG funds but rather could be dependent on factors such as the ESG market maturity, differences in investor sentiments, and the extent of ESG regulations.

5.2 Evidence of portfolio rebalancing around disclosure

The significant return gap we identify indicates abnormal trading around disclosure, consistent with traditional window dressing. However, this does not imply systematic reallocation from brown to green holdings.

The significant pre-disclosure return gap of -0.036% per day is consistent with the results from U.S. mutual funds (Parise & Rubin, 2025) in terms of sign, where they report a significant pre-disclosure return gap of -0.014% per day. A negative return gap is consistent with the fund manager engaging in unobservable actions such as window dressing, as realized fund returns are lower than implied by the disclosed holdings portfolio.

The negative return gap we observe is consistent with two effects related to traditional window dressing. First, window dressing requires costly trading. Funds must rebalance substantial portions of their portfolio over a short horizon, which leads them to accrue

significant trading costs which in turn depress realized returns (Parise & Rubin, 2025). Second, a common form of traditional window dressing includes selling stocks with poor recent performance ahead of reporting (Lakonishok et al., 1991). This mechanism also generates a negative pre-disclosure return gap, as the disclosed portfolio mechanically reflects securities with stronger recent performance than the portfolio actually held during the pre-disclosure period.

Our study identifies a larger negative return gap than the U.S. study. However, this does not necessarily imply that the funds in our sample engage in portfolio rebalancing or window dressing to a greater extent. Previous studies have shown that mutual funds outside the U.S. generally incur larger costs associated with trading (Ferreira, Keswani, Miguel & Ramos, 2013). This can be attributed to the lower liquidity and higher fragmentation of European and other non-U.S. markets, increasing trading costs through wider bid-ask spreads. Thus, the larger pre-disclosure return gap identified in our study is consistent with the documented higher trading fees of EU funds.

Furthermore, as shown in Table A.III in the appendix, our pre-disclosure return gap decrease in magnitude by 25% when removing fiscal year ends from our sample, looking only at intermediate report dates. This shows that part of the underperformance we observe can specifically be attributable to the year-end disclosure event. Gibson, Safieddine, and Titman (2000) find that funds sell holdings that have accumulated negative returns ahead of tax reporting to realize losses and reduce tax liabilities. This form of forced selling would thus also contribute to the underperformance of the realized fund returns. EU funds generally have their tax report date at the end of the fiscal year, compared to U.S. mutual funds that often have their tax year-end on October 31st (Gibson et al., 2000). Consequently, our study captures this effect substantially, while the U.S. study does not capture it as much.

Part of the return gap we observe may also reflect another form of window dressing, namely, window dressing to conceal risk. This is a practice identified by Musto (1999), where funds take on more risk outside of disclosure and then present a less risky portfolio than they actually hold. This mechanism could explain why we do not observe a positive post-disclosure return gap like Parise and Rubin (2025), who analyze a period where higher risk exposure is associated with greater returns. Their study looks at data from 2016 to 2022, a period of predominantly bullish stock market returns that would allow fund managers to outperform their disclosed portfolio by taking on more risk. In contrast, the 2021 to 2025 sample period is characterized by more turbulent market conditions, which would appear less favorable to a fund manager taking on substantial risk. Under such conditions, concealed risk-taking outside disclosure periods may instead lead to weaker realized returns relative to the disclosed portfolio. Table A.IV in the appendix supports this interpretation. It shows that for the period 2017-2020, the funds from our sample had a less negative return gap pre-disclosure and a positive return gap post-disclosure, consistent with the findings reported by Parise and Rubin (2025) for a comparable market period. The same funds, therefore, show opposite return gap patterns in different market conditions, suggesting that hidden risk-taking pays off in bullish markets but backfires in turbulent ones.

Taken together, the return gap evidence indicates that funds in our sample are actively trading around disclosure dates in a way that is consistent with traditional window dressing, such as reallocation towards stocks with strong recent performance and tax harvesting. To assess whether this trading reflects systematic reallocation from brown to green stocks, we consider the ESG beta analysis.

5.3 Absence of green window dressing

The ESG beta coefficients shown in Table IV are both positive and significant, indicating that the model successfully captures ESG exposure while controlling for market movements. This suggests that the sample funds maintain greater ESG exposure than the broader market portfolio, which is consistent with the sustainable investment objectives of Article 9 funds. Due to the high correlation between the ESG index and the market portfolio, the magnitudes of the betas are not directly interpretable (Parise & Rubin, 2025). The relevant measure to consider is the ESG beta difference. The estimated beta difference is small, 0.06, and statistically indistinguishable from zero. Parise and Rubin (2025) identify a significant negative ESG beta difference, which they use to justify the claim that the funds in their sample engage in green window dressing. We cannot compare magnitudes directly with their study, as the magnitude of the beta difference is highly dependent on the underlying indices used for factor construction. However, as our beta difference is small, insignificant, and opposite in sign to what green window dressing would predict, these findings do not identify significant green window dressing in the sample.

Crucially, even funds engaged in the most aggressive disclosure-related trading do not exhibit significant ESG beta differences, supporting the interpretation that the observed rebalancing is not directed toward inflating ESG exposure. As shown in Table A.V in the appendix, the funds that engage in the most rebalancing around disclosure, proxied by their pre-disclosure return gap, do not significantly change their ESG index exposure from their reported portfolio. The test splits the sample into halves, based on the magnitude of their pre-disclosure return gap. Neither subgroup exhibits a statistically significant beta difference, showing that a larger return gap is not linked to greater ESG beta difference for our sample. Taken together, these results suggest that the funds in our sample engage in trading around disclosure dates, but without materially altering their aggregate ESG exposure.

5.4 Why green window dressing may be absent in the EU

The absence of green window dressing in our sample is consistent with three features of the EU sustainable investment environment that distinguish it from the U.S. market studied by Parise and Rubin (2025): higher average ESG commitment among EU firms, stronger investor scrutiny of substantive ESG performance, and a more demanding regulatory framework. These are discussed below.

First, the high baseline ESG commitment of European firms reduces the difference between a fund's unconstrained portfolio and a portfolio optimized for ESG characteristics, thereby reducing incentives for green window dressing. The correlation between ESG and market returns is substantially higher in the EU than in the U.S., as

shown in Table A.VI in the appendix. This suggests that if an EU fund were to alter holdings to appear greener at mandatory disclosure dates, we can expect a less substantial holding alteration necessary compared to the U.S., as EU firms generally exhibit stronger ESG performance (Cicchiello, Marrazza & Perdichizzi, 2023). Moreover, the high correlation between market and ESG returns also contributes to a decreased incentive for fund managers to green window dress, as it is implied that there are smaller return differences between ESG-oriented stocks and the broader market, suggesting limited return-oriented benefits to rebalance away from green holdings. This may also contribute to reducing the detectability of green window dressing, which is discussed in our limitations.

Second, European investors appear to place greater emphasis on substantive ESG performance than on ESG labeling alone. Existing research has shown that European investors are significantly more concerned about ESG factors when selecting investments than U.S. investors are (Amel-Zadeh & Serafeim, 2018). Further research focused on PRI signatories, an ESG certification framework for investment funds, has shown that European and other non-U.S. funds adopting the classification are associated with a significant increase in ESG-related activity (Gibson et al., 2022). For U.S. funds, no such increase in ESG activity is observed, indicating that U.S. funds use the classification primarily as a signaling device, more closely aligned with greenwashing. Thus, European investors and fund managers place greater emphasis on substantive ESG impact which may contribute to lower incentives for green window dressing.

Third, the strict ESG regulation in the EU raises the cost of green window dressing in ways not directly comparable to the U.S. regulatory environment. Over the past decade, several regulatory frameworks have been introduced, including SFDR, the EU Taxonomy Regulation, and the Corporate Sustainability Reporting Directive (European Union, 2022). Article 9 funds face the most significant requirements, as they are obliged to pursue sustainable investment objectives. Together, these regulatory frameworks increase pressure on EU fund managers to avoid misrepresenting ESG exposure, where potential consequences can include fines from National Competent Authorities (NCAs), and for the SFDR Article 9 funds, forced reclassification to Article 8 or Article 6 (European Commission, 2025). This regulatory pressure is further illustrated by the fact that many funds have voluntarily downgraded their SFDR classification from Article 9 to Article 8 in recent years (Scheitza & Busch, 2024), potentially to avoid regulatory scrutiny. Additionally, lawsuits from investors can be filed if they feel misled regarding the ESG content of the fund. In contrast, ESG claims among the U.S. funds studied by Parise and Rubin (2025) are primarily evaluated through third-party ESG ratings such as Morningstar sustainability scores, which impose limited direct consequences for misleading ESG claims. The actual regulations that control ESG branding of U.S. funds are ambiguous, and funds may avoid meaningful penalties for greenwashing due to ambiguity in ESG classification standards (Gibson et al., 2022). Consequently, the regulatory consequences for greenwashing in the EU substantially increase the expected costs of engaging in green window dressing, decreasing incentives for managers to engage in the practice.

5.5 Implications

The findings of this study have important implications for the growing body of research on sustainable investing, policymakers concerned with greenwashing and investors with ESG preferences. The absence of evidence of green window dressing in Europe suggests that the findings of Parise and Rubin (2025) are not universally applicable but instead could depend on factors such as regulation and ESG maturity in the market. The lack of evidence for green window dressing in the EU is also important for regulators, as it suggests that strict ESG legislation may deter systematic green window dressing. Thus, the EU's SFDR, CSRD and Taxonomy Regulation may serve as a benchmark for other regulators seeking to encourage sustainable investment practices while limiting opportunities for deceptive ESG practices. For individual investors, the absence of detected green window dressing in SFDR Article 9 funds may strengthen the credibility and informational value of the SFDR classifications, as our findings are consistent with the view that SFDR Article 9 funds adhere to the investment objectives of their classification.

5.6 Limitations and future research

The method from Parise and Rubin (2025) that we use to detect green window dressing entails limitations that should be considered when interpreting our results. First, the beta estimation is highly sensitive to noise in the data. As our analysis concerns disclosure-related effects, we are limited to a short estimation window after each disclosure. This, combined with the high correlation between the ESG index and the market factor in the regression, amplifies the impact of any noise on the resulting beta estimates. In the European context, the fragmented nature of the market introduces additional noise through imperfect exchange rate conversions. Incomplete holdings coverage for some funds further increases estimation sensitivity. Together, these factors can contribute to the wide standard errors and statistical insignificance of our beta difference.

A second consideration is sample size. EU funds are only required to disclose holdings semi-annually, compared to the quarterly disclosure required of U.S. funds. As a result, our dataset contains roughly half the number of fund-period observations compared to Parise and Rubin's (2025) study, despite covering a similar number of funds over a period of similar length. This reduces the statistical power of our tests and contributes to the larger standard errors we observe. Additionally, our sample consists exclusively of Article 9 funds, which are pre-selected for strong ESG characteristics. While this is appropriate for our research question, it limits the generalizability of our findings to the broader universe of ESG-oriented funds.

Finally, while we find no evidence of a significant beta difference, the absence of statistically significant evidence should not be interpreted as definitive proof that green window dressing does not occur in our sample. Our cross-sectional analysis of the beta gap, combined with the institutional and regulatory features of the EU market discussed above, supports the interpretation that Article 9 funds do not systematically engage in green window dressing. However, our methodology can only identify the absence of detectable effects, not rule out the phenomenon entirely.

Our paper provides a useful foundation for research on green window dressing in Europe, and there are several extensions we believe merit further investigation. First, we suggest widening the scope to include SFDR Article 8 funds. These funds do not face Article 9's strict requirement of a sustainable investment objective and are only required to promote environmental or social characteristics (European Commission, 2025). The less stringent requirements could possibly give Article 8 fund managers stronger incentives to engage in green window dressing, as the regulatory scrutiny is less severe. Analyzing Article 8 funds would also substantially expand the available sample, improving statistical power and reducing the noise-related concerns discussed above. Alternatively, one could replicate the test utilizing only SFDR Article 8 funds and compare the results with our study. Given that Article 9 funds already face very strong ESG obligations, one could argue that Article 9 funds have limited scope for further ESG signaling through disclosure-related trading, thus making a comparison particularly interesting.

Second, future research with access to proprietary holdings data would be able to identify green window dressing with substantially greater precision. Daily holdings data would allow researchers to observe exact portfolio changes around disclosure dates rather than inferring them from return-based methods, removing much of the noise inherent in the beta estimation approach. This would give clearer insight into fund managers' trading behavior and incentives.

Despite these limitations, our methodology follows established approaches in the literature and provides a useful framework for analyzing potential green window dressing behavior in the European market. As data availability improves and ESG disclosure regulation evolves, future research can build on our framework to deepen our understanding of how regulation, market structure, and investor preferences shape ESG-related fund behavior.

6 Conclusion

This study does not find evidence that EU-domiciled SFDR Article 9 mutual funds engage in green window dressing. The analysis identifies a statistically significant pre-disclosure return gap, but a non-significant ESG beta gap, consistent with the presence of traditional window dressing behavior. Our findings, therefore, suggest that while fund managers adjust holdings around mandatory reporting dates, these adjustments are not driven by attempts to temporarily inflate ESG exposure. The study applies the methodology developed by Parise and Rubin (2025), who investigate green window dressing in U.S. funds. Their study observes a significant return gap as well as a significant negative beta difference, finding evidence of green window dressing among U.S. mutual funds. Our findings, therefore, stand in contrast to Parise and Rubin's (2025), where we argue that the absence of evidence of green window dressing in the EU may reflect the EU's strong ESG maturity and strict ESG regulations, such as the SFDR framework. The EU also has an already high ESG exposure of the average fund, reflected in the high ESG index to market correlation, further reducing the incentive for green window dressing.

This paper contributes to the literature by providing the first empirical test of green window dressing prevalence among EU-domiciled funds, as well as investigating the

external validity of the U.S.-based findings in a different regulatory and market environment. Our results suggest that green window dressing is not a universal phenomenon but instead depends on regulatory pressure and the ESG maturity of the market. These findings add to the growing body of research on sustainable finance regulation and greenwashing. While this study is subject to certain limitations, it provides a solid framework for future research on the prevalence of green window dressing in different fund categories or regulatory environments.

References

Allcott, Hunt, Mark L. Egan, Paul Smeets and Hanbin Yang. 2026, The Effects of Regulating Greenwashing: Evidence from Europe's Sustainable Finance Disclosure Regulation, *Harvard Business School*, Working Paper 26-045.

Amel-Zadeh, Amir and George Serafeim. 2018, Why and How do Investors Use ESG Information: Evidence from a Global Survey. *Financial Analysts Journal*, 74(3), 87-103.

Andrikogiannopoulou, Angie, Filippos Papakonstantinou, Hao Song and Chunyu Zu. 2025, Do Mutual Funds Dress their Greens? Working Paper, presented at the Global Research Alliance for Sustainable Finance and Investment.

Ascari, Guido, Dennis Bonam and Andra Smadu. 2024, Global Supply Chain Pressures, Inflation, and Implications for Monetary Policy, *Journal of International Money and Finance*, 142.

Cameron, A. Colin, Jonah B. Gelbach and Douglas L. Miller. 2011, Robust Inference with Multiway Clustering, *Journal of Business & Economic Statistics,* 29(2), 238-249.

Cicchiello, Antonella Francesca, Ferdinando Marrazza and Salvatore Perdichizzi. 2023, Non-financial Disclosure Regulation and Environmental, Social, and Governance (ESG) Performance: The case of EU and US firms, *Corporate Social Responsibility and Environmental Management*, 30(3), 1121-1128.

Cremers, Martijn, Miguel A. Ferreira, Pedro Matos and Laura Starks. 2016, Indexing and Active Fund Management: International Evidence. *Journal of Financial Economics*, 120(3), 539-560.

European Commission. 2025, *Proposal to Amend the Sustainable Finance Disclosure Regulation (SFDR),* COM(2025) 841 final.

European Union. 2022, *Directive (EU) 2022/2464 on corporate sustainability reporting (CSRD)*.

Ferreira, Miguel A., Aneel Keswani, Antonio F. Miguel, and Sofia B. Ramos. 2013, The Determinants of Mutual Fund Performance: A Cross-Country Study, *Review of Finance,* 17(2), 483-525.

Gibson, Rajna, Simon Glossner, Philipp Krueger, Pedro Matos, and Tom Steffen. 2022, Do Responsible Investors Invest Responsibly? *Review of Finance,* 26(6), 1389-1432.

Gibson, Scott, Assem Safieddine and Sheridan Titman. 2000, Tax-Motivated Trading and Price Pressure: An Analysis of Mutual Fund Holdings, *Journal of Financial and Quantitative Analysis*, 35(3), 369-386.

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

Heeb, Florian, Julian F. Kölbel, Falko Paetzold and Stefan Zeisberger. 2023, Do Investors Care About Impact? *Review of Financial Studies*, 36(5), 1737-1787.

Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng. 2008, Unobserved Actions of Mutual Funds, *Review of Financial Studies,* 21(6), 2379-2416.

Kräussl, Roman, Tobi Oladiran, and Denitsa Stefanova. 2024, A Review on ESG Investing: Investors' Expectations, Beliefs and Perceptions, *Journal of Economic Surveys,* 38(2), 476-502.

Lakonishok, Josef, Andrei Shleifer, Richard Thaler, and Robert Vishny. 1991, Window Dressing by Pension Fund Managers, *American Economic Review,* 81(2), 227-231.

Lowry, Michelle, Pingle Wang and Kelsey D. Wei. 2026, Are All ESG Funds Created Equal? Only Some Funds Are Committed, *Review of Financial Studies*, 39(1), 79-113.

McKibbin, Warwick J., Marcus Noland and Geoffrey Shuetrim. 2025, The Global Economic Effects of Trump's 2025 Tariffs, *Peterson Institute for International Economics,* Working Paper 25-13.

Musto, David K. 1999, Investment Decisions Depend on Portfolio Disclosures, *Journal of Finance,* 54(3), 935-952.

Parise, Gianpaolo, and Mirco Rubin. 2025, Green Window Dressing, *Journal of Finance*, 80(6), 3555-3588.

Scheitza, Lisa, and Timo Busch. 2024, SFDR Article 9: Is it All About Impact? *Finance Research Letters*, 62.

Stanford Graduate School of Business Library. 2025, "Business Databases", Stanford University, https://gsb-research-help.stanford.edu/library/faq/285672, Accessed 09 May 2026.

Data sources

French, Kenneth R. 2026, "Kenneth R. French - Data Library", https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html, Accessed 23 April 2026.

LSEG I/B/E/S. 2025, "Institutional Brokers' Estimate System database", accessed via Wharton Research Data Services (WRDS).

Morningstar Direct. 2025, “Morningstar Direct database”, https://morningstardirect.morningstar.com, Accessed 23 April 2026.

Morningstar, Inc. n.d., “Share Class Types”, https://morningstardirect.morningstar.com/clientcomm/Share_Class_Types.pdf, Accessed 01 May 2026.

Morningstar, Inc. 2021, “Morningstar Primary Share in Market Indicator”, https://advisor.morningstar.com/ReleaseNewsLive/releasePopUp.aspx?Id=657&type=Release%20Notes&name=Licensed%20Data, Accessed 01 May 2026.

Refinitiv Workspace. 2025, “Refinitiv Financial Markets Database”.

S&P Global. 2025, “Compustat Fundamentals Annual Database”, accessed via Wharton Research Data Services (WRDS).

Wharton Research Data Services (WRDS). 2026, “Wharton Research Data Services.” WRDS, https://wrds-www.wharton.upenn.edu/, Accessed 22 April 2026.

Appendix A: Additional analysis

Table A.I
Correlation of realized fund returns with different market factors

This table reports the correlation of the realized daily returns of our pooled fund sample with different market factors during the sample period (March 2021 to November 2025). Europe correlation is the correlation with Kenneth R. French's European market factor. Developed correlation is the correlation with Kenneth R. French's Developed market factor.

Factor	Europe corr.	Developed corr.	Diff. (EU − Developed)
Fund returns	0.7184	0.6588	0.0596

Table A.II
ESG betas across alternative event windows

This table reports the estimated mean exposure of the realized and counterfactual portfolios to an ESG factor. This is calculated in the days after disclosure, excluding the report date and the day after, over varying window lengths of 5, 10, 15 and 20 days. Fund returns are regressed against the Morningstar Europe Sustainability Index while controlling for the market factor, using OLS regression. The mean ESG beta difference is calculated as the difference between the realized and counterfactual ESG betas.

	(1)	(2)	(3)	(4)
Window length (trading days):	5	10	15	20
Realized ESG beta	0.4108	0.2197	0.4259**	0.3328**
	(0.7589)	(0.2322)	(0.1974)	(0.1404)
Counterfactual ESG beta	0.6480	0.2985	0.3747**	0.2764***
	(0.5637)	(0.1888)	(0.1481)	(0.1049)
Beta difference (realized − counterf.)	−0.2372	−0.0788	0.0512	0.0564
	(0.2306)	(0.0999)	(0.0931)	(0.0871)
Observations	1,680	1,680	1,680	1,680

Notes: Standard errors in parentheses, double-clustered by fund and period with bootstrap (B = 1,000). *** p < 0.01, ** p < 0.05, * p < 0.10.

Table A.III
Return gap analysis, mid-year disclosures only (10-day window)

This table replicates the results of Table II, while excluding year-end disclosures. It reports the mean daily return gap in the ten days leading up to disclosure of holdings, $t \in [t - 10, t - 1]$, and the ten days after disclosure, excluding the disclosure date and the day after, $t \in [t + 2, t + 11]$. Year-end disclosures have been excluded, and the return gap is only calculated for mid-year disclosures. The return gap is calculated as the difference between each fund's realized and counterfactual returns.

	Coefficient	**SE**	**n**
Pre-disclosure return gap	−0.0272	(0.0195)	9,190
Post-disclosure return gap	−0.0090	(0.0180)	9,140

Notes: Coefficients reported in percent (%) per day. Standard errors in parentheses, double-clustered by fund and date with bootstrap (B = 1,000). *** p < 0.01, ** p < 0.05, * p < 0.10.

Table A.IV
Return gap analysis, 2017–2020 (10-day window)

This table replicates the results of Table II, but over the time period January 2017 to December 2020. It reports the mean daily return gap in the ten days leading up to disclosure of holdings, $t \in [t - 10, t - 1]$, and the ten days after disclosure, excluding the disclosure date and the day after, $t \in [t + 2, t + 11]$. The fund sample consists of the funds from the main sample that are also operating during the 2017-2020 period. The return gap is calculated as the difference between each fund's realized and counterfactual returns.

	Coefficient	**SE**	**n**
Pre-disclosure return gap	−0.0188	(0.0213)	11,071
Post-disclosure return gap	0.0208	(0.0146)	9,670

Notes: Coefficients reported in percent (%) per day. Standard errors in parentheses, double-clustered by fund and date with bootstrap (B = 1,000). *** p < 0.01, ** p < 0.05, * p < 0.10.

Table A.V
ESG betas by pre-disclosure return gap half (15-day window)

This table replicates the results of Table II but split over two subsamples based on the magnitude of the pre-disclosure return gap per fund. It shows the estimated mean exposure of the realized and counterfactual portfolios to an ESG factor. This is calculated in the 15 days after disclosure, excluding the report date and the day after, $t \in [t + 2, t + 16]$. Fund returns are regressed against the Morningstar Europe Sustainability Index while controlling for the market factor, using OLS regression. The mean ESG beta difference is calculated as the difference between the realized and counterfactual ESG betas. H1 shows the half of the fund sample with the largest (most negative) pre-disclosure return gap and H2 shows the half of the fund sample with the smallest pre-disclosure return gap.

Half	Mean pre-gap (%)	Specification	Coefficient	SE	n
H1 (largest gap)	−0.0617	Realized ESG beta	0.3717*	(0.2135)	851
		Counterf. ESG beta	0.3404*	(0.1787)	851
		Difference (R − C)	0.0313	(0.1037)	851
H2 (smallest gap)	−0.0117	Realized ESG beta	0.4816**	(0.1885)	829
		Counterf. ESG beta	0.4098***	(0.1278)	829
		Difference (R − C)	0.0717	(0.0943)	829

Notes: Mean pre-gap reported in percent (%) per day. Standard errors in parentheses, double-clustered by fund and period with bootstrap (B = 1,000). *** p < 0.01, ** p < 0.05, * p < 0.10.

Table A.VI
Correlation of regional ESG indices with regional market factor

This table reports the correlation of the regional ESG indices with their respective market factor during the sample period (March 2021 to November 2025). US correlation is the correlation between the Morningstar U.S. Sustainability Index and Kenneth R. French's U.S. market factor. EU correlation is the Morningstar Europe Sustainability Index and Kenneth R. French's European market factor.

Factor	US correlation	EU correlation	Difference (US − EU)
Mkt	0.9460	0.9816	−0.0356

Appendix B: AI and academic integrity disclosure

Part of the code used in this thesis is adapted from the replication materials provided in Parise and Rubin (2025). AI tools have been used to assist in coding and to refine language. Claude was used to assist in the coding process, including debugging, commenting and cleaning the script for the replication package. ChatGPT was used to check grammar in the text, assisting in ensuring coherent sentence flow. Please note that the design of the research, analysis of the data, and interpretation of the results are the products of the authors' own independent work.