SCIENTIFIC CULTURE
www.sci-cult.com

SCIENTIFIC CULTURE, Vol. 12, No. 4, (2026), pp. 11232-11245
Open Access. Online & Print

DOI: 10.5281/zenodo.124261073

ESG QUALITY, PRICE VOLATILITY, AND MARKET GOVERNANCE: DISENTANGLING THE INFORMATION DISTORTION AND RISK-REDUCTION CHANNELS IN UK EQUITY MARKETS

[1]Dr. Yiu Fai Chan, [2]Lawrence M. Ngoe, [3]Dr. William Tannor, [4]Shailendra Baghel, [5]Dr. Charles Nzeh and [6]Dr. Giamene Odom

[1]*University of Salford, United Kingdom*
[1]*y.f.chan@salford.ac.uk*
[1]*ORCID: 0009-0001-8436-3460*
[2]*Business and Management, Greater Manchester Business School, University of Greater Manchester, United Kingdom*
[2]*L.Ngoe@greatermanchester.ac.uk*
[2]*ORCID: 0000-0002-8655-1908*
[3]*Global Banking School (GBS), Manchester Campus, United Kingdom*
[3]*wtannor@globalbanking.ac.uk*
[3]*ORCID: 0000-0003-2458-7224*
[4]*De Montfort University, Dubai, UAE*
4shailendra.baghel@dmu.ac.uk
[4]*ORCID: 0009-0003-4692-7733*
[5]*University of Greater Manchester, Bolton, United Kingdom*
[5]*C.Nzeh@greatermanchester.ac.uk*
[5]*ORCID: 0009-0009-8131-0265*
[6]*University of Lancashire, Preston, United Kingdom*
[6]*godom@lancashire.ac.uk*
[6]*ORCID: 0000-0002-2943-9821*

Received: 29/12/2025
Accepted: 17/04/2026

Corresponding Author: Dr. Yiu Fai Chan
(Email: y.f.chan@salford.ac.uk)

ABSTRACT

This paper shows that environmental and social (ES) quality affects stock price volatility through two opposing channels whose dominance depends systematically on externality salience. The market-governance channel of Chen, Gupta, and Starmans (2026, Journal of Financial Economics) predicts that ESG investor trading distortions reduce price informativeness, generating a positive ES quality–volatility association for negative externality firms under high externality salience. The competing risk-reduction channel predicts a negative association through operational risk mitigation. A formal comparative statics analysis of the dual-channel model shows that $\partial Volatility/\partial ES\_Quality = \beta_1 + \beta_3 \times Salience$, where $\beta_1 < 0$ reflects risk-reduction dominance at baseline and $\beta_3 > 0$ captures salience amplification of the CGS mechanism. We test this prediction using a panel of 140 firm-year observations across 20 FTSE-listed firms (2018–2024). A COVID-19 difference-in-differences design provides evidence consistent with salience-driven between-group effects: the

cross-group volatility differential expands by 4.8 percentage points during the pandemic period ($p < 0.05$), with partial post-COVID persistence. The mechanism test—an ES quality × salience interaction for negative externality firms—yields a positive coefficient ($\beta_3 = +0.0019$, $p < 0.10$) consistent with CGS channel strengthening under elevated salience, representing a 61% attenuation of the baseline risk-reduction slope (−0.0031 to −0.0012). This interaction result is robust across three alternative salience proxies (COVID binary, Google Trends ESG search intensity, sector ESG ETF trading volume), each capturing a distinct dimension of investor attention. We characterise our findings as patterns consistent with salience-amplified trading distortions under maintained assumptions, rather than causal identification, and specify the research design—annual panel ESG data, active ESG fund ownership, and the UK TCFD quasi-experiment—required for causal channel separation.

KEYWORDS: Market governance, ESG, price volatility, information distortion, risk reduction, externality salience, FTSE, difference-in-differences, interaction test

JEL Classification: G12, G14, G30, G34, M14, Q56

1. INTRODUCTION

A firm's environmental and social (ES) quality affects the information content of its stock price through two opposing mechanisms with fundamentally different governance implications. Identifying which mechanism dominates—and under what conditions—is both theoretically significant and empirically tractable. This paper develops and tests the conditions for channel dominance.

The information distortion channel, formalised in Chen, Gupta, and Starmans (2026)—hereafter CGS—operates through informed investor trading behaviour. Socially concerned investors deviate from trading on financial fundamentals when they observe significant externalities: an investor employing negative screening may abstain from buying despite observing strong financial performance. This deviation reduces the effort informativeness of stock prices—the precision of the signal about managerial effort that shareholders use to design incentive contracts. The asset pricing consequence (CGS Proposition 5) is that better ES quality, by reducing the frequency of externality realisations that trigger trading distortions, produces more informed trading on financial fundamentals and therefore higher price volatility. The risk-reduction channel operates through fundamentals rather than trading behaviour: better ES quality reduces firm-specific operational risk, lowering cash flow variance and directly reducing price volatility, independent of any informed trading mechanism. The two channels make opposite within-group predictions about the ES quality–volatility relationship for negative externality firms.

We contribute a formal comparative statics result that unifies the two channels under externality salience as the moderating variable. Salience governs the intensity with which informed investors condition their trading on externality information. When salience is high, the CGS trading distortion mechanism is more pronounced, pushing the ES quality–volatility relationship toward positive territory. When salience is low, the distortion is negligible and risk-reduction dominates, producing the negative association. Our comparative statics show that $\partial\text{Volatility}/\partial\text{ES\_Quality} = \beta_1 + \beta_3 \times \text{Salience}$, where $\beta_1 < 0$ (risk-reduction baseline effect) and $\beta_3 > 0$ (salience amplification of the CGS channel). This mapping directly links the theoretical framework to an empirically estimable regression equation.

We test these predictions using a panel of 140 firm-year observations across 20 FTSE-listed firms over 2018–2024, partitioned into negative externality firms (energy, mining, tobacco) and positive externality firms (healthcare, education, data analytics). Our empirical strategy has three components. First, a difference-in-differences design exploiting COVID-19 as a high-salience shock for negative externality firms. Second, an ES quality × salience interaction test within the negative externality group—the key mechanism test that discriminates between the two channels. Third, validation across three alternative salience proxies capturing distinct dimensions of investor attention: a COVID binary indicator, a Google Trends ESG search intensity index, and a sector ESG exchange-traded fund (ETF) trading volume measure. Convergence across these three proxies substantially strengthens the credibility of the salience mechanism.

Our findings are as follows. First, negative externality firms exhibit higher unconditional price volatility (0.2776 versus 0.2592, $p < 0.05$), consistent with their lower mean ES quality (70.16 versus 85.49). Second, the COVID-19 DiD yields a positive and significant differential effect on negative externality firm volatility ($\beta_3 = +0.0482$, $p < 0.05$), robust to controls and subsample exclusions. Third—the key mechanism result—the ES quality × salience interaction is positive and marginally significant across all three salience proxies (β_3 ranging from +0.0017 to +0.0022, $p < 0.10$ in all specifications), consistent with patterns that would obtain under salience-amplified trading distortions as maintained in the CGS framework. Economically, the interaction implies that the baseline risk-reduction slope of −0.0031 attenuates by 61% to −0.0012 during high-salience periods. Relative to the within-group standard deviation of annual volatility (0.1029), this attenuation corresponds to a 1.8 percentage point shift in annual price volatility per unit of ES quality improvement—an economically meaningful magnitude given that mean annual volatility for the group is 27.8%.

We characterise these findings with appropriate discipline. The interaction results are consistent with patterns that would obtain under salience-amplified trading distortions, under maintained assumptions about the exogeneity of our salience proxies to idiosyncratic firm volatility. They do not constitute causal identification of the CGS mechanism. The research design in Section 7 specifies the data requirements and quasi-experimental variation needed for causal identification.

Our contribution to the literature is threefold. First, we provide the first dual-channel comparative statics framework for the ESG quality–volatility relationship, formally deriving the salience-interaction prediction from first principles. Second, we test this prediction using three complementary salience proxies across 140 FTSE firm-year observations, finding convergent evidence consistent with the framework. Third, we provide a precise specification of the causal identification agenda, positioning this paper as Stage 1 of a two-stage research programme.

2. RELATED LITERATURE AND CHANNEL DEVELOPMENT

2.1 The CGS Information Distortion Channel

Holmström and Tirole (1993) establish the market governance channel of stock prices. In their framework, informed traders incorporate private information about firm value into prices through their order flow, generating a likelihood ratio (φ^*) of observed order flow under managerial effort versus shirking. Shareholders exploit this likelihood ratio to design efficient compensation contracts: higher φ^* (more informative prices) implies lower incentive costs. Kyle (1985) provides the microstructural foundation; Edmans (2009) extends this to blockholder trading; Maug (1998) analyses the liquidity-control trade-off.

CGS introduce socially concerned informed investors into this framework. An investor with social concern intensity $\gamma \geq 0$ values share ownership at $x(F + \gamma E)$, where F is the financial payoff and E the externality. When γ is sufficiently large, the investor deviates from trading on financial fundamentals: she may abstain from buying despite $F = 1$ when $E = \eta < 0$ (negative screening), reducing informed trading intensity τ^* and thereby φ^*. The central result for our purposes is CGS Proposition 5: price volatility increases with p_E (the probability of a low, zero externality—our ES quality proxy) for negative externality firms, because higher p_E means the informed investor's trading is more frequently undistorted and more informative about financial fundamentals.

Empirical evidence consistent with the CGS information distortion mechanism is provided by Goldstein, Kopytov, Shen, and Xiang (2022), who show that ESG investor heterogeneity reduces price informativeness, and Hitzemann et al. (2024), who document that sustainable investing dampens stock price sensitivity to earnings announcements. Cao, Titman, Zhan, and Zhang (2023) show that ESG institutional trading generates return patterns consistent with informed trading distortions.

2.2 The Risk-Reduction Channel

The risk-reduction channel operates through cash flow fundamentals. Better ES quality reduces the probability of ESG-related adverse shocks—regulatory penalties, reputational crises, stranded asset write-downs, supply chain disruptions—that generate large negative cash flow realisations. Lins, Servaes, and Tamayo (2017) show that high-CSR firms exhibited smaller stock price declines during the 2008–2009 financial crisis, attributing this to social capital that reduces exposure to trustdestroying shocks. Albuquerque, Koskinen, and Zhang (2019) derive that ESG differentiation reduces systematic risk through product market channels, predicting lower betas and lower volatility for high-ESG firms. Derrien, Krueger, Landier, and Yao (2025) document that negative ES events reduce firm value through cash flow channels, implying that ES quality improvements stabilise fundamentals. The risk-reduction channel therefore predicts ∂Volatility/∂ES_Quality < 0 for both firm types—the opposite of the CGS prediction for negative externality firms.

2.3 Externality Salience as the Moderating Variable

We introduce externality salience as the key moderating variable governing channel dominance. Salience captures the degree to which a firm's externality is salient to informed investors' trading decisions at a given point in time. When salience is high, the CGS trading distortion is more intense because investors are more likely to incorporate externality information into trading decisions, reducing the weight placed on financial fundamentals. When salience is low, the distortion is negligible, and risk-reduction dominates.

Salience is distinct from, but correlated with, several observables. It is not identical to media coverage (which may be endogenous to volatility), not identical to regulatory pressure (which operates through fundamentals rather than trading distortions), and not identical to ESG fund ownership (which determines the prevalence of ESG-motivated investors rather than the salience of their information). We measure salience through three proxies that capture different dimensions: the COVID-19 period as a discrete high-salience shock, Google Trends ESG search intensity as a measure of retail and institutional investor attention, and sector ESG ETF trading volume as a measure of market-level ESG-motivated trading activity. Convergence across these three proxies, despite their different measurement approaches, provides triangulating evidence on the salience mechanism.

2.4 The UK Institutional Context as a High-ESG Equilibrium

We explicitly position the UK as a high-ESG institutional equilibrium: a market in which both channels operate with sufficient intensity to be detectable, but in which the institutional features predict specific patterns of channel dominance. Four features of the UK market are relevant.

First, ESG investor prevalence is high: Edmans, Gosling, and Jenter (2024) report that 77% of UK active equity managers incorporate ES performance into stock selection. This high prevalence means the CGS baseline distortion (γ_0) is not negligible, even if risk-reduction dominates at baseline. Second, categorical sectoral exclusion is common: many UK ESG funds exclude the energy and mining sectors categorically, independent of within-sector ES quality variation. This reduces within-group variation in CGS distortion intensity, weakening the within-group mechanism relative to the between-group mechanism. Third, mandatory TCFD reporting (FTSE 100 from April 2022) improves externality information precision for treated firms, which CGS Section 6.4 predicts affects the ES quality–volatility relationship non-monotonically. Fourth, within-group ES quality variation for FTSE firms is compressed relative to global samples—large-cap firms have invested substantially in ESG management, narrowing within-group dispersion. These features collectively predict: risk-reduction dominates within groups at baseline; CGS patterns are detectable between groups and under high salience; and the UK is a conservative test environment (understating the CGS mechanism relative to global or US samples).

3. THEORETICAL FRAMEWORK: COMPARATIVE STATICS OF THE DUAL-CHANNEL MODEL

3.1 The Salience-Moderated Dual-Channel Model

We develop the dual-channel model formally. Let $Volatility_{it}$ denote the annual price volatility of firm i in period t. We decompose this into CGS and risk-reduction components, moderated by salience:

$Volatility_{it} = V_!(ES\_Quality_i, Salience_t) + V_R(ES\_Quality_i) + \varepsilon_{it}$ (1)

where $V_!$ is the CGS component (information distortion effect on volatility) and V_R is the risk-reduction component. We specify linear approximations:

$V_! = (\gamma_0 + \gamma_1 Salience_t) \times ES\_Quality_i$ (2)

$V_R = \delta \times ES\_Quality_i, \quad \delta < 0$ (3)

where $\gamma_0 \geq 0$ is the baseline CGS distortion effect (proportional to the unconditional prevalence of ESG-motivated trading and the baseline externality salience), $\gamma_1 > 0$ captures the salience amplification of the CGS component, and $\delta < 0$ is the risk-reduction channel slope. Substituting into Equation 1 and taking the partial derivative with respect to ES quality:

$\partial Volatility/\partial ES\_Quality = (\gamma_0 + \gamma_1 Salience_t) + \delta$ (4)

Equation 4 is the paper's key comparative statics result. It maps directly to the estimable regression equation:

$\text{Volatility}_{it} = \alpha + \beta_1\ \text{ES\_Quality}_i + \beta_2\ \text{Salience}_t + \beta_3\ (\text{ES\_Quality}_i \times \text{Salience}_t) + \gamma X_{it} + \varepsilon_{it}$ (5)

where the estimated coefficient $\beta_1 = \gamma_0 + \delta$ is the net ES quality–volatility slope at zero salience. If risk-reduction dominates at baseline (δ numerically larger than γ_0), then $\beta_1 < 0$. The coefficient $\beta_3 = \gamma_1$ is the salience amplification of the CGS component: the degree to which the ES quality–volatility slope becomes more positive as salience increases. A positive and significant β_3 is therefore the empirical signature of salience-amplified trading distortions consistent with the CGS mechanism, under the maintained assumption that salience is approximately exogenous conditional on controls.

The marginal effect of ES quality at any salience level is: ME(ES_Quality | Salience) = $\beta_1 + \beta_3 \times$ Salience. At low salience (Salience = 0), ME = β_1, which is predicted to be negative (risk-reduction dominates). As salience increases, ME shifts toward zero and potentially becomes positive (CGS dominates). The salience level at which the marginal effect changes sign is: Salience* = $-\beta_1/\beta_3$, representing the 'tipping point' at which the CGS channel begins to outweigh risk-reduction.

3.2 Testable Predictions

P1 (Cross-group gap): *Under maintained assumptions that ES quality differences across groups reflect differences in externality probability* p_E*, negative externality firms (lower ES quality) exhibit higher mean price volatility than positive externality firms.*

P2 (Baseline within-group slope): *At baseline salience,* $\beta_1 < 0$ *for both firm type groups, consistent with risk-reduction channel dominance when externality salience is low.*

P3 (DiD salience shock): *A high-salience shock (COVID-19) increases the cross-group volatility differential, consistent with salience-driven amplification of between-group mechanisms.*

P4 (Mechanism test—key): *For negative externality firms,* $\beta_3 > 0$ *in Equation 5: the ES quality–volatility slope becomes less negative (more positive) under high salience, consistent with salience-amplified trading distortions of the CGS type under maintained exogeneity assumptions.*

P4a (Proxy convergence): *P4 holds across multiple salience proxies capturing distinct dimensions of investor attention. Convergence across proxies reduces the concern that the interaction result reflects proxy-specific measurement.*

We note the maintained assumption required for causal interpretation of β_3: that the salience proxy is uncorrelated with idiosyncratic firm-level volatility shocks conditional on controls. This assumption is most plausible for aggregate salience measures (Google Trends ESG index; ETF trading volume) that are common across firms in a sector rather than firm-specific. It is less plausible for firm-level media coverage measures, which is why we avoid firm-specific media coverage as a primary proxy.

4. DATA, SAMPLE, AND METHODOLOGY

4.1 Sample Construction

Our sample comprises 20 FTSE-listed firms over 2018–2024, classified by primary externality type. Negative externality firms (N = 10): BP plc, Shell plc (oil and gas; carbon emissions), Rio Tinto Group, Anglo American plc, Glencore plc (mining; land use and pollution), Imperial Brands plc, British American Tobacco plc (tobacco; public health externalities), SSE plc (utilities with fossil fuel exposure), Diageo plc (alcohol), and Unilever plc (supply chain environmental impact). Positive externality firms (N = 10): AstraZeneca plc, GSK plc (pharmaceuticals; public health R&D), Halma plc (safety technology), Pearson plc, RELX plc (education and information externalities), Experian plc (financial inclusion), SEGRO plc (sustainable logistics), InterContinental Hotels Group plc, Whitbread plc (employment training), and Mondi plc (sustainable packaging). The balanced panel contains 140 firm-year observations across seven years, with three sub-periods: pre-COVID (2018–2019; N = 40), COVID (2020–2021; N = 40), and post-COVID (2022–2024; N = 60).

4.2 Price Volatility

Annual price volatility (Volatility_{it}) is the annualised standard deviation of daily log returns: $\sigma_{it} = \text{SD}(\log(P_t/P_{t-1})) \times \sqrt{252}$. Daily adjusted closing prices are retrieved via the Yahoo Finance API (yfinance v0.2+, auto_adjust=True). We require at least 200 trading day observations per firm-year. Idiosyncratic volatility is the residual standard deviation from a daily market model regression against the FTSE All-Share index, removing common market movements. We use idiosyncratic volatility in robustness checks to confirm that results are not driven by market-wide volatility correlated with ESG profiles.

4.3 ES Quality

ES quality (ES_Quality_i) is 100 minus the Sustainalytics ESG Risk Score (Morningstar, April 2024). This score is time-invariant in our dataset—the primary limitation of our design. As a consequence, within-group regression estimates identify cross-firm differences in permanent ESG profiles against cross-firm differences in mean annual volatility, rather than within-firm dynamics as in a full panel with annual ESG scores. We note this limitation explicitly in every table footnote. Future work using annual Refinitiv ESG data would enable within-firm identification.

4.4 Salience Proxies: Three-Proxy Triangulation Strategy

Our identification strategy for the salience interaction rests on three proxies capturing distinct dimensions of investor attention. Convergence across proxies is the key credibility argument: if the interaction is an artefact of measurement noise in any single proxy, it should not survive across three conceptually distinct measures.

Proxy S1—COVID Binary (H_Salience): The binary indicator $H\_Salience_t$ = 1 for 2020–2021 and 0 otherwise captures the COVID-19 period as a discrete high-salience shock. During this period, oil demand collapse, energy transition acceleration, and intensified public health scrutiny dramatically increased investor attention to the externalities of energy, mining, and tobacco firms. This proxy is the most exogenous of the three, as the COVID shock was not caused by firm-level ESG quality variation. Its limitation is measurement precision: it assigns a single indicator to a two-year period of varying salience intensity.

Proxy S2—Google Trends ESG Search Intensity (GT_Salience): Following Da, Engelberg, and Gao (2011), who show that Google search volume measures investor attention, we construct an annual index of Google Trends search intensity for the query terms 'ESG investing UK', 'sustainable investing UK', and 'green investing UK', averaged and standardised to mean zero and unit variance. This measure captures the intensity of retail and institutional investor interest in ESG, which directly proxies the attention allocated to externality information in trading decisions. Unlike the COVID binary, this proxy varies continuously across all seven years and captures ESG-specific attention rather than general market stress.

Proxy S3—Sector ESG ETF Trading Volume (ETF_Salience): We construct an annual measure of aggregate trading volume in the three largest UK-domiciled ESG equity ETFs (iShares MSCI UK ESG Enhanced, SPDR S&P UK Dividend Aristocrats ESG, and Invesco MSCI Europe ESG Universal), standardised to mean zero and unit variance. High ETF trading volume indicates high ESG-motivated trading activity in UK equity markets, directly proxying the prevalence of the market-level behaviour that the CGS mechanism requires. This measure captures ESG market activity rather than attention (proxied by GT_Salience), providing a distinct operational dimension of salience.

The three proxies are positively correlated (mean pairwise correlation = 0.61) but distinct: their correlations with market-wide volatility (VIX) differ substantially (H_Salience: r = 0.71; GT_Salience: r = 0.24; ETF_Salience: r = 0.18), confirming that GT_Salience and ETF_Salience are less confounded by general market stress than the COVID binary. If the salience interaction result is robust across all three proxies, its interpretation is less susceptible to the alternative explanation that the interaction merely captures general market turmoil correlated with sector-level ESG profiles.

Table 1: Salience Proxies—Definitions, Sources, and Correlations

Proxy	Name	Definition	Source	r(VIX)	r(S2)	Dimension captured
S1	H_Salience (COVID binary)	=1 for 2020–2021; 0 otherwise	Constructed	.71	—	Crisis shock; most exogenous
S2	GT_Salience (Google Trends)	Standardised annual Google Trends ESG search index (UK)	Google Trends API	.24	1.00	Investor attention; ESG-specific
S3	ETF_Salience (ETF volume)	Standardised annual aggregate ESG ETF trading volume (UK)	Yahoo Finance API	.18	.58	Market-level ESG trading activity

Note: *r(VIX) = Pearson correlation of annual proxy with FTSE 100 implied volatility index. r(S2) = correlation with GT_Salience. Lower VIX correlation indicates proxy is less confounded by general market stress.*

4.5 Estimation

We estimate three specifications corresponding to the three predictions. The DiD specification (P3) is:

$$Volatility_{it} = \alpha + \beta_1\, Neg_i + \beta_2\, S1_t + \beta_3\, (Neg_i \times S1_t) + \gamma X_{it} + \varepsilon_{it} \quad (6)$$

The within-group mechanism test (P4), estimated separately for each externality group, is Equation 5 from the theoretical section. We estimate this three times, once for each salience proxy (S1, S2, S3), and report all specifications in Table 4. For P2 (baseline slope), we estimate Equation 5 with Salience set to zero (or equivalently, pool all years and test β_1 alone). All specifications use firm-level clustered standard errors. Pooled OLS without firm fixed effects is the primary specification for within-group slope estimates, since ES_Quality is time-invariant. Year fixed effects are included throughout.

5. EMPIRICAL RESULTS

5.1 Descriptive Statistics

Table 2 presents descriptive statistics. The cross-group volatility differential (+0.0184, $p < 0.05$) consistent with P1 is documented in Panel B. Panel C reveals the temporal structure: the differential is insignificant pre-COVID (+0.012, $p > 0.10$), large during COVID (+0.081, $p < 0.01$), and partially persistent post-COVID (+0.016, $p < 0.10$). All three salience proxies peak in the COVID period (Panel D), confirming that the COVID period represents the highest-salience episode in our sample.

Panel of 20 FTSE-listed firms, 140 firm-year observations, 2018–2024

Table 2: Descriptive Statistics

Variable	N	Mean	SD	Min	Max	Neg−Pos Diff	Sig.
Panel A: Full Sample (N=140)							
Annual Volatility	140	0.2684	0.0974	0.1590	0.6340		
ES Quality (100−Sustainalytics)	140	77.83	10.21	58.50	88.60		
Idiosyncratic Volatility	140	0.2322	0.0934	0.1498	0.5924		
Panel B: By Externality Type (Neg/Pos)							
Annual Volatility	70/70	0.2776 / 0.2592	0.1029 / 0.0914	0.1817 / 0.1590	0.5762 / 0.6340	+0.0184	**
ES Quality Score	70/70	70.16 / 85.49	9.24 / 1.98	58.50 / 83.20	83.60 / 88.60	−15.33	***
Env. Score (raw)	70/70	25.33 / 7.90	10.87 / 1.53	11.80 / 5.80	42.10 / 10.40	+17.43	***
Idiosyn. Volatility	70/70	0.2431 / 0.2213	0.0981 / 0.0887	0.1612 / 0.1487	0.5418 / 0.5931	+0.0217	**
Panel C: Cross-Group Volatility Gap by Period							
Pre-COVID (2018–19)	40	0.221	0.069	0.163	0.378	+0.012	n.s.
COVID (2020–21)	40	0.394	0.136	0.219	0.634	+0.081	***
Post-COVID (2022–24)	60	0.233	0.054	0.159	0.349	+0.016	*
Panel D: Salience Proxy Values by Period							
S1: H_Salience (COVID binary)		0.29	0.45	0	1	Peak: 2020–21	
S2: GT_Salience (standardised)		0.00	1.00	−1.42	+1.84	Peak: 2021	
S3: ETF_Salience (standardised)		0.00	1.00	−1.31	+1.62	Peak: 2020–21	

Note: **** p<0.01, ** p<0.05, * p<0.10, n.s. not significant. Panel B rows show values for negative / positive groups respectively. ESG Risk Scores from Sustainalytics (April 2024); time-invariant. Idiosyncratic volatility is the residual standard deviation from a daily CAPM regression against the FTSE All-Share.*

5.2 Difference-in-Differences: COVID as High-Salience Shock (P3)

Table 3 presents the DiD results estimating Equation 6. The coefficient of interest is β_3—the differential effect of the COVID period on negative versus positive externality firm volatility.

Dependent variable: Annual price volatility. N = 140 firm-year observations.

Table 3: DiD Results — COVID Shock and Cross-Group Volatility Gap

	(1) Baseline	(2) +Controls	(3) Excl. Oil	(4) Excl. Health	(5) Idiocyn. Vol.	(6) Post-COVID only
Neg (group indicator)	+0.0142	+0.0139	+0.0121	+0.0156	+0.0118	+0.0081

	(0.0091)	*(0.0088)*	*(0.0094)*	*(0.0097)*	*(0.0076)*	*(0.0062)*
H_Salience (COVID period)	+0.1482***	+0.1391***	+0.1203***	+0.1618***	+0.1241***	n.a.
	(0.0248)	*(0.0231)*	*(0.0261)*	*(0.0227)*	*(0.0198)*	*n.a.*
Neg × H_Salience (β_3)	+0.0482**	+0.0461**	+0.0284*	+0.0531**	+0.0394**	+0.0161*
	(0.0214)	*(0.0208)*	*(0.0151)*	*(0.0234)*	*(0.0181)*	*(0.0093)*
Controls	No	Yes	No	No	No	No
N	*140*	*140*	*126*	*126*	*140*	*60*
R^2	0.448	0.491	0.412	0.461	0.413	0.194

Note: **** p<0.01, ** p<0.05, * p<0.10. Firm-level clustered standard errors in parentheses. Controls include leverage, ROA, log market capitalisation, and market beta. Model 3 excludes BP and Shell. Model 4 excludes AstraZeneca and GSK. Model 5 uses idiosyncratic volatility. Model 6 restricts to 2022–2024 to test post-COVID persistence; β_3 is the Neg group indicator (Salience = 0 in this period). Pre-trend test (unreported): 2018 and 2019 interaction terms are +0.0021 and −0.0018 (both p > 0.50), consistent with parallel pre-trends.*

The DiD estimate (β_3 = +0.0482, $p < 0.05$) is consistent with P3. This effect is robust across all five robustness specifications. Model 3 (excluding oil majors) yields a reduced but still positive and significant estimate (+0.0284, $p < 0.10$), suggesting the differential is not entirely driven by the oil price collapse in the energy sector. Model 5 (idiosyncratic volatility) yields +0.0394 ($p < 0.05$), ruling out the alternative that the result is driven by market-wide volatility movements correlated with sector classification. Model 6, restricting to the post-COVID period (2022–2024) and testing whether the Neg group differential persists, yields +0.0161 ($p < 0.10$): roughly one-third of the COVID differential persists post-period, consistent with partial structural change in ESG investor behaviour rather than a purely transitory shock.

We note the appropriate interpretive constraint: both the CGS and risk-reduction channels predict a positive β_3 in the DiD. The DiD is consistent with P3 but does not discriminate between channels. The interaction test in Section 5.3 provides the discriminating evidence.

5.3 The Core Mechanism Test: ES Quality × Salience Interaction with Three Proxies (P4)

Table 4 presents the key mechanism test. The interaction coefficient β_3 in Equation 5 is the primary object of interest. We report six specifications: the baseline within-group slope without interaction (Models 1–2), and the salience interaction using each of the three proxies for negative externality firms (Models 3–5), and positive externality firms (Model 6) using the primary proxy.

Dependent variable: Annual price volatility. Estimated within each externality group.

Primary result: β_3(ES_Quality × Salience) for negative externality firms across three salience proxies.

Table 4: Core Mechanism Test — ES Quality × Salience Interaction

	(1) Neg. Baseline	**(2) Pos. Baseline**	**(3) Neg.×S1 (COVID)**	**(4) Neg.×S2 (Google)**	**(5) Neg.×S3 (ETF Vol.)**	**(6) Pos.×S1 (COVID)**
ES_Quality (β_1)	−0.0031**	−0.0018*	−0.0031**	−0.0029**	−0.0030**	−0.0021*
	(0.0013)	*(0.0008)*	*(0.0014)*	*(0.0013)*	*(0.0013)*	*(0.0011)*
Salience proxy (β_2)			−0.0412	−0.0211	−0.0198	+0.0218
			(0.0384)	*(0.0291)*	*(0.0274)*	*(0.0271)*
ES × Salience (β_3) ★			+0.0019*	+0.0017*	+0.0022*	−0.0004
			(0.0011)	*(0.0009)*	*(0.0013)*	*(0.0008)*
Net slope (High Sal.)	—	—	−0.0012	−0.0012	−0.0008	−0.0025
Salience tipping point	–	–	*ES≈153*	*ES≈171*	*ES≈136*	*n.a.*
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	*70*	*70*	*70*	*70*	*70*	*70*
R^2	0.383	0.160	0.421	0.409	0.416	0.174

Note: **** p<0.01, ** p<0.05, * p<0.10. Firm-level clustered standard errors. ★ = key mechanism test coefficient. S1 = COVID binary; S2 = standardised Google Trends ESG search index; S3 = standardised sector ESG ETF trading volume. Net slope = $\beta_1 + \beta_3$ at S=1 (binary) or S=+1 SD (continuous). Salience tipping point = $-\beta_1/\beta_3$; values > 100 (max ES Quality in sample = 88.6) indicate the CGS channel does not fully dominate* *within the sample range. All models use pooled OLS; ES_Quality is time-invariant.*

Baseline results (Models 1–2) confirm P2: the within-group ES quality–volatility slope is negative for both firm types (−0.0031 and −0.0018, both statistically significant), consistent with risk-reduction channel dominance at baseline.

The mechanism test (Models 3–5) provides the paper's key finding. Across all three salience proxies, the interaction coefficient β_3 is positive and marginally significant: +0.0019 ($p < 0.10$) for S1 (COVID binary), +0.0017 ($p < 0.10$) for S2 (Google Trends), and +0.0022 ($p < 0.10$) for S3 (ETF trading volume). This convergence across three proxies capturing conceptually distinct dimensions of investor attention—crisis shock, attention/search behaviour, and market-level ESG trading activity—substantially strengthens the credibility of the interaction result. It is unlikely that a single measurement artefact would produce a positive interaction coefficient across all three measures, each of which has a different correlation structure with market-wide volatility (Table 1).

The economic magnitude of the interaction is meaningful. At baseline salience (S = 0), the marginal effect of ES quality on volatility is β_1 = −0.0031. During the COVID high-salience period (S1 = 1), the marginal effect shifts to $\beta_1 + \beta_3 = -0.0031 + 0.0019 = -0.0012$. This represents a 61% attenuation of the baseline risk-reduction slope. Translating to volatility units: a one standard deviation improvement in ES quality (9.24 points) is associated with 0.029 lower annual volatility at baseline, compared with 0.011 lower annual volatility during high-salience periods—a difference of 1.8 percentage points in annual price volatility. Given that mean annual volatility in the negative externality group is 27.8% and the within-group standard deviation is 10.3%, this 1.8 percentage point shift in the ES quality–volatility tradeoff represents approximately 17% of the within-group volatility standard deviation—an economically non-trivial magnitude. The salience tipping point—the ES quality level at which the CGS channel fully offsets risk-reduction—exceeds the maximum ES quality in our sample (88.6) for all three proxies, confirming that the CGS channel does not fully dominate within our sample range even at peak salience. This is consistent with our characterisation of the UK market as a context in which both channels operate, with neither fully dominating.

For positive externality firms (Model 6), the interaction coefficient is negative and insignificant (−0.0004, $p > 0.10$). This non-result is interpretable within the framework: the risk-reduction channel predicts negative slopes for positive externality firms regardless of salience, and the CGS channel's salience prediction for positive externality firms is ambiguous (positive screening creates different dynamics than negative screening). The absence of a significant positive interaction for the positive externality group is therefore consistent with the framework rather than contradictory to it.

5.4 Alternative Specifications and Robustness

We conduct four additional robustness checks, reported in Appendix B. First, replacing ES Quality with tercile dummies (Low/Medium/High) reveals a monotonic negative relationship at baseline with attenuation during high-salience periods, confirming the interaction result is not artefactual to the cardinal scaling of Sustainalytics scores. Second, excluding borderline-classified firms (Diageo and Unilever) strengthens the interaction coefficient marginally (+0.0022 across proxies, $p < 0.10$), confirming that classification ambiguity does not drive the result. Third, using idiosyncratic volatility as the dependent variable preserves the sign and approximate magnitude of the interaction, confirming the result is not driven by market-wide volatility correlated with ESG profiles. Fourth, estimating the interaction using only the pre- and post-COVID years (excluding 2020–2021) with GT_Salience and ETF_Salience as continuous proxies—thereby testing whether the interaction exists outside the COVID window—yields β_3 = +0.0014 ($p < 0.15$), suggestive of a weaker but directionally consistent interaction in lower-salience years.

6. DISCUSSION: CHANNEL DOMINANCE CONDITIONS AND THE UK CONTEXT

6.1 Synthesising the Evidence

Table 5 synthesises the evidence against each prediction using calibrated language that distinguishes consistent patterns from causal identification.

Table 5: Evidence Synthesis

	Prediction	Result	Assessment	Interpretive constraint
P1	Cross-group: Neg > Pos volatility	Diff = +0.018 ($p<0.05$)	Consistent	Not channel-specific; both channels predict this direction given ES quality differential.
P2	Baseline within-group slope < 0	Neg: −0.003**; Pos: −0.002*	Consistent	Consistent with risk-reduction dominance at baseline under maintained assumption of

				negligible CGS distortion at low salience.
P3	DiD COVID shock: $\beta_3 > 0$	$\beta_3 = +0.048$ ($p<0.05$)	Consistent	Not fully channel-discriminating. Robust to exclusion of oil majors and healthcare firms.
P4	ES×Salience interaction: $\beta_3 > 0$ (Neg. group)	S1: +0.0019*; S2: +0.0017*; S3: +0.0022*	Suggestive	Consistent with salience-amplified trading distortions under CGS mechanism. Marginal significance ($p<0.10$). Convergent across 3 proxies reduces measurement artefact concern.
P4a	Proxy convergence across 3 measures	All three positive and significant	Consistent	Key credibility result. Each proxy has different VIX correlation, reducing alternative explanation from general market stress.

Note: *'Consistent' = pattern in predicted direction but not causally identified. 'Suggestive' = consistent with mechanism under maintained exogeneity assumptions; confidence limited by sample size and marginal statistical significance.*

6.2 Four Conditions for Channel Dominance

Our results indicate four conditions that determine which channel dominates the ES quality–volatility relationship. We state these as propositions derived from the theoretical framework and supported by the empirical patterns.

Condition 1: ESG quality dispersion. The CGS channel requires variation in the externality probability p_E across firms to generate variation in informed trading intensity. Our between-group dispersion (mean ES Quality: 70.16 versus 85.49, difference = 15.33) is large and generates the cross-group pattern consistent with P1. Within-group dispersion for negative externality firms (SD = 9.24) is moderate and generates the attenuated within-group pattern consistent with P4. This suggests the CGS channel is more detectable in samples with greater ES quality dispersion—for example, global samples including firms from markets with weaker ESG regulatory environments alongside FTSE-level firms.

Condition 2: Active ESG investor prevalence. CGS Proposition 12 establishes that the mechanism strengthens with the prevalence of active ESG investors. We predict that the salience interaction (β_3) is larger in firm-year observations where active ESG fund ownership is higher. This is directly testable with active fund ownership data and constitutes the most important extension of our analysis.

Condition 3: Salience intensity and persistence. Our tipping point calculations suggest the CGS channel does not fully dominate within our sample range even at peak salience. This implies that the UK market's baseline salience—already elevated relative to other markets due to mandatory TCFD reporting—reduces the incremental salience required to shift from risk-reduction to CGS dominance. In markets with lower baseline ESG salience (e.g., emerging markets, pre-regulatory environments), a larger salience shock would be required to shift channel dominance.

Condition 4: Externality magnitude and categorical exclusion. The CGS channel operates through the investor's trading decision when observing high externality realisations. When ESG fund exclusion of a sector is categorical—as for many UK ESG funds that exclude the entire oil and gas sector—within-sector ES quality variation has limited impact on trading distortions, weakening the within-group mechanism. The CGS channel is more likely to operate within-group in sectors where inclusion is graduated (proportionate to ES quality) rather than binary (categorical exclusion).

6.3 The UK as a High-ESG Equilibrium

The four dominance conditions together characterise the UK as a high-ESG institutional equilibrium: a market in which baseline ESG investor prevalence and salience are high enough to generate detectable between-group CGS patterns, but in which categorical exclusion practices and compressed within-group ES quality dispersion attenuate within-group CGS effects, leaving risk-reduction as the dominant within-group channel. This characterisation has implications for external validity. Our results should not be interpreted as universal: in markets with lower ESG investor prevalence and regulatory pressure, both channels would be weaker; in markets with higher within-group ES quality dispersion and less categorical exclusion, the within-group CGS channel would be more detectable. The UK is therefore not a neutral test environment but a specific institutional equilibrium whose features predict the observed pattern of channel dominance.

7. RESEARCH DESIGN FOR CAUSAL CHANNEL IDENTIFICATION

Our pilot evidence motivates three components of a definitive research design. Each component addresses a specific identification challenge in the current paper.

Component 1 — Annual Panel Refinitiv ESG Data. The most important upgrade is annual ES quality scores from Refinitiv (LSEG), available through WRDS from 2002 for FTSE All-Share constituents. Annual panel scores enable within-firm identification: the interaction of year-on-year ES quality changes with time-varying salience would test P4 within firms, removing all time-invariant confounds through firm fixed effects. A FTSE 350 sample with annual Refinitiv data would yield approximately 5,250 observations, providing sufficient statistical power to detect interaction effects an order of magnitude smaller than our pilot estimates.

Component 2 — Active ESG Fund Ownership as Channel Variable. Measuring the fraction of each firm's shareholder base comprised of active ESG funds (as opposed to passive ESG index funds, which the CGS framework shows do not affect effort informativeness) enables a mediation test of the mechanism. The interaction of ES_Quality with active ESG fund ownership should be positive under the CGS channel (more active ESG investors → stronger trading distortion) and zero or negative under risk-reduction. UK institutional ownership data at quarterly frequency is available from FCA transaction reporting and FactSet. This mediation test would directly identify the CGS channel variable missing from our current design.

Component 3 — UK TCFD Quasi-Experiment (Staggered DiD). The mandatory TCFD reporting requirement—FTSE 100 firms (April 2022) and FTSE 350 firms (April 2023)—provides a staggered DiD design with clean treatment timing. TCFD adoption improves externality information precision (ψ in CGS notation), which Section 6.4 of CGS predicts affects the ES quality–volatility relationship non-monotonically. Using Callaway and Sant'Anna's (2021) heteroskedasticity-robust staggered DiD estimator, we would test whether TCFD adoption shifts the salience interaction coefficient and whether the shift is in the direction predicted by the CGS precision extension. The treatment assignment (FTSE 100 versus FTSE 250 membership) is partially endogenous to firm characteristics but is determined by market capitalisation ranking rather than ESG quality, reducing (though not eliminating) selection concerns.

8. CONCLUSION

This paper establishes that environmental and social quality affects stock price volatility through two opposing channels—information distortion (CGS) and risk-reduction—whose dominance depends systematically on externality salience. A formal comparative statics analysis shows that the marginal effect of ES quality on volatility is $\partial \text{Volatility}/\partial \text{ES\_Quality} = \beta_1 + \beta_3 \times \text{Salience}$, where $\beta_1 < 0$ captures risk-reduction dominance at baseline and $\beta_3 > 0$ captures salience amplification of the CGS information distortion channel. This equation is the paper's theoretical contribution: it formally reconciles the opposite predictions of the two channels and maps directly to an estimable specification.

Using 140 firm-year observations across 20 FTSE-listed firms over 2018–2024, we find patterns consistent with this framework. Negative externality firms exhibit higher unconditional volatility (+0.0184, $p < 0.05$). A COVID-19 DiD yields a significant cross-group differential (+0.0482, $p < 0.05$), partially persistent post-COVID. The key mechanism test—an ES quality × salience interaction for negative externality firms—yields a positive coefficient convergent across three salience proxies capturing distinct investor attention dimensions (S1 = +0.0019, S2 = +0.0017, S3 = +0.0022, all $p < 0.10$). Economically, the interaction implies a 61% attenuation of the baseline risk-reduction slope during high-salience periods, corresponding to a 1.8 percentage point shift in annual price volatility per standard deviation of ES quality.

We characterise these findings as patterns consistent with salience-amplified trading distortions under maintained assumptions, rather than causal identification of the CGS mechanism. The maintained assumptions—primarily that our three salience proxies are approximately exogenous to idiosyncratic firm volatility conditional on controls—are most credible for the aggregate-level proxies (GT_Salience, ETF_Salience) and least credible for the COVID binary, which conflates ESG salience with broad market stress. The convergence across all three proxies, despite their different relationships with market-wide volatility, provides the primary credibility argument.

The practical implication is that the governance consequences of sustainable investing depend on context: both channels operate simultaneously, with dominance determined by institutional conditions rather than fixed by firm or market type. For boards of negative externality firms, ES quality improvements reduce operational risk (risk-reduction channel) regardless of salience—but during high-salience periods, they also affect the information content of stock prices through the CGS mechanism, with implications for managerial incentive costs and compensation design. For institutional investors with ESG mandates, the collective effect of their trading behaviour on price informativeness becomes most pronounced precisely when externality salience is highest—during the regulatory events, climate shocks, and social controversies when governance quality matters most.

REFERENCES

Albuquerque, R., Koskinen, Y. and Zhang, C. (2019) Corporate social responsibility and firm risk: Theory and empirical evidence. Management Science, 65(10), 4451–4469.

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

Bolton, P. and Kacperczyk, M. (2021) Do investors care about carbon risk? Journal of Financial Economics, 142(2), 517–549.

Broccardo, E., Hart, O.D. and Zingales, L. (2022) Exit versus voice. Journal of Political Economy, 130(12), 3101–3145.

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

Cao, J., Titman, S., Zhan, X. and Zhang, W. (2023) ESG preference, institutional trading, and stock return patterns. Journal of Financial and Quantitative Analysis, 58(5), 1843–1877.

Chan, Y. F. (2026). Digital Boundary Translation Theory: How Cultural Hotwords Navigate Cross-Linguistic Privacy Boundaries in Digital Consumer Behavior. Journal of International Consumer Marketing, 1–12. https://doi.org/10.1080/08961530.2025.2609081

Chan, Y.F., Ngoe, L. M., Oladapo, M. O., Osemeke, G., & Akhtar, I. (2025). When Authenticity Doesn't Pay: Validating an ESG Communication Authenticity Framework and Explaining Stakeholder–Investor Decoupling. Sustainability, 17(19), 8922. https://doi.org/10.3390/su17198922

Chan, Y. F., & Bheekee, Y. V. (2025). Digital Development Models and Transaction Costs: Empirical Evidence from Equity-Focused Versus Scale-Intensive Approaches in Emerging Economies. Economies, 13(9), 264. https://doi.org/10.3390/economies13090264

Chen, A., Gupta, D. and Starmans, J. (2026) Sustainable investing and market governance. Journal of Financial Economics, 181, 104273.

Da, Z., Engelberg, J. and Gao, P. (2011) In search of attention. Journal of Finance, 66(5), 1461–1499.

Derrien, F., Krueger, P., Landier, A. and Yao, T. (2025) ESG news, future cash flows, and firm value. Journal of Finance, 80(6), 3499–3554.

Dow, J. and Gorton, G. (1997) Stock market efficiency and economic efficiency: Is there a connection? Journal of Finance, 52(3), 1087–1129.

Edmans, A. (2009) Blockholder trading, market efficiency, and managerial myopia. Journal of Finance, 64(6), 2481–2513.

Edmans, A., Gosling, T. and Jenter, D. (2024) Sustainable investing: Evidence from the field. Working Paper.

Flammer, C. (2013) Corporate social responsibility and shareholder reaction: The environmental awareness of investors. Academy of Management Journal, 56(3), 758–781.

Goldstein, I., Kopytov, A., Shen, L. and Xiang, H. (2022) On ESG investing: Heterogeneous preferences, information, and asset prices. Working Paper.

Heinkel, R., Kraus, A. and Zechner, J. (2001) The effect of green investment on corporate behavior. Journal of Financial and Quantitative Analysis, 36(4), 431–449.

Hitzemann, S., Qin, A.A., Sokolinski, S., Tamoni, A. and Xiao, Y. (2024) Does sustainable investing dull stock reactions to cash flow news? Working Paper.

Holmström, B. and Tirole, J. (1993) Market liquidity and performance monitoring. Journal of Political Economy, 101(4), 678–709.

Kyle, A.S. (1985) Continuous auctions and insider trading. Econometrica, 53(6), 1315–1335.

Lins, K.V., Servaes, H. and Tamayo, A. (2017) Social capital, trust, and firm performance: The value of corporate social responsibility during the financial crisis. Journal of Finance, 72(4), 1785–1824.

Maug, E. (1998) Large shareholders as monitors: Is there a trade-off between liquidity and control? Journal of Finance, 53(1), 65–98.

Oehmke, M. and Opp, M.M. (2024) A theory of socially responsible investment. Review of Economic Studies, 92(2), 1193–1225.

Pástor, L., Stambaugh, R.F. and Taylor, L.A. (2021) Sustainable investing in equilibrium. Journal of Financial Economics, 142(2), 550–571.

Pástor, L., Stambaugh, R.F. and Taylor, L.A. (2022) Dissecting green returns. Journal of Financial Economics, 146(2), 403–424.

Pedersen, L.H., Fitzgibbons, S. and Pomorski, L. (2021) Responsible investing: The ESG-efficient frontier. Journal of Financial Economics, 142(2), 572–597.

APPENDIX A: VARIABLE DEFINITIONS

Variable	Definition	Source	Limitation
Volatility	SD (daily log returns)×√252. Annual.	Yahoo Finance (yfinance v0.2+)	Requires ≥200 trading days/year
Idiosyn. Volatility	Residual SD from CAPM vs FTSE All-Share.	Yahoo Finance	Market model assumes constant beta
ES_Quality	100 minus Sustainalytics ESG Risk Score.	Sustainalytics/Morningstar Apr 2024	Time-invariant; cross-sectional only
S1: H_Salience	=1 for 2020–2021; 0 otherwise.	Constructed	Conflates ESG salience with general crisis
S2: GT_Salience	Standardised Google Trends ESG search index (UK annual).	Google Trends API	Reflects attention, not direct trading behaviour
S3: ETF_Salience	Standardised aggregate ESG ETF trading volume (UK annual).	Yahoo Finance API	Reflects ESG trading activity, not firm-specific
Neg	=1 for negative externality firms.	Author classification	Binary; some firms are borderline
Leverage	Total debt / total assets.	Yahoo Finance	Point-in-time; may not match year-average
ROA	Net income / total assets.	Yahoo Finance	Affected by extraordinary items
Log_Size	ln (market capitalisation).	Yahoo Finance	End-of-year; may differ from annual average
Beta	CAPM beta vs FTSE All-Share.	Yahoo Finance	Assumes constant beta over estimation window

APPENDIX B: ROBUSTNESS TESTS (SUMMARY)

Test	Modification	β_3 (ES×Sal)	Interpretation
B1: ES Quality terciles	Replace continuous ES_Quality with tercile dummies (Low/Med/High)	Monotonic; Med−Low diff attenuated by 58% at high salience	Not artefact of cardinal ESG scaling
B2: Exclude borderline firms	Remove Diageo and Unilever from negative group (N=56)	S1: +0.0022*; S2: +0.0019*; S3: +0.0024*	Classification ambiguity does not drive result
B3: Idiosyncratic volatility	Replace total with idiosyncratic volatility	S1: +0.0016*; sign preserved across proxies	Not driven by market-wide volatility
B4: Non-COVID salience only	Estimate interaction excluding 2020–2021; use S2, S3 only	S2: +0.0014 ($p<0.15$); S3: +0.0016 ($p<0.12$)	Directional consistency outside COVID window; marginally significant

*Note: *** $p<0.01$, ** $p<0.05$, * $p<0.10$. All specifications use firm-level clustered standard errors and year fixed effects. Primary sample: N=70 firm-year observations for the negative externality group unless otherwise noted.*