UCLouvain

LOUVAIN
School of Management

Louvain School of Management

Monetary Policy Stance and Corporate ESG Performance: Evidence from S&P 500 Firms, 2007–2020

Author(s): William Cleenewerck de Crayencour
Supervisor(s): Prabal Shrestha
Academic year 2025-2026
Dissertation for the master of
Master in Management, focus on International Finance
Daytime schedule

Abstract :

This thesis examines whether a tightening of the monetary policy stance is associated with lower corporate ESG performance, whether any reduction falls evenly across the environmental, social and governance pillars, and whether board independence moderates the response. It uses an unbalanced panel of 1,667 firm-years covering 416 S&P 500 firms between 2007 and 2020, drawn from the ESG report corpus of Thewissen, Shrestha, Arslan-Ayaydin and Beibei (2026) and merged with accounting, market and rating data. The monetary policy stance is measured by the December to December change in the Wu-Xia shadow federal funds rate. A tightening of one percentage point is associated with a fall of 0.115 points in the MSCI weighted average key issue score, though the association does not survive substitution of the industry-adjusted score. The social and governance pillars decline while the environmental pillar does not, but the ordering reverses under a second rating provider. Board independence shows no moderating effect. Three of the four findings turn on how ESG performance is measured rather than on how firms behaved, which is the study's principal contribution and a caution for work drawing behavioural inferences from vendor ratings.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
Louvain School of Management
Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve
Boulevard Emile Devreux 6, 6000 Charleroi, Belgique
Chaussée de Binche 151, 7000 Mons, Belgique
www.uclouvain.be/lsm

Statements on responsible use of generative AI	Your answer
The content of my/our master's thesis is my/our original output, even if I/we used generative AI to help prepare it.	☒Yes ☐No
I/we master the theories, tools and methods that I/we use for the thesis.	☒Yes ☐No
I/we verified that the output presented in the thesis is accurate and valid, and that it contains no hallucination, no false reference, etc.	☒Yes ☐No
I/we acknowledge that I/we master the final output and that I am/we are able to explain it and answer questions about it.	☒Yes ☐No
I/we made sure not to provide generative AI tools with access to confidential data or information	☒Yes ☐No

While writing this master's thesis, several artificial intelligence tools were used to assist specific aspects of the research and writing process:

- Claude (Anthropic): used to make the writing as clear and readable as possible; to assist with the writing of R scripts used to build the datasets and estimate the models.
- NotebookLM: used to compare papers and reports
- Grammarly: used for a spelling and grammar check.

In every case, the author reviewed and verified all outputs to ensure accuracy. The hypotheses, analytical decisions, methodological choices and interpretation of the results are the author's own.

	Author last name, first name(s)	Date	Signature
1	Cleenewerck de Crayencour, William	11/08/2026	

I would like to express my sincere gratitude to my supervisor, Professor Prabal Shrestha, for his guidance throughout this thesis. His availability, his precise feedback at every stage of the work, and his willingness to discuss the methodological choices in detail have shaped this research considerably. I am also grateful to him for granting me access to the ESG report dataset underlying Thewissen, Shrestha, Arslan-Ayaydin and Beibei (2026), on which this study is built and without which this empirical study would not have been possible. Finally, I would like to thank my family and friends for their support and encouragement throughout my studies.

Table of Contents

1. Introduction 1
2. Literature review 2
2.1 ESG and why firms engage in it 2
2.2 Determinants of ESG engagement and disclosure 3
2.3 The policy environment and monetary policy 4
2.4 ESG activity under economic and financial pressure 6
2.5 Board independence and the ESG committee 9
3. Hypotheses 11
4. Data and descriptive statistics 13
4.1 Data and sample 13
4.2 Variable definitions 15
4.3 Descriptive statistics 16
4.4 Correlation matrix 19
5. Model 20
5.1 Baseline specification 20
5.2 Identification of the moderating effect 21
5.3 Estimation and inference 22
5.4 Fixed effects and unobserved heterogeneity 23
6. Results 23
6.1 Monetary policy and aggregate ESG performance 23
6.2 The result under industry adjustment 24
6.3 Composition of ESG effort across the pillars 25
6.4 Board independence as a moderator 26
6.5 The sustainability committee 27
6.6 Fixed and random effects 29
7. Robustness 29
8. Discussion 33
8.1 What the estimates show 33
8.2 The industry adjustment and interpretation of Hypothesis 1a 33
8.3 The pillar results and the divergence between providers 34
8.4 Why board independence shows no moderating effect 35

8.5 Limitations 36
8.6 Implications and directions for further work 37
9. Conclusion 37
Bibliography 39
Appendices 42

1. Introduction

Production is organised inside firms because doing so costs less than contracting for every transaction in the market, which is the reason for the firm's boundaries given by Coase (1937). Friedman (1970) draws the consequence for what managers may do with the resources under their control, holding that a manager conducts the business in accordance with the desires of its owners, which are generally to make as much money as possible while conforming to the basic rules of society. Environmental and social spending is, on that account, the use of someone else's money for a purpose they did not choose, while Freeman (1984) sets against it a firm answerable to all the groups it affects and is affected by, holding obligations to employees, customers, suppliers and communities alongside those owed to shareholders.

Firms have meanwhile made such spending standard practice, and Tsang, Frost and Cao (2023) report that eighty percent of large and mid-cap firms worldwide published a sustainability report in 2020 against twelve percent in 1993, with more than half of firms now paying for independent assurance of what they publish. Behind that reporting sits expenditure on environmental programmes, community investment and supply chain initiatives, none of which is contractually fixed in the way that payroll and debt service are.

Public authorities shape that expenditure through disclosure rules addressed to ESG directly. Still, they also shape it through monetary policy, which alters the price and availability of finance without mentioning ESG at all. Policy does not reach all corporate investment equally, and the firms whose investment responds most are identified by their financial position, whether by low default risk in Ottonello and Winberry (2020) or by youth and the absence of dividends in Cloyne, Ferreira, Froemel and Surico (2019).

If a tightening does not reach all investments equally, it may not reach all ESG investments equally either, and Hong, Kubik and Scheinkman (2012) provide a reason for expecting ESG to be exposed, finding that responsibility spending is more sensitive to financial slack than capital expenditure or research and development. This thesis therefore asks whether a tightening of the monetary policy stance reduces corporate ESG performance, whether any reduction falls evenly across the environmental, social and governance pillars, and whether board independence changes how firms respond, using a panel of 1,667 firm-years across 416 S&P 500 firms between 2007 and 2020.

2. Literature review

2.1 ESG and why firms engage in it

Environmental, social and governance activity has moved over the past three decades from a marginal concern of a small number of firms to something close to a standard feature of large-company practice. Tsang et al. (2023), reviewing the field, report survey evidence that eighty percent of large and mid-cap firms worldwide published a sustainability report in 2020, against twelve percent in 1993, and that more than half of firms now pay for independent third-party assurance of what they publish. Behind that reporting sits real expenditure, since firms fund environmental programmes, community investment, employee and supply chain initiatives, and the governance apparatus assembled to oversee them.

What distinguishes this expenditure from most other claims on a firm's cash is that none of it belongs to the operating core of the business. Environmental programmes and community investment are not contractually fixed in the way that payroll, supplier obligations and debt service are, and they can be deferred, scaled back or quietly discontinued without immediate operational consequence. ESG engagement is discretionary in a strong sense, in that it is a choice managers make in light of their circumstances rather than an obligation they discharge, and one renewed each year rather than settled once.

Legitimacy theory reads ESG engagement, and the disclosure that accompanies it, as an effort to demonstrate that the firm's conduct aligns with what society expects of it, thereby securing the informal licence it needs to operate. Impression management accounts begin from the same observation and reach a less generous conclusion. They hold that firms may engage and disclose in order to construct a favourable image rather than to inform, emphasising what reflects well on them and softening what does not (Tsang et al., 2023). Eliwa, Aboud and Saleh (2021) show how difficult the two readings are to separate in practice, finding that European lenders reward both ESG performance and ESG disclosure with a lower cost of debt, which gives firms a financial reason to pursue either, while cautioning that the disclosure they measure may partly reflect image building rather than substance.

Most large firms now publish standalone sustainability reports alongside their financial statements, but the two documents are governed very differently. Financial reporting is mandatory, standardised and audited. In contrast, ESG reporting has, for most of the period studied here, been voluntary, loosely regulated, and composed largely of narrative text rather than numbers (Tsang et al., 2023; Thewissen et al., 2026). Managers retain wide discretion over

whether to report at all, over what to include, and over how to frame it. ESG therefore involves two discretionary choices stacked one on the other: what the firm does and what it says about what it does. Both sit outside the operating core and are consequently available to be adjusted when circumstances change.

2.2 Determinants of ESG engagement and disclosure

Tsang et al. (2023) group the motives for ESG reporting into six groups: signalling, reputational insurance, the reduction of information asymmetry, managerial attributes, stakeholder demand and the cost of disclosing. These are reasons rather than predictions, and the empirical work that follows has been concerned less with why firms might disclose than with which firms actually do, concentrating on three families of firm-level characteristics, namely the board, the ownership of the firm and its financial condition.

Lagasio and Cucari (2019), aggregating 24 empirical studies in the first meta-analysis of the question, find that board independence, board size and the presence of women directors all raise the level of voluntary ESG disclosure, while board ownership and CEO duality do not, and that the evidence on the frequency of board meetings remains unsettled.

Ownership is the second established family of determinants. Bancel and Glavas (2022) organise the evidence around the investment horizon of the owner, with long-horizon institutional investors having reason to press a sustainability agenda, while short-horizon investors treat the same spending as a cost. Al Amosh and Khatib (2022) find that board independence improves disclosure quality and counteracts the negative effects of blockholder ownership, which is the moderating role taken up later in this review.

Hong, Kubik and Scheinkman (2012) study S&P 500 firms between 1991 and 2008 and show that less financially constrained firms score higher on corporate social responsibility. This result holds across every constraint measure they employ. During the internet bubble of 1996 to 2000, which relaxed financing constraints even for firms outside the technology sector, previously constrained firms raised their social responsibility scores relative to their previously unconstrained peers, and the convergence reversed once the bubble ended. Apart from that episode, the responsibility scores of constrained firms move more with their own equity valuation and cost of capital than the scores of unconstrained firms do.

Social responsibility spending is more sensitive to financial slack than capital expenditure and research and development are, a comparison Hong et al. (2012) draw within their own results.

ESG effort is not merely one discretionary use of funds among others, since it is among the first to give way when funds are scarce.

Jankensgård (2015), using Swedish disclosure data covering 2007 to 2012, reports that financially strong firms disclose more than weaker ones, that firms raising new equity or debt disclose more than firms that do not, and that the association between disclosure and financing events is strongest among financially weak firms. Taken together with Hong et al. (2012), this describes both of the discretionary choices identified earlier as responsive to the firm's financial position rather than fixed by policy.

What the established evidence has in common, across all three families, is that it explains differences between firms rather than changes within them, since board composition, ownership structure and financial strength vary far more across a cross-section than within a single firm from one year to the next. The model in Hong et al. (2012) allows a firm's financing capacity to be shifted by aggregate shocks as well as idiosyncratic ones, and their first strategy uses one. But identification in both cases rests on the difference between constrained and unconstrained firms rather than on the common movement itself, which the second strategy removes explicitly using time dummies.

Variation in financing conditions that originates outside the firm and moves every firm at once is therefore treated in this literature as something to be set aside rather than as something to be studied. Whether ESG behaviour responds to it, and whether it responds evenly across firms, remains open.

2.3 The policy environment and monetary policy

Regulation is the clearest case of a determinant of ESG behaviour that originates outside the firm, and two studies show that changes in the rules governing disclosure alter what firms do and what their disclosure is worth to them.

Allman and Won (2021), using the adoption of Directive 2014/95/EU, find that US firms exposed to it through their EU activities reduced underinvestment and raised additional debt afterwards. Yao et al. (2019), using the introduction of China's environmental disclosure measures, find cash flow sensitivity of cash falling by over thirteen percent, with the effect concentrated among high-polluting firms.

Both studies examine a mandate addressed to ESG disclosure itself, requiring firms to report more than they had reported before, so what they establish is that policy made outside the firm

changes ESG reporting and the financial consequences that follow from it, rather than the ESG activity firms undertake. Each also rests on a single event.

Neither speaks to policy that leaves ESG rules untouched while changing the financial conditions under which firms decide how much ESG activity to fund. Monetary policy is of that second kind. It is not addressed to ESG at all, it alters the price and availability of finance for every firm simultaneously, and it moves repeatedly and by varying amounts rather than once.

Bernanke and Gertler (1995) describe the credit channel, in which the direct effect of policy on interest rates is amplified by changes in the external finance premium, meaning the difference in cost between funds raised externally through debt or equity and funds generated internally through retained earnings. That premium reflects imperfections in credit markets, and a tightening that raises open market rates tends to raise it as well. The balance sheet channel operates through the effect of policy on borrowers' net worth, cash flow and liquid assets, and the bank lending channel operates through the supply of credit from banks.

Both routes bear on ESG effort because of what the determinants literature established about its sensitivity to financial slack, and about its responding more sharply to that slack than capital expenditure or research and development do (Hong et al., 2012). The credit channel describes how policy moves precisely that slack.

Ottonello and Winberry (2020) find that firms with low default risk, meaning those with low debt burdens and high distance to default, are the most responsive to monetary shocks, and interpret this through a model in which low-risk firms face a flatter marginal cost curve for external finance. Cloyne, Ferreira, Froemel and Surico (2019), using United States and United Kingdom data, reach a conclusion that runs the other way, finding the investment response concentrated among younger firms that pay no dividends, with net worth falling for all groups after a tightening but borrowing declining only for that group.

The two papers identify opposite ends of the financial spectrum as most sensitive, and the disagreement between them is genuine. What they agree on matters more for present purposes: the response to monetary policy depends on the firm's financial position rather than being common to all firms.

Bauer, Offner and Rudebusch (2024), studying European equity markets, find that brown firms, measured by carbon emission levels or intensities, are more negatively affected by tighter

monetary policy than green firms are. They examine a credit channel explanation resting on differential reliance on external finance and find no evidence for it. The difference is not accounted for by leverage, profitability or capital tangibility either. They note that several other recent studies reach the same conclusion.

Kansoy and Stasiulaitis (2025), examining roughly one hundred and sixty Federal Reserve announcements between 2005 and 2025, find that high ESG firms are protected against surprises about the immediate policy rate but are more vulnerable to surprises about the future policy path. They also report that the relationship inverted around the Paris Agreement, with high ESG firms more vulnerable to contractionary surprises within their industries beforehand and protected afterwards.

Salzmann (2025), using a European panel, finds that macroeconomic variables including the central bank policy rate and short term interbank rates significantly affect both ESG and conventional indices, and Mohammad, Hussain and Adnan (2025) find that higher ESG scores weaken the risk taking channel of monetary policy transmission among South Asian financial institutions.

ESG characteristics evidently matter for how monetary policy propagates, but ESG occupies the same position across these four studies. In Bauer et al. (2024), in Kansoy and Stasiulaitis (2025) and in Salzmann (2025), ESG is a fixed characteristic of a firm or an index, and the outcome measured is a price, a return or a sensitivity.

Two bodies of work therefore stop just short of meeting. The determinants literature has shown that a firm's financial condition shapes how much ESG effort it makes, but has studied only variation in financial condition that differs across firms at a point in time. The monetary policy literature supplies variation in financial conditions that is external to the firm, common to all firms and repeated over time, but treats ESG as an explanatory characteristic rather than as an outcome that policy might change.

2.4 ESG activity under economic and financial pressure

What the determinants literature did not establish is how ESG effort responds when the pressure arrives from outside the firm and reaches many firms at once, a question on which three studies bear and do not agree. Harrison and Berman (2016) ask whether firms alter their social responsibility programmes as economic growth changes, expecting, on the logic set out earlier, that discretionary activity will be cut when conditions deteriorate. They study 837 firms over

1995 to 2009 using the annual change in real GDP per capita, and they measure social performance with KLD ratings, which are the ancestors of the MSCI ESG ratings in use today.

Concerns record negative outcomes such as pollution incidents, legal suits and violations, which arise largely from neglect. In contrast, strengths record positive programmes such as employee benefits, charitable contributions and pollution prevention, which require resources. Harrison and Berman keep the two apart rather than combining them, on the expectation that they respond differently to economic conditions, and they do.

Concerns rise when growth is weak, significantly so in community relations, employee relations, the environment, and product safety and quality, while diversity concerns show no relationship with the cycle. Firms neglect more when the economy is weak and restore vigilance when it recovers.

Positive programmes move the opposite way to the authors' own prediction, increasing when growth is weak and decreasing when it recovers, with significant effects in diversity, which is much the strongest, in the environment, and more weakly in product safety and quality. At the same time, employee and community strengths show no significant relationship. Both sides therefore rise together in a downturn, and in the environment the same firms simultaneously expand positive programmes and accumulate more negative incidents. Harrison and Berman describe the strengths result as surprising and suggest that firms may be using visible positive activity strategically to protect their reputations when conditions are difficult, stating explicitly that they offer this reading with reservations and that further work is needed before it can be asserted with confidence.

Sakunasingha, Jiraporn and Uyar (2018) approach the question through a single severe episode, the Great Recession of 2008, on the reasoning that an adverse shock forces firms to prioritise and thereby reveals what they treat as essential. Firms did not reduce overall CSR investment during the recession, but the composition moved substantially, falling in community, employee, environment, human rights and product, and rising in corporate governance and diversity, which the authors read as evidence that firms regard some categories as closer to the strategic core of the business than others.

Hong et al. (2012), discussed earlier as the strongest firm-level evidence that financial slack drives ESG effort, point toward a third pattern, one that concerns the level rather than the mix, with less constrained firms scoring higher overall and their responsibility spending proving more sensitive to slack than capital expenditure or research and development. In a robustness

exercise, they rerun the tests on the six individual components of the score and obtain similar signs across almost all categories, with corporate governance as the exception.

Corporate governance rises in Sakunasingha et al. (2018) while five other categories fall, and behaves unlike the rest in Hong et al. (2012), while diversity rises during the recession in the former and is the most strongly counter-cyclical of all the strengths in Harrison and Berman (2016). Different shocks, different periods and different empirical strategies isolate the same two components as behaving unlike the others.

Harrison and Berman (2016) use economic growth and Sakunasingha et al. (2018) use a recession, both of which are contractions that combine falling demand, uncertainty, declining revenue and tighter finance into a single movement. In contrast, Hong et al. (2012) isolate financing capacity and hold the remainder constant by construction. A general contraction reduces the resources available for discretionary spending while at the same time raising the value of appearing to behave well. This is the mechanism Harrison and Berman invoke to explain why positive programmes expand exactly when neglect increases, and a change in the cost of finance that comes without a collapse in demand carries no comparable reputational trigger.

A study measuring only the total reported by Sakunasingha et al. (2018) would have concluded that the Great Recession left corporate social responsibility untouched, and would have missed a reallocation across seven categories moving in two directions. An aggregate rating cannot distinguish between a firm that did not respond and a firm whose responses cancelled, since both produce a flat score, which makes component-level evidence a requirement for interpreting a null result rather than an extension of the analysis.

The evidence assembled here establishes that ESG activity responds to conditions outside the firm rather than being fixed, that the response differs by component with governance and diversity repeatedly separating from the rest, that positive activity and negative incidents can move together rather than trading off, and that a composite measure can conceal all of it. What it does not establish is how ESG activity responds to a change in financing conditions that is external to the firm, common across firms, and unaccompanied by a general economic contraction, since the cyclical studies supply the external and common shock but confound financing pressure with everything else a contraction brings, including the reputational incentive that appears to push positive activity the other way, while the financing study isolates the right mechanism but identifies it from variation specific to each firm.

2.5 Board independence and the ESG committee

Board independence has historically been assessed through the absence of financial, family or professional ties between a director and the company on whose board the director sits, and two theoretical accounts explain why the presence of such directors should affect a firm's social and environmental conduct (Ortas, Álvarez & Zubeltzu, 2017). Because both predict the same association under ordinary conditions, the distinction between them is easily overlooked, becoming consequential only as circumstances change.

Agency theory holds that independent directors mitigate conflicts of interest between shareholders and managers and provide shareholders with a protection mechanism against managerial discretion (Ortas et al., 2017), which is the governance expression of the view of corporate purpose set out by Friedman (1970). Liao, Luo and Tang (2015) apply the same reasoning to environmental reporting, describing independent directors as a monitoring mechanism restricting the opportunistic behaviour of top executives. Independence, on this reading, disciplines managers in the interest of the providers of capital, and ESG activity is subject to that discipline like any other use of corporate resources.

Stakeholder theory (Freeman, 1984) begins instead from the position that a company is an open system affected by, and affecting, agents both inside and outside it, so that boards with greater independence are more likely to weigh interests other than those of managers and majority shareholders (Ortas et al., 2017). Liao et al. (2015) frame their own results in these terms, arguing that an independent board can mediate between financial and non-financial demands under limited resources.

Ortas et al. (2017) aggregate 87 published papers addressing corporate social performance, understood as the actual social, environmental and economic results of the organisation rather than its accountability to society, and report a positive and significant association with board independence, with an average correlation of 0.1258 and a confidence interval of 0.0946 to 0.1566. The association is weaker among studies using externally produced data than among those using company self-reported figures, which the authors attribute to social desirability bias. Independent boards are therefore associated with stronger social performance, though the effect is highly heterogeneous, which the authors read as evidence that additional moderating variables operate within the relationship.

Which of the two accounts is correct cannot be settled by this evidence, because both predict the same positive association under ordinary conditions. Jizi (2017), using FTSE 350 firms,

advances both within a single paper, describing independence as directing scarce resources toward value maximising social projects in one passage and as reflecting a wider outlook less confined to financial measures in another. An independent board acting for the providers of capital should, under the agency account, withdraw discretionary non-core spending more readily when finance becomes expensive, whereas under the stakeholder account a board attentive to interests beyond the shareholder should be more willing to sustain such spending through a period of pressure. On that question the literature is close to silent.

The one direct attempt at that comparison is made by Ortas et al. (2017), whose hypothesis holds that the link between independence and social performance weakens in bear market periods. They reason that boards prioritise financial performance over social performance when conditions are adverse. They find the association positive and significant in every period after 2002 and growing in strength over time, and state that it contradicts the expectation that firms in adverse economic settings reduce their attention to social issues in order to cut costs. Their comparison is drawn across studies conducted at different times rather than within firms, so it cannot be separated from the growth in ESG attention over the same decades.

Al Amosh and Khatib (2022) place board independence in a moderating role, as it dampens the negative influence of blockholder ownership on ESG disclosure among Jordanian industrial firms. Mohammad et al. (2025), discussed earlier, similarly treat ESG characteristics as moderating the transmission of monetary policy.

Baraibar-Diez and Odriozola (2019) ground a dedicated board committee for corporate social responsibility in stakeholder theory, on the reasoning that such a committee creates a governance body capable of addressing stakeholder needs and delegates the monitoring of sustainability issues to members with relevant expertise. Studying 197 listed firms in Spain, France, Germany and the United Kingdom between 2005 and 2015, they find that firms with a committee record significantly higher environmental, social, governance and economic performance scores across all four countries, with corporate governance in Germany and the economic score in Spain as the two exceptions.

The effect they observe is broadly homogeneous across the four scores rather than concentrated in any one of them, which contrasts with the component-level divergence documented in the cyclical literature reviewed earlier, and around ninety percent of their sample had established such a committee by 2014, which indicates how rapidly the practice became standard among large listed firms and how little a simple indicator of presence may distinguish between them.

Elmghaamez, Nwachukwu and Ntim (2024), treating sustainability committees as moderators of the relationship between ESG disclosure and firm performance across 500 multinational enterprises in 40 countries, report results that vary with the performance measure selected, and record that the prior literature on such committees is contradictory.

Board independence is reliably, if modestly, associated with stronger social performance, across markets and across financial and non-financial firms, and a dedicated committee is associated with stronger performance on evidence that is newer and more contested. What this literature does not establish is which of the two accounts underlies the association, and consequently whether either governance feature makes a firm's ESG behaviour more or less sensitive when the financial conditions under which ESG decisions are made are altered from outside the firm.

3. Hypotheses

The literature reviewed above establishes that ESG engagement and disclosure are discretionary choices, that they respond to the financial condition of the firm making them, and that monetary policy alters that condition for all firms simultaneously through the cost and availability of external finance. What follows sets out what this study expects to observe when those three findings are brought together, and identifies in each case the argument that would produce the opposite result.

Contractionary monetary policy raises the external finance premium and reduces the net worth and cash flow available to firms, which narrows the resources available for spending that is not contractually fixed. Hong, Kubik and Scheinkman (2012) show that responsibility spending falls when financing capacity tightens, and that it does so more sharply than capital expenditure or research and development, placing ESG effort among the first claims on cash to be released rather than among the last. Applying that mechanism to a change in financing conditions that originates outside the firm and reaches all firms at once gives the first expectation.

> **H1a.** A tightening of the monetary policy stance is associated with a reduction in the aggregate ESG performance of firms.

Working against this is the reputational argument that Harrison and Berman (2016) advance to explain why positive social programmes expanded during periods of weak economic growth in their sample. If managers read a monetary tightening as a general deterioration in conditions,

the incentive to signal good conduct could offset the reduction in available resources. The reason for expecting the first effect to dominate is that a change in the cost of finance is far less publicly salient than a recession, and therefore carries much weaker reputational pressure while carrying the same financial pressure.

An aggregate rating nets positive actions against negative incidents and averages across areas that the literature shows do not move together. This means that a firm reducing effort in some areas while maintaining or increasing it in others would register little change in its overall score. Sakunasingha et al. (2018) found precisely this pattern during the Great Recession, with total investment unchanged and seven categories moving in two directions. Hong et al. (2012) report that the governance component does not behave like the others when financing conditions change. Governance activity also differs from environmental and social activity, resting largely on board structures, policies and procedures rather than on continuing programme expenditure, which makes it less exposed to a contraction in existing funds.

> **H1b.** A tightening of the monetary policy stance is associated with a change in the composition of ESG effort across the environmental, social and governance pillars, with governance declining least.

Baraibar-Diez and Odriozola (2019) provide the counterweight here, having found the effect of a governance intervention broadly homogeneous across the four scores they examine rather than concentrated in any of them. Their variation is the cross-sectional presence of a committee rather than a shock, however, so their finding speaks to whether pillars differ in their determinants rather than to whether they differ in their response to pressure.

Independent directors are less closely tied to management and are drawn from a wider range of backgrounds, which the governance literature associates with attention to interests beyond the immediate financial position of the firm and with a longer view of what serves it (Liao et al., 2015; Ortas et al., 2017). Liao et al. describe an independent board as able to balance financial against non-financial goals under limited resources, which is the condition a monetary contraction creates. Al Amosh and Khatib (2022), in the only study in this literature that places independence in a moderating role, find it buffers the influence of other governance factors that push ESG disclosure downward.

> **H2.** Board independence weakens the relationship described in H1a, such that firms with more independent boards reduce ESG effort less in response to a monetary tightening.

The opposing prediction follows from the agency reading of ESG spending, under which some corporate goodness represents managers consuming a private benefit at the shareholders' expense, a possibility that Hong et al. (2012) acknowledge directly. On that reading, an independent board should be quicker rather than slower to withdraw such spending when finance becomes expensive, which would strengthen rather than weaken the relationship in H1a. Both readings appear in the same sources, and neither is settled by the existing evidence, so the direction stated above rests on the moderating evidence rather than on theory alone.

A dedicated sustainability committee is examined as a secondary moderator alongside board independence, on the expectation that it operates in the same direction and more weakly. Because committee status changes within firms too rarely across this sample to be treated as a time varying characteristic, it is used as a fixed attribute of the firm, and the corresponding test therefore concerns whether firms that maintain such a committee respond differently rather than whether establishing one changes how a firm responds.

4. Data and descriptive statistics

4.1 Data and sample

The analysis uses an unbalanced panel of 1,667 firm-year observations covering 416 firms listed in the S&P 500 index over the period 2007 to 2020. The unit of observation is the firm-year, and a firm enters the panel in a given year when it published a standalone ESG report in that year and when the accounting, market and rating data required by the model are available.

The ESG reports underlying the sample are drawn from the corpus assembled by Thewissen et al. (2026), who collect 7,565 standalone ESG reports published by 1,715 United States firms between 1998 and 2023. The present study restricts attention to S&P 500 constituents and to the years 2007 to 2020. The lower bound reflects the point at which standalone ESG reporting becomes sufficiently widespread among large United States firms for a panel to be meaningful, and it places the start of the sample immediately before the financial crisis, so that the full sequence of policy easing, the zero lower bound period and the subsequent normalisation is observed. The upper bound reflects the coverage of the rating data.

ESG performance is measured using MSCI ESG ratings. The main dependent variable is the weighted average key issue score (hereafter the weighted average score), which measures a firm’s rating in absolute terms and is not benchmarked against industry peers. It is available

for 1,192 of the 1,667 firm-years, so the tests of Hypothesis 1a are conducted on that panel. The industry-adjusted score, which expresses the same underlying rating relative to the range observed within the firm's industry in the same year, is reported alongside it throughout as a robustness measure and is available for the full sample. The three pillar scores are drawn from the same source and cover the same 1,192 firm-years, on which the tests of Hypothesis 1b are conducted.

Accounting and market data are taken from Compustat and CRSP. Governance data, specifically the proportion of independent directors, board size and the presence of a sustainability committee, are collected from Bloomberg. Board independence is available for 1,601 firm-years, or 96 percent of the sample. Bloomberg records these fields with a reporting date fixed at 31 October rather than at the firm's fiscal year end, and the consequences of this timing convention are examined in the robustness analysis.

Over much of the sample, the effective federal funds rate is uninformative about the stance of policy, as the target range stayed between 0% and 0.25% from December 2008 to December 2015. Half of the period studied in this thesis falls within the zero lower bound, during which the Federal Reserve acted through asset purchases and forward guidance rather than through the policy rate.

Shadow rates were developed for precisely this difficulty, coinciding with the observed policy rate when it is unconstrained and taking negative values when it is not, so that a single continuous series describes the stance across the whole sample period. The idea originates with Black (1995), and Krippner (2012) sets out the case for reading such rates as measures of the policy stance at the zero lower bound. Anderl and Caporale (2023) identify the model of Wu and Xia (2016) as one of the three standard approaches, and de la Horra, Perote and de la Fuente (2021) use that rate to measure the United States monetary policy stance in a firm-level panel over a period largely overlapping the one studied here.

The stance is accordingly measured by the shadow federal funds rate (hereafter named shadow rate) of Wu and Xia (2016). The main regressor is the December to December change in this rate, so that a positive value denotes a tightening of the stance over the year.

The panel is unbalanced by construction, since firms enter and leave according to their reporting behaviour. The number of observations rises from 25 in 2007 to 262 in 2019. Then it falls to 63 in 2020, the decline in the final year reflecting incomplete coverage of reports

published late in the period rather than a fall in reporting activity. On the firm side, 140 of the 416 firms appear in a single year, and only one firm is observed in all fourteen years, with a median of three observations per firm. This structure has direct consequences for the choice of fixed effects, which are set out in the following section.

Firms in the financial sector are retained. Financial institutions account for 186 observations and insurance and real estate for a further 51, together 14.2 percent of the sample. Excluding them is common in studies of corporate investment, where the interpretation of leverage and capital structure differs for intermediaries. Still, the outcome studied here is ESG performance rather than investment, and the industry adjustment embedded in the rating already benchmarks each firm against its own sector. Ellili (2023) similarly retains financial firms in an ESG setting. The remainder of the sample is spread across fourteen industries, the largest being durable manufacturers with 332 observations, computers with 232 and retail with 182.

4.2 Variable definitions

Table 1 lists each variable used in the analysis together with a short definition and its source.

Table 1. Variable definitions

Variable	Definition	Source
ESG	MSCI weighted average key issue score, measuring ESG performance in absolute terms on a scale from 0 to 10.	MSCI
ESG (adj.)	MSCI industry-adjusted score, expressing the firm's weighted average key issue score relative to the range observed in its industry, on a scale from 0 to 10.	MSCI
E	MSCI environmental pillar score, on a scale from 0 to 10.	MSCI
S	MSCI social pillar score, on a scale from 0 to 10.	MSCI
G	MSCI governance pillar score, on a scale from 0 to 10.	MSCI
ΔSHADOW	Change in the Wu-Xia shadow federal funds rate from December of the previous year to December of the current year, in percentage points; a positive value denotes tightening.	Wu and Xia (2016)
INDEP	Proportion of the board of directors classified as independent, in percent.	Bloomberg

Variable	Definition	Source
COMMITTEE	Indicator equal to one if the firm has a corporate social responsibility or sustainability committee at board level, and zero otherwise, measured at the firm's first appearance in the panel.	Bloomberg
SIZE	Natural logarithm of total assets at the end of the fiscal year.	Compustat
ROA	Income before extraordinary items divided by total assets at the end of the previous fiscal year, in percent.	Compustat
LEV	Short-term plus long-term debt divided by total assets at the end of the fiscal year, in percent.	Compustat
VOLAT	Standard deviation of return on assets over the preceding five fiscal years.	Compustat
RD	Research and development expenditure divided by sales, in percent, set to zero for firms that report no research and development expenditure.	Compustat
NoA	Natural logarithm of one plus the number of analysts issuing an earnings forecast for the firm in the fiscal year.	I/B/E/S

Note. The table lists each variable used in the analysis with its definition and source. All continuous control variables are winsorised at the first and ninety-ninth percentiles.

4.3 Descriptive statistics

Table 2 reports the distribution of each variable over the estimation sample. Variables that enter the model in logarithms are reported both in levels and in logarithms, so that the economic magnitudes remain interpretable.

Table 2. Descriptive statistics

Variable	N	Mean	Median	SD	Min	P25	P75	Max
ESG (weighted average)	1,192	4.983	4.895	1.065	1.480	4.279	5.600	9.450
ESG (industry-adjusted)	1,667	5.297	5.400	2.228	0.000	3.730	6.900	10.000
E	1,192	5.611	5.391	2.105	0.000	4.196	6.926	10.000
S	1,192	4.616	4.567	1.650	0.000	3.646	5.602	9.967
G	1,192	5.715	5.567	1.607	0.417	4.767	6.535	10.000
ΔSHADOW	1,667	0.035	-0.288	1.261	-3.049	-0.804	0.986	2.678
INDEP	1,601	85.797	88.889	8.341	33.333	81.818	91.667	100.000
Board size	1,603	11.155	11.000	2.101	5.000	10.000	12.000	24.000
COMMITTEE	1,598	0.371	0.000	0.483	0.000	0.000	1.000	1.000
Total assets (USD bn)	1,667	72.5	18.0	237.7	0.2	6.7	48.1	2,687.4
SIZE (log)	1,667	9.871	9.797	1.475	6.658	8.810	10.781	14.081
ROA	1,667	6.634	5.922	5.857	-11.327	2.769	9.764	23.675
LEV	1,667	27.307	26.083	15.830	0.000	14.962	37.925	69.124
VOLAT	1,667	2.958	2.028	2.948	0.084	1.051	3.761	15.204
RD	1,667	4.235	0.474	6.893	0.000	0.000	5.396	28.175
Analysts (n)	1,667	15.711	15.000	7.182	3.000	11.000	20.000	47.000
NoA (log)	1,667	2.712	2.773	0.477	1.386	2.485	3.045	3.638

Note. ROA, LEV, VOLAT and RD are reported after winsorisation at the first and ninety-ninth percentiles. Total assets and analyst coverage are shown in raw form, while SIZE (log) and NoA (log) are computed on their winsorised values. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

The weighted average score, the main dependent variable, averages 4.98 with a median of 4.90 and a standard deviation of 1.07. The aggregate ESG score, reported alongside it, averages 5.30 with a median of 5.40 and a standard deviation of 2.23, spanning the full range from zero to ten. Across the three pillars, governance records the highest average at 5.72 and social the lowest at 4.62, with environmental performance in between at 5.61. Social performance is also

the least dispersed of the three, with a standard deviation of 1.65 against 2.11 for the environmental pillar. The gap between the pillars is relevant to Hypothesis 1b, which concerns the composition of ESG effort rather than its level, and the descriptive pattern establishes that the pillars differ enough for a reallocation between them to be detectable.

The change in the shadow rate averages 0.035 percentage points across firm-years, with a median of -0.288 and a standard deviation of 1.26. The maximum (2.68 percentage points) corresponds to the exit from the zero lower bound and is examined separately in the robustness analysis. Because the variable takes only fourteen distinct values, one per year, its dispersion across the 1,667 firm-years reflects the composition of the panel as much as the underlying series.

Board independence averages 85.8 percent with a median of 88.9 percent, and half the sample lies between 81.8 percent and 91.7 percent. High and relatively homogeneous board independence is characteristic of large United States firms in this period, and reflects the listing requirements introduced in the wake of the Sarbanes-Oxley Act. Meaningful variation nonetheless remains, with the lowest observed value at 33.3 percent. Still, the concentration of the distribution implies that the moderating effect tested in Hypothesis 2 is identified from a relatively narrow band of the possible range. Boards contain 11.2 members on average, and 37.1 percent of firm-years correspond to firms that maintain a board-level sustainability committee.

Turning to the control variables, total assets are strongly right-skewed, with a mean of 72.5 billion dollars against a median of 18.0 billion and a maximum above 2.6 trillion, which is why this variable has been logged. Research and development intensity has a median of 0.47 percent against a mean of 4.24 percent, reflecting the fact that a little over half the sample reports no research and development expenditure at all, so the variable is best read as separating research-intensive industries from the rest.

Two variables required attention before estimation. Earnings volatility reached 553 in its raw form against a median of 2.03, driven by a small number of firms with near-zero asset bases in the denominator, and leverage exceeded 100 percent of total assets for a few observations. Rather than delete the affected firm-years, which the panel cannot well afford given that 140 firms are observed only once, all continuous control variables were winsorised at the first and ninety-ninth percentiles. The effect is visible in Table 2: the maximum of earnings volatility

falls from 553 to 15.2 and that of leverage from 100.1 to 69.1, while every median is unchanged.

4.4 Correlation matrix

Table 3 reports pairwise Pearson correlations between the dependent variable, the policy variable and the explanatory variables used in the analysis, computed over the estimation sample.

Table 3. Correlation matrix

	ESG	ΔSHAD.	INDEP	SIZE	ROA	LEV	VOLAT	RD	NoA
ESG	1.000								
ΔSHADOW	-0.149	1.000							
INDEP	-0.007	0.030	1.000						
SIZE	0.062	0.003	0.229	1.000					
ROA	0.110	-0.039	-0.047	-0.127	1.000				
LEV	-0.012	0.021	0.077	-0.065	-0.031	1.000			
VOLAT	-0.090	-0.054	-0.091	-0.212	0.006	0.107	1.000		
RD	0.118	-0.029	-0.009	-0.035	0.145	-0.102	0.324	1.000	
NoA	0.044	-0.002	0.072	0.454	0.159	-0.065	0.078	0.124	1.000

Note. Correlations are computed on the 1,192 firm-years of the main estimation sample, except for those involving board independence, which use the 1,601 observations for which that variable is available. Control variables are winsorised at the first and ninety-ninth percentiles, and firm size and analyst coverage enter in logarithms. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

No pair of explanatory variables is correlated strongly enough to raise a concern about multicollinearity. The largest value in the matrix is 0.454, between firm size and analyst coverage, which is expected since larger firms attract more analysts, and it remains well below the level at which coefficient estimates become unstable. The two variables nonetheless capture distinct constructs: size proxies for resources and public visibility, while analyst coverage proxies for the intensity of external monitoring by capital market participants. Retaining both is therefore preferred to dropping either.

The correlations involving the dependent variable are informative about what follows, though they should be read with care, since they contain no controls and no fixed effects. ESG performance correlates positively with profitability at 0.110 and with research and development intensity at 0.118. The association with firm size is positive but weak at 0.062, which is lower than the literature on ESG determinants would lead one to expect. The simple correlation between the change in the shadow rate and the weighted average score is -0.149, against -0.052 for the industry-adjusted score, while the two measures correlate similarly with every other control variable. The sign is the one Hypothesis 1a predicts, and although the coefficient is modest in absolute terms, it is the strongest association the dependent variable displays in the matrix. Neither fact carries evidential weight on its own.

5. Model

5.1 Baseline specification

Hypothesis 1a states that a tightening of the monetary policy stance is associated with a retreat in aggregate ESG performance. It is tested with the following equation, estimated on the 1,192 firm-years for which the main dependent variable is available:

$$ESG_{ijt} = \alpha + \beta\, \Delta SHADOW_t + \gamma' X_{ijt} + \delta_j + \varepsilon_{ijt} \qquad (1)$$

The dependent variable is the MSCI weighted average score of firm i, belonging to industry j, in year t, which measures ESG performance in absolute terms rather than relative to industry peers. The industry-adjusted score is used as a robustness measure and is available for the full panel of 1,667 firm-years. The regressor of interest is the December to December change in the Wu-Xia shadow federal funds rate, which is positive when the stance tightens over the year. The vector X contains firm size, return on assets, leverage, earnings volatility, research and development intensity and analyst coverage, all defined in Table 1. The term δ_j denotes industry fixed effects, which allow the average level of ESG performance to differ across sectors and absorb any permanent sectoral characteristic, such as exposure to environmental regulation or the typical composition of a firm's stakeholder base.

The equation contains no year fixed effects. This follows from the nature of the regressor rather than from a modelling preference: the change in the shadow rate takes a single value in each year, so it is perfectly collinear with a full set of year indicators, and its coefficient cannot be estimated once they are included. The estimation therefore relies on variation over time in the

policy stance, and β measures the association between that stance and ESG performance across the fourteen years of the sample.

The consequence is that β absorbs any other influence that moves at annual frequency and is correlated with the policy stance. Monetary policy does not tighten at random: it tightens when output is growing and inflation is rising, and it eases when conditions deteriorate. The coefficient is therefore reported as an association rather than as an identified causal effect, and the discussion returns to this point. The direction of the resulting bias is, however, predictable. Since policy tightens in good times, and firms have more discretionary resources for ESG effort in good times, the omitted business cycle should push β upward, that is, against the negative sign that Hypothesis 1a predicts. A negative estimate is therefore conservative with respect to the hypothesis, while a positive estimate would be difficult to interpret as evidence against it.

Hypothesis 1b concerns the composition of ESG effort rather than its level. It is tested by estimating the same equation three times, replacing the aggregate score with the environmental, social and governance pillar scores. These regressions are run on the same panel as Hypothesis 1a. Hypothesis 1b predicts that the three coefficients differ, and specifically that the governance pillar declines by less than the environmental and social pillars.

5.2 Identification of the moderating effect

Hypothesis 2 states that board independence dampens the relationship established in Hypothesis 1a. Testing a moderating effect requires a different specification, and it also resolves the identification problem described above.

When a regressor takes a common value for every firm in a given year, its average effect cannot be separated from other aggregate conditions. Its differential effect can. If firms differ along some dimension that varies within the year, and if the response to the policy variable differs systematically along that dimension, then the interaction between the two is identified even when the policy variable itself is absorbed. This is the approach taken by Ottonello and Winberry (2020) in the context of corporate investment. The dimension along which firms are compared here is board independence, which is the moderator that Hypothesis 2 concerns:

$$ESG_{ijt} = \alpha + \beta_1\left(\Delta SHADOW_t \times INDEP_{ijt}\right) + \beta_2 INDEP_{ijt} + \gamma' X_{ijt} + \delta_j + \tau_t + \varepsilon_{ijt} \quad (2)$$

The term τ_t denotes year fixed effects. The main effect of the policy variable is absorbed by these and is therefore omitted from the equation. Board independence is centred on its sample mean, so that in specifications where the main effect of the policy variable is estimable it can

be read at the average level of board independence rather than at a board with no independent directors. This value does not occur in the sample. In the equation above, the centring does not affect β_1.

The coefficient β_1 is the parameter of interest for Hypothesis 2. Because the year fixed effects absorb everything that is common to all firms within a year, β_1 is not vulnerable to the omitted aggregate factor that limits the interpretation of β in the baseline equation. Output growth, credit conditions, market sentiment and any other macroeconomic influence are removed regardless of whether they are observed. What remains is the difference in how firms with more and less independent boards respond to the same policy environment, and this specification therefore carries the identifying content of the analysis.

The interpretation of β_1 should nonetheless be stated precisely. It identifies heterogeneity in the response, not the average response. A finding that firms with more independent boards retreat less does not by itself establish that firms retreat, since the common component of any retreat is absorbed by the year fixed effects. The two layers are complementary: the baseline equation speaks to the level of the relationship without identifying it, and the interaction identifies a difference without speaking to the level.

The presence of a board-level sustainability committee is examined as a secondary moderator, entering through its interaction with the policy variable in the same specification. It is treated as a time-invariant firm attribute, fixed at the value recorded when the firm first appears in the panel. This reflects the structure of the data rather than a belief that committee status never changes. Only 85 of the 416 firms are ever observed on both sides of the threshold, most of them with a single year on one side, and the recorded transitions are concentrated in the years from 2018 onwards. Treating the variable as time-varying would identify the effect from a small number of firms within a narrow window.

5.3 Estimation and inference

All equations are estimated by ordinary least squares with the fixed effects absorbed. Standard errors are clustered at the firm level, which allows the residuals of the same firm to be correlated over time and permits arbitrary heteroskedasticity across firms. Clustering at the firm level is standard in panel settings of this kind and addresses the most likely form of dependence, namely persistence in a firm's ESG performance from one year to the next.

One limitation of this choice should be acknowledged. Because the policy variable takes only fourteen distinct values, one for each year, residuals may also be correlated across firms within

a year in ways that firm-level clustering does not capture. Inference on the coefficient of the policy variable in the baseline equation should therefore be read with some caution. The issue is much less severe for the interaction term, since the year fixed effects already absorb the common component within each year.

The full control coefficients for every specification reported are given in the appendices.

5.4 Fixed effects and unobserved heterogeneity

The main specification uses industry rather than firm fixed effects, following Thewissen et al. (2026) on the same corpus. A firm fixed effects estimator uses only variation within firms over time, and 140 of the 416 firms appear in a single year and therefore contribute nothing to it, with the median firm observed three times. A firm fixed effects specification is nonetheless reported in the robustness analysis.

Random effects estimation is not used. It would treat the unobserved firm-specific component as a random draw rather than removing it, which requires that component to be uncorrelated with the explanatory variables. That assumption is implausible here, since the unobserved characteristics driving ESG performance are unlikely to be independent of firm size, leverage and board composition.

6. **Results**

6.1 Monetary policy and aggregate ESG performance

Table 4 reports the estimation of equation (1). Column 1 gives the main specification, in which the dependent variable is the MSCI weighted average score. The coefficient on the change in the shadow rate is -0.115 with a clustered standard error of 0.024, significant at the one percent level. A tightening of one percentage point is therefore associated with a fall of 0.115 points on a measure whose standard deviation is 1.07, or roughly eleven percent of a standard deviation. Expressed against the variation actually observed in the policy series, a one standard deviation tightening of 1.26 percentage points is associated with a fall of 0.145 points, or 0.14 of a standard deviation of the dependent variable. Hypothesis 1a is supported.

This specification is estimated on the 1,192 firm-years for which the weighted average score is available, covering 315 firms, rather than on the full panel of 1,667. The restriction follows from the coverage of the rating extract, rather than from any selection made. The industry-

adjusted score is reported on both samples in columns 2 and 3, so that the consequences of the restriction can be seen.

Among the control variables, profitability enters positively and earnings volatility negatively; both significant at the one percent level. More profitable firms record higher ESG scores and firms with less predictable earnings record lower ones, which is consistent with the slack resources argument presented in the literature review. Firm size, leverage, research intensity and analyst coverage are not significant in this specification.

Table 4. Monetary policy and aggregate ESG performance

	(1) Weighted average score	(2) Industry-adjusted score	(3) Industry-adjusted score
ΔSHADOW	-0.1147***	-0.0323	-0.0554
	(0.0235)	(0.0566)	(0.0456)
Controls	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
Year FE	No	No	No
Observations	1,192	1,192	1,667
Firms	315	315	416
Adjusted R^2	0.208	0.152	0.097

Note. The dependent variable in column 1 is the MSCI weighted average score, which is not benchmarked against industry peers. Columns 2 and 3 use the MSCI industry-adjusted score, over the same firm-years as column 1 and over the full panel respectively. ΔSHADOW is the December to December change in the Wu-Xia shadow federal funds rate, in percentage points, so that a positive value denotes a tightening of the monetary policy stance. Controls are firm size, return on assets, leverage, earnings volatility, research intensity and analyst coverage, as defined in Table 1. Their coefficients are reported in full in Appendices 1 and 2. Year fixed effects are not included, since the policy variable takes a single value per year and would be absorbed by them. Standard errors clustered at the firm level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

6.2 The result under industry adjustment

The finding in column 1 does not survive the substitution of the industry-adjusted score. Over the same 1,192 firm-years, the same specification returns a coefficient of -0.032 with a standard error of 0.057, which is not distinguishable from zero. On the full panel of 1,667 firm-years, the estimate is -0.055 with a standard error of 0.046 and likewise insignificant. The result

reported above is therefore not robust to the choice of ESG measure, and this qualification should be carried through the interpretation of Hypothesis 1a.

The industry-adjusted results are reported throughout the robustness analysis alongside the main specification, and the reasons for the divergence between the two measures are examined in the discussion.

6.3 Composition of ESG effort across the pillars

Hypothesis 1b holds that a tightening of the policy stance changes the composition of ESG effort across the three pillars, and specifically that governance declines by less than the environmental and social dimensions. Table 5 reports equation (1) estimated separately on each pillar, together with the aggregate score restricted to the same 1,192 observations for comparability.

Table 5. Monetary policy and the three ESG pillars

	Aggregate (weighted average)	Aggregate (industry-adjusted)	Environmental	Social	Governance
ΔSHADOW	-0.1147***	-0.0323	0.0851	-0.1802***	-0.2519***
	(0.0235)	(0.0566)	(0.0438)	(0.0379)	(0.0356)
Controls	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Observations	1,192	1,192	1,192	1,192	1,192
Adjusted R^2	0.208	0.152	0.206	0.193	0.134

Note. Controls are firm size, return on assets, leverage, earnings volatility, research intensity and analyst coverage, as defined in Table 1, and their coefficients are reported in full in Appendix 3. Standard errors clustered at the firm level in parentheses. *** $p < 0.001$. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

The three pillars respond very differently, which supports the central claim of Hypothesis 1b that the composition of ESG effort changes. The environmental pillar carries a coefficient of 0.085, positive and marginally outside conventional significance at a p-value of 0.052. The social pillar declines by 0.180 and the governance pillar by 0.252, both significant at the one percent level. The two aggregate measures behave differently over these same observations. The weighted average score falls by 0.115, while the industry-adjusted score shows no

movement at all, at -0.032. Set against pillar estimates that run from +0.085 to -0.252, both aggregates understate the reallocation taking place beneath them, and the industry-adjusted measure hides it entirely.

The direction of the ranking, however, is the reverse of what Hypothesis 1b predicted. The prediction was that governance would prove most resilient and that the environmental and social dimensions would bear the reduction. In these estimates, governance falls by the most, social falls by an intermediate amount and environmental does not fall at all. The broader proposition that policy tightening reallocates ESG effort across dimensions is therefore supported, while the predicted ordering is not. Whether that ordering reflects firm behaviour or the construction of the MSCI ratings is examined in the robustness analysis.

Because the policy variable takes a single value per year and the sample is concentrated in the later years of the period, the possibility that these coefficients capture a secular trend rather than a response to policy requires attention. Adding a linear time trend to each pillar regression leaves the governance coefficient at -0.236 and the social coefficient at -0.126, both still significant at the one percent level, while the environmental coefficient little changed at 0.077. The trend term is itself significant for the social pillar at -0.099, indicating a genuine downward drift in social scores over the period that accounts for roughly a third of the raw estimate. Still, it is small and insignificant for governance at -0.030. The correlation between the policy variable and the calendar year is 0.199. The pillar results are therefore not an artifact of trending scores, although the social coefficient should be read as partly reflecting a trend.

6.4 Board independence as a moderator

Hypothesis 2 holds that board independence dampens the relationship between the policy stance and ESG performance. Columns 1 and 2 of Table 6 report the two specifications set out in the model section. In the specification with year fixed effects, which carries the identifying content, the interaction between the policy change and centred board independence is 0.0032 with a standard error of 0.0049. The estimate is in the predicted direction, since a positive interaction implies a smaller decline for firms with more independent boards, but it is roughly two-thirds of its standard error. The specification with industry fixed effects only produces a nearly identical estimate of 0.0029, so the finding is not sensitive to the treatment of aggregate conditions.

The economic magnitude confirms that the estimate is negligible rather than merely imprecise. The interquartile range of board independence in the sample is 9.85 percentage points.

Multiplying the interaction coefficient by this range implies that a firm at the third quartile of board independence responds to a one percentage point tightening by 0.031 points of ESG score less than a firm at the first quartile. Against a dependent variable with a standard deviation of 2.23, that difference is close to zero on any interpretation.

The main effect of board independence is also small and insignificant, at 0.012 with a standard error of 0.012 in the specification with year fixed effects. In this sample, board independence is associated neither with the level of ESG performance nor with its sensitivity to monetary policy. Hypothesis 2 is not supported.

6.5 The sustainability committee

The presence of a board-level sustainability committee, measured at the firm's first appearance in the panel, is examined as a secondary moderator. Column 3 of Table 6 reports the estimation. Firms that had such a committee when they entered the sample score 0.603 points higher on the industry-adjusted ESG scale, significant at the five percent level, which is approximately 0.27 of a standard deviation. The interaction with the policy change is -0.045 with a standard error of 0.119 and is not significant.

Table 6. Board independence and the sustainability committee as moderators

	(1) Independence, industry FE	(2) Independence, industry + year FE	(3) Committee, industry + year FE
ΔSHADOW	-0.0613	-	-
	(0.0460)	-	-
INDEP (centred)	0.0136	0.0121	-
	(0.0121)	(0.0124)	-
ΔSHADOW × INDEP	0.0029	0.0032	-
	(0.0050)	(0.0049)	-
COMMITTEE	-	-	0.6026*
	-	-	(0.2664)
ΔSHADOW × COMMITTEE	-	-	-0.0449
	-	-	(0.1188)
Controls	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
Year FE	No	Yes	Yes
Observations	1,601	1,601	1,610
Firms	398	398	398
Adjusted R^2	0.114	0.123	0.116

Note. Controls are firm size, return on assets, leverage, earnings volatility, research intensity and analyst coverage, as defined in Table 1, and their coefficients are reported in full in Appendix 4. The dependent variable is the MSCI industry-adjusted aggregate ESG score throughout. Cells marked - indicate coefficients that are not estimated: the main effect of the policy variable is absorbed by the year fixed effects in columns 2 and 3, and each moderator appears only in the columns in which it is tested. Standard errors clustered at the firm level in parentheses. * $p < 0.05$. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

The pattern is therefore that a governance structure dedicated to sustainability is associated with the level of ESG performance but not with its response to the monetary environment. It is worth noting that 19.6 percent of firms had a committee at their first appearance in the panel,

against 37.1 percent of firm-years across the sample as a whole, which reflects widespread adoption of such committees over the period rather than a discrepancy between the two figures.

6.6 Fixed and random effects

The Hausman test returns conflicting results depending on the treatment of the errors. The classical formulation rejects random effects, with a chi-squared statistic of 17.44 and a p-value of 0.015, while the cluster-robust version does not, at 7.62 and 0.368. Since the clustering used throughout explicitly rejects the assumption the classical test relies on, the robust version is the more appropriate.

The empirical case for fixed effects is therefore weaker than the classical test suggests, and the choice rests on the theoretical argument set out in the model section. The failure to reject should be read as an absence of evidence against random effects rather than as evidence in their favour, given a panel in which 140 of 416 firms are observed only once.

7. **Robustness**

Five checks examine whether the results depend on decisions taken in constructing the sample and the specification. Four vary the specification and are reported in Table 7. The fifth replaces the rating provider and, since it concerns the pillar scores only, is reported separately in Table 8.

Table 7. Robustness checks

	Baseline	Mean-to-mean change	2015 excluded	Independence next year	Firm fixed effects
H1a, unadjusted score	-0.1147***	-0.1703***	-0.1250***	-	-0.0928***
	(0.0235)	(0.0345)	(0.0337)	-	(0.0172)
H1a, adjusted score	-0.0554	-0.0605	-0.0160	-	-0.0319
	(0.0456)	(0.0611)	(0.0642)	-	(0.0405)
Environmental pillar	0.0851	0.0652	0.1197*	-	0.0880*
	(0.0438)	(0.0581)	(0.0572)	-	(0.0371)
Social pillar	-0.1802***	-0.2443***	-0.2211***	-	-0.1343***
	(0.0380)	(0.0493)	(0.0501)	-	(0.0324)
Governance pillar	-0.2519***	-0.3858***	-0.2089***	-	-0.2787***
	(0.0356)	(0.0537)	(0.0469)	-	(0.0386)
H2 interaction	0.0032	0.0084	0.0060	0.0086	-0.0029
	(0.0049)	(0.0064)	(0.0071)	(0.0051)	(0.0039)

Note. Each cell reports the coefficient on the policy variable, or on its interaction with board independence in the final row, with the clustered standard error below. Controls, fixed effects and clustering are as in the corresponding main specification. The board independence check alters only the interaction model, so the remaining rows are unchanged from the baseline and are not repeated. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

The first check replaces the December to December change in the shadow rate with the change in its annual mean. This measure smooths the within-year path of policy and is less sensitive to the timing of moves late in the calendar year. Every coefficient retains its sign, and the two pillar results strengthen considerably, governance moving from -0.252 to -0.386 and the unadjusted aggregate from -0.115 to -0.170. The environmental coefficient loses the significance it had. The conclusions are unaffected.

The second check excludes 2015, the year in which the exit from the effective lower bound produces a recorded change of 2.68 percentage points that is partly an artifact of how the

shadow rate is constructed. The unadjusted aggregate and the social and governance pillars all remain significant at the one percent level or better, and the social coefficient strengthens. The governance coefficient falls from -0.252 to -0.209, indicating that 2015 contributes to but does not generate the result. The coefficient on the adjusted aggregate falls close to zero, which reinforces the reading that this measure carries little of the policy signal.

The third check addresses the Bloomberg reporting convention, under which governance fields carry a date of 31 October rather than the firm's fiscal year end. Substituting the following year's value of board independence leaves the interaction at 0.0086 with a standard error of 0.0051, marginally larger than the baseline but still not significant at the five percent level. The conclusion regarding Hypothesis 2 does not change.

The fourth check replaces industry with firm fixed effects, so that identification comes only from movements within firms over time. This is the most demanding specification and the least well suited to the panel, since 140 of the 416 firms are observed in a single year and drop out entirely, leaving 276 firms and 1,527 observations for the aggregate model. The unadjusted aggregate remains significant at -0.093, and both the social and governance pillars remain significant at the one percent level, with governance slightly stronger than in the baseline. The environmental coefficient becomes significant and positive. The interaction changes sign to -0.0029 and remains insignificant.

A fifth check replaces the MSCI pillar scores with those of Bloomberg. Because Bloomberg covers 695 firm-years against MSCI's 1,192, the comparison is restricted to the 650 firm-years for which both providers report all three pillars, so that the estimates differ only by provider. The two give opposite accounts. On Bloomberg, the environmental and social pillars fall significantly, while governance does not. On MSCI, over the same firms and years, governance falls and environmental does not. The reason is that the providers barely agree at pillar level, correlating at 0.055, 0.221 and 0.299 for the three pillars respectively. The social pillar declines significantly under both, and that is the only pillar result common to them.

Table 8. Monetary policy and the three ESG pillars under two rating providers

	Environmental		Social		Governance	
	MSCI	Bloomberg	MSCI	Bloomberg	MSCI	Bloomberg
ΔSHADOW	0.0070	-0.2014***	-0.0825*	-0.1847***	-0.2018***	-0.0246
	(0.0525)	(0.0436)	(0.0391)	(0.0459)	(0.0309)	(0.0174)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	650	650	650	650	650	650
Adjusted R^2	0.259	0.421	0.285	0.195	0.186	0.179
Correlation		0.055		0.221		0.299

Note. The dependent variables are the environmental, social and governance pillar scores of MSCI and of Bloomberg, each on a scale from 0 to 10. ΔSHADOW is the December to December change in the Wu-Xia shadow federal funds rate, in percentage points, so that a positive value denotes a tightening of the monetary policy stance. Controls are firm size, return on assets, leverage, earnings volatility, research intensity and analyst coverage, as defined in Table 1. The sample is restricted to the 650 firm-years for which both providers report all three pillars, so that the two sets of estimates differ only by provider. The final row reports the pairwise correlation between the two providers' scores for the same pillar over these observations. Standard errors clustered at the firm level are reported in parentheses. * $p < 0.05$, *** $p < 0.001$. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

Taken together, the checks show that the two central findings are not fragile. The association between the policy stance and the unadjusted aggregate score is significant at the one percent level or better in all four variants, and the null on the industry-adjusted aggregate and on the board independence interaction are likewise stable. The social pillar declines significantly in every check and under both rating providers, which makes it the most robust of the pillar results. The governance and environmental results are less settled. Governance is significant in all four MSCI variants but not under Bloomberg. The environmental estimate is insignificant in the baseline, positive and significant when 2015 is excluded and under firm fixed effects, and negative and significant under Bloomberg. No reading of the environmental pillar is advanced here.

8. Discussion

8.1 What the estimates show

A tightening of the monetary policy stance is associated with a reduction in aggregate ESG performance on the unadjusted score, at -0.115 and significant at the one percent level, but not on the industry-adjusted score, where the estimate is indistinguishable from zero on either sample. The three pillars move by different amounts under both rating providers, and the social pillar declines under both, while the ordering of the three reverses when the provider changes. Neither board independence nor a sustainability committee dampens the response to policy. The committee is associated with a higher level of ESG performance on the industry-adjusted score, at 0.603 points, but not on the unadjusted one.

8.2 The industry adjustment and interpretation of Hypothesis 1a

The support for Hypothesis 1a rests on the measure adopted. The industry-adjusted score expresses a firm's rating relative to the range observed among its industry peers in the same year, so that movements shared within a sector are largely removed, while the unadjusted score retains them. Over identical firms and years, the association is -0.115 and significant on the unadjusted measure, against -0.032 and insignificant on the adjusted one, a difference of roughly a factor of seven in standard deviation terms. The two correlate at 0.75, and the adjusted score is more dispersed, with a standard deviation of 2.21 against 1.07, so a difference in units would make the coefficient of the adjusted score larger instead of smaller. Neither unrelated constructs nor scale can therefore account for the gap.

Two readings are possible. The first is that the retreat is real and largely common within sectors, since a change in the policy stance has an impact on the cost of finance for every firm at once, in which case the adjustment removes the variation the hypothesis concerns. The second is that the association reflects industry-specific dynamics, meaning sectoral movements coincide with the policy cycle without originating in it, in which case the adjustment removes a confound rather than the effect of interest.

The design cannot distinguish between them, since separating the two would require variation in the policy stance independent of sectoral conditions, which the aggregate specification does not supply. The support for Hypothesis 1a is therefore conditional on the measure, and this qualification should be carried into any use made of the result. More generally, peer-relative ratings are appropriate for questions about determinants that vary within an industry, and less for influences that reach entire sectors at once.

8.3 The pillar results and the divergence between providers

Hypothesis 1b made two claims. The first claim, that a tightening changes the composition of ESG effort rather than reducing it uniformly, is supported and does not depend on the rating provider: under both MSCI and Bloomberg, the pillars move by different amounts, and under both the social pillar declines.

The aggregate industry-adjusted score over the same observations shows no movement at all. This is the pattern Sakunasingha et al. (2018) document during the Great Recession: a composite rating cannot distinguish a firm that did not respond from a firm whose responses cancelled, and that component-level evidence is therefore required to interpret a null result rather than an extension of the analysis.

The second claim, that governance would prove the most resilient, cannot be assessed with these data. On MSCI, the ranking is the reverse of the prediction. On Bloomberg, over the same firms and years, the prediction holds. The ordering of the pillars is therefore a property of the rating provider rather than of firm behaviour, and no interpretation of why one pillar falls more than another can be sustained when a second provider reverses the order. Disagreement of this kind is documented rather than unexpected: Thewissen et al. (2026) construct a measure of rating disagreement from exactly this divergence between providers. A related divergence appears in Bauer et al. (2024), who find the differential response of green and brown stocks to monetary policy significant when greenness is measured by carbon emissions but not when it is measured by ESG data provider scores.

Plausible behavioural accounts of a pillar ranking can be constructed, and several are consistent with one provider or the other. Environmental commitments are frequently embedded in physical capital, in regulatory obligations and in multi-year public targets, which would make them slow to reverse. Social programmes correspond more closely to discretionary operating expenditure in training, community investment and supply chain management, and Hong et al. (2012) find that responsibility spending responds more to financial slack than capital expenditure or research and development. Governance ratings incorporate executive compensation arrangements and long-term contractual commitments, whose cost is itself sensitive to the rate of interest. Each of these is reasonable, and none can be settled here, because the evidence required to distinguish them is precisely the pillar ranking on which the two providers do not agree.

The conclusion is methodological. Pillar-level claims about which dimensions of conduct firms protect under pressure rely on an ordering that a change of provider may reverse, and require corroboration across providers before they can be advanced.

8.4 Why board independence shows no moderating effect

The absence of any moderating effect of board independence is the clearest null in the analysis. The interaction is estimated at 0.0032 against a standard error of 0.0049, and the implied difference in response between firms at the first and third quartiles of board independence is 0.03 points of ESG score for a one percentage point tightening, which is negligible under any interpretation. On the unadjusted score, the same interaction is 0.002 with a standard error of 0.0026, so the null does not depend on the measure.

The most immediate explanation lies in the composition of the sample. Board independence averages 85.8 percent with half the firms falling between 81.8 and 91.7 percent, a concentration that follows from the listing requirements applying to large United States firms over this period. A moderator with little variation cannot produce much differential response, and the test is correspondingly weak even though the estimate is precise.

A second explanation concerns the construct rather than its variation. The hypothesis rested on the argument that independent directors bring a longer horizon and a broader conception of the firm's obligations, and are therefore less willing to sacrifice ESG commitments under financial pressure. What the measure captures, however, is the absence of an employment or commercial relationship with the firm, which is a statement about a director's formal position rather than about their orientation towards stakeholders. Independence in this sense is compatible with a strictly shareholder-oriented view of the director's role, which is what agency theory advances, and the empirical literature on independence and social performance is itself mixed, with negative and null associations reported alongside positive ones (Ortas et al., 2017).

The result for the sustainability committee points the same way, and adds a further instance of the measurement problem set out above. On the industry-adjusted score, a committee is associated with 0.6 points of higher ESG performance, but on the unadjusted score, the same estimate is 0.01 with a standard error of 0.138. What is common to both measures is the absence of any moderating effect: committee firms do not respond differently to a change in the policy stance. Governance arrangements therefore appear to mark the level of a firm's ESG commitment, on some measures at least, but do not protect against macroeconomic pressure. Baraibar-Diez and Odriozola (2019) report the same level association and read it as the effect

of the committee. However, the present design cannot separate that reading from the possibility that firms establish such committees because they already prioritise these matters.

Because committee status is fixed at the value recorded when a firm first enters the panel, a firm that established one later is treated as having none throughout. Measurement error of this kind biases an estimated interaction towards zero, so the absence of a moderating effect is consistent with an effect this design cannot detect.

Taken together, however, the analysis produces several null estimates, but it does not establish that monetary policy has no impact on corporate ESG behaviour. The unadjusted aggregate score and two of the three pillars all move significantly with the policy stance, so the underlying relationship is present in the data. What the nulls establish is narrower and more specific: that the relationship is not detectable in a peer-benchmarked aggregate measure, and that board independence does not condition it.

Nor do the nulls establish that the theoretical arguments motivating the hypotheses are wrong. The design tests those arguments through particular measures on a particular sample, and a null is equally consistent with the argument being correct but the measure being unsuited to detecting it.

8.5 Limitations

The measure of monetary policy captures the stance of policy rather than an identified exogenous shock. Policy tightens when economic conditions are strong, and the estimates therefore cannot separate the effect of policy from the effect of the conditions that prompted it. This limitation cuts in a specific direction: since firms have more discretionary resources for ESG effort in good times, and policy tightens in good times, the omitted business cycle should bias the estimated coefficient upward, that is, against the negative sign the hypotheses predict. The negative estimates reported here are conservative for that reason.

The policy variable takes a single value in each of the fourteen years of the sample. Standard errors clustered at the firm level allow for correlation within firms over time but not for correlation across firms within a year, and the significance levels attached to the policy coefficient are therefore likely to be overstated. The panel is markedly unbalanced. Of 416 firms, 140 appear in a single year, and the median firm is observed three times. Firms with longer reporting histories are considerably larger, with median assets of 14.1 billion against 6.2 billion for firms appearing once. Coverage is also thin at the ends of the period, with 25 observations in 2007 and 63 in 2020 against 262 in 2019.

The sample consists of large United States firms that publish standalone ESG reports. The findings should not be extended to smaller firms, to firms that do not report, or to other institutional settings without further evidence, particularly given that the argument advanced above about board independence turns on characteristics specific to large United States boards.

Two data limitations are worth recording. Board independence could not be obtained for 66 firm-years, so the tests of Hypothesis 2 rest on 1,601 of 1,667 observations. The MSCI industry-adjusted score appears in two of the source files with values that do not always agree, correlating at 0.92 with a median difference of zero but with individual discrepancies in both directions, which is consistent with the two extracts having been taken at different times.

8.6 Implications and directions for further work

For boards, neither a high proportion of independent directors nor a dedicated sustainability committee alters how a firm's ESG performance responds to a tightening of the policy stance, and the estimates are precise enough to exclude a difference of any economic consequence. Boards that wish to protect these programmes through periods of expensive finance should look to protected budgets and multi-year commitments rather than to composition alone. For investors and analysts using ESG ratings as inputs, the divergence documented here matters directly: a portfolio screened on one provider's pillar scores is not the same portfolio as one screened on another's.

Whether the response to monetary policy differs across sectors is a natural question that the present design does not address, and one that a specification allowing the policy coefficient to vary by industry could answer directly.

Decomposing each pillar into its underlying key issues, and doing so for more than one rating provider, would show whether the divergence documented here originates in the issues each provider selects, or in how it scores them.

Finally, replacing the shadow rate with an identified monetary policy surprise, derived from high-frequency movements in interest rate futures around policy announcements, would allow a causal rather than an associational claim.

9. Conclusion

This thesis asked whether a tightening of the monetary policy stance reduces corporate ESG performance, whether any reduction falls evenly across the environmental, social and governance pillars, and whether board independence conditions the response. It answered those

questions on an unbalanced panel of 1,667 firm-years drawn from 416 S&P 500 firms between 2007 and 2020, using the December to December change in the shadow rate to measure the stance.

A tightening of one percentage point is associated with a fall of 0.115 points in the MSCI weighted average score, roughly a tenth of its standard deviation and significant at the one percent level. The same specification applied to the industry-adjusted score returns an estimate indistinguishable from zero. Across the pillars, the social and governance dimensions decline while the environmental dimension does not, but the ordering reverses when the same firms and years are rated by Bloomberg, and only the decline in the social pillar survives a change of provider. Board independence exhibits no moderating effect, and the estimate is precise enough to exclude one of any economic consequence. A sustainability committee is associated with a higher level of ESG performance on one measure and with nothing on the other, and with no difference in response to policy on either.

Three of these four findings turn on how ESG performance is measured rather than on how firms behaved. Whether a relationship with monetary policy appears at all depends on whether the rating is benchmarked against industry peers; which pillar firms protect depends on which provider is asked; whether governance arrangements are associated with stronger ESG performance depends on the same choice of aggregate. This is the principal contribution of the analysis, and it is a caution rather than a result. Behavioural inferences drawn from vendor ratings, and particularly from their components, rest on constructs on which the vendors themselves do not agree.

What the evidence does support is narrow and worth stating plainly. Firms subject to a tightening of the policy stance record lower ESG performance in absolute terms, and they do not reduce their effort uniformly. Because the shadow rate measures the stance of policy rather than an identified shock, these are associations and not causal effects. Policy tightens when conditions are strong, and firms have more to spend on discretionary activity when conditions are strong, so the omitted business cycle works against the negative association reported here. The estimates are conservative for that reason.

Bibliography

Al Amosh, H., & Khatib, S. F. A. (2022). Ownership structure and environmental, social and governance performance disclosure: The moderating role of the board independence. *Journal of Business and Socio-Economic Development,* 2(1), 49-66. https://doi.org/10.1108/JBSED-07-2021-0094

Allman, E., & Won, J. (2021). The Effect of ESG Disclosure on Corporate Investment Efficiency. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3816592

Anderl, C., & Caporale, G. M. (2023). Shadow rates as a measure of the monetary policy stance: Some international evidence. *Scottish Journal of Political Economy*, 70(5), 399-422. https://doi.org/10.1111/sjpe.12343

Bancel, F., & Glavas, D. (2022). Ownership and sustainability: The type of shareholder matters. *Bankers, Markets & Investors*, *169*(2), 31-38. https://doi.org/10.54695/bmi.169.0031

Baraibar-Diez, E., & Odriozola, M. D. (2019). CSR Committees and Their Effect on ESG Performance in UK, France, Germany, and Spain. *Sustainability*, 11(18), 5077. https://doi.org/10.3390/su11185077

Bauer, M. D., Offner, E. A., & Rudebusch, G. D. (2024). *Green stocks and monetary policy shocks: Evidence from Europe* (Working Paper No. 215). IMFS Working Paper Series. https://www.econstor.eu/handle/10419/308032

Bernanke, B. S., & Gertler, M. (1995). Inside the Black Box: The Credit Channel of Monetary Policy Transmission. *Journal of Economic Perspectives*, *9*(4), 27-48. https://doi.org/10.1257/jep.9.4.27

Black, F. (1995). Interest Rates as Options. *The Journal of Finance*, 50(5), 1371-1376. https://doi.org/10.1111/j.1540-6261.1995.tb05182.x

Cloyne, J., Ferreira, C., Froemel, M., & Surico, P. (2019). *Monetary Policy, Corporate Finance and Investment* (SSRN Scholarly Paper No. 3361718). Social Science Research Network. https://doi.org/10.2139/ssrn.3361718

Coase, R. H. (1937). The Nature of the Firm. *Economica*, 4(16), 386-405. https://doi.org/10.1111/j.1468-0335.1937.tb00002.x

De La Horra, L. P., Perote, J., & De La Fuente, G. (2021). Monetary policy and corporate investment: A panel-data analysis of transmission mechanisms in contexts of high uncertainty. *International Review of Economics & Finance*, 75, 609-624. https://doi.org/10.1016/j.iref.2021.04.035

Eliwa, Y., Aboud, A., & Saleh, A. (2021). ESG practices and the cost of debt: Evidence from EU countries. *Critical Perspectives on Accounting, SI: Accounting and Social Impact (Part I)*, *79*, 102097. https://doi.org/10.1016/j.cpa.2019.102097

Ellili, N. O. D. (2023). Impact of corporate governance on environmental, social, and governance disclosure: Any difference between financial and non-financial companies?

Corporate Social Responsibility and Environmental Management, 30(2), 858-873. https://doi.org/10.1002/csr.2393

Elmghaamez, I. K., Nwachukwu, J., & Ntim, C. G. (2024). ESG disclosure and financial performance of multinational enterprises: The moderating effect of board standing committees. *International Journal of Finance & Economics*, 29(3), 3593-3638. https://doi.org/10.1002/ijfe.2846

Freeman, R. E. (1984). *Strategic Management: A Stakeholder Approach*. Pitman.

Friedman, M. (1970, September 13). A Friedman doctrine - The Social Responsibility of Business Is to Increase Its Profits. *The New York Times*. https://www.nytimes.com/1970/09/13/archives/a-friedman-doctrine-the-social-responsibility-of-business-is-to.html

Harrison, J. S., & Berman, S. L. (2016). Corporate Social Performance and Economic Cycles. *Journal of Business Ethics*, *138*(2), 279-294. https://doi.org/10.1007/s10551-015-2646-9

Hong, H., Kubik, J. D., & Scheinkman, J. A. (2012). Financial Constraints on Corporate Goodness (Working Paper No. 18476). National Bureau of Economic Research. https://doi.org/10.3386/w18476

Jankensgård, H. (2015). The Relationship between Voluntary Disclosure, External Financing and Financial Status. *Journal of Business Finance & Accounting*, 42(7-8), 860-884. https://doi.org/10.1111/jbfa.12120

Jizi, M. (2017). The Influence of Board Composition on Sustainable Development Disclosure. *Business Strategy and the Environment*, 26(5), 640-655. https://doi.org/10.1002/bse.1943

Kansoy, F., & Stasiulaitis, D. (2025). *Green Shields: The Role of ESG in Uncertain Time* (arXiv:2506.02143). arXiv. https://doi.org/10.48550/arXiv.2506.02143

Krippner, L. (2012). Measuring the Stance of Monetary Policy in Zero Lower Bound Environments (SSRN Scholarly Paper No. 2112758). Social Science Research Network. https://doi.org/10.2139/ssrn.2112758

Lagasio, V., & Cucari, N. (2019). Corporate governance and environmental social governance disclosure: A meta-analytical review. *Corporate Social Responsibility and Environmental Management*, 26(4), 701-711. https://doi.org/10.1002/csr.1716

Liao, L., Luo, L., & Tang, Q. (2015). Gender diversity, board independence, environmental committee and greenhouse gas disclosure. *The British Accounting Review*, 47(4), 409-424. https://doi.org/10.1016/j.bar.2014.01.002

Mohammad, K. U., Hussain, M. M., & Adnan, N. U. H. (2025). Risk-taking channel of monetary policy and the role of ESG and political instability: Evidence from South Asia. *International Economics and Economic Policy*, *22*(3), 49. https://doi.org/10.1007/s10368-025-00674-8

Ortas, E., Álvarez, I., & Zubeltzu, E. (2017). Firms' Board Independence and Corporate Social Performance: A Meta-Analysis. *Sustainability*, 9(6), 1006. https://doi.org/10.3390/su9061006

Ottonello, P., & Winberry, T. (2020). Financial Heterogeneity and the Investment Channel of Monetary Policy. *Econometrica*, *88*(6), 2473-2502. https://doi.org/10.3982/ECTA15949

Sakunasingha, B., Jiraporn, P., & Uyar, A. (2018). Which CSR activities are more consequential? Evidence from the Great Recession. *Finance Research Letters*, *27*, 161-168. https://doi.org/10.1016/j.frl.2018.02.003

Salzmann, E. (2025). *ESG Indices and Macroeconomic Factors: A European Panel Data Analysis* (SSRN Scholarly Paper No. 5291809). Social Science Research Network. https://doi.org/10.2139/ssrn.5291809

Thewissen, J., Shrestha, P., Arslan-Ayaydin, Ö., & Beibei, Y. (2026). Topics That Matter: ESG Reports and Information Asymmetry. *Journal of Financial Research*, 1-54. https://doi.org/10.1111/jfir.70053

Tsang, A., Frost, T., & Cao, H. (2023). Environmental, Social, and Governance (ESG) disclosure: A literature review. *The British Accounting Review*, *55*(1), 101149. https://doi.org/10.1016/j.bar.2022.101149

Wu, J. C., & Xia, F. D. (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking*, 48(2-3), 253-291. https://doi.org/10.1111/jmcb.12300

Yao, S., Hong, Y., & Lin, C.-M. (2019). Environmental Information Disclosure and Financial Constraint. *Asia-Pacific Journal of Financial Studies*, 48(5), 666-689. https://doi.org/10.1111/ajfs.12277

Appendices

Appendix 1. Monetary policy and the industry-adjusted ESG score on the full panel, full control coefficients

	Industry-adjusted score
ΔSHADOW	-0.0554
	(0.0456)
SIZE	0.2273*
	(0.0939)
ROA	0.0363*
	(0.0145)
LEV	-0.0021
	(0.0074)
VOLAT	-0.0852*
	(0.0330)
RD	0.0580*
	(0.0280)
NoA	-0.1642
	(0.2287)
Industry FE	Yes
Year FE	No
Observations	1,667
Firms	416
Adjusted R^2	0.097

Note. The table reports the estimation of equation (1) on the full panel of 1,667 firm-years, with all control coefficients shown. The dependent variable is the MSCI industry-adjusted score. ΔSHADOW is the December to December change in the shadow federal funds rate, in percentage points, so that a positive value denotes a tightening of the monetary policy stance. Industry fixed effects are included; year fixed effects are not, since the

policy variable takes a single value per year. Standard errors clustered at the firm level in parentheses. * p < 0.05, ** p < 0.01, *** p < 0.001. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

Appendix 2. Adjusted and unadjusted ESG measures compared, full control coefficients

	Unadjusted score	Industry-adjusted score
ΔSHADOW	-0.1147***	-0.0323
	(0.0235)	(0.0566)
SIZE	0.0952	0.1997*
	(0.0514)	(0.1012)
ROA	0.0212**	0.0357*
	(0.0078)	(0.0163)
LEV	0.0011	0.0018
	(0.0037)	(0.0084)
VOLAT	-0.0527**	-0.1277**
	(0.0180)	(0.0415)
RD	0.0080	0.0707*
	(0.0158)	(0.0355)
NoA	-0.0004	0.0939
	(0.1603)	(0.2722)
Industry FE	Yes	Yes
Observations	1,192	1,192
Adjusted R^2	0.208	0.152

Note. The table compares the estimation of equation (1) with the MSCI weighted average key issue score and with the MSCI industry-adjusted score as dependent variables, over the same firm-years, with all control coefficients shown. ΔSHADOW is the December to December change in the shadow federal funds rate, in percentage points, so that a positive value denotes a tightening of the monetary policy stance. Industry fixed effects are included; year fixed effects are not, since the policy variable takes a single value per year. Standard errors clustered at the firm level in parentheses. * p < 0.05, ** p < 0.01, *** p < 0.001. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

Appendix 3. Monetary policy and the three ESG pillars, full control coefficients

	Aggregate (weighted average)	Aggregate (industry-adjusted)	Environmental	Social	Governance
ΔSHADOW	-0.1147***	-0.0323	0.0851	-0.1802***	-0.2519***
	(0.0235)	(0.0566)	(0.0438)	(0.0379)	(0.0356)
SIZE	0.0952	0.1997*	0.4514***	0.0448	-0.2347***
	(0.0514)	(0.1012)	(0.0873)	(0.0692)	(0.0654)
ROA	0.0212**	0.0357*	0.0387*	0.0231	-0.0006
	(0.0078)	(0.0163)	(0.0183)	(0.0128)	(0.0086)
LEV	0.0011	0.0018	0.0154**	-0.0049	-0.0106**
	(0.0037)	(0.0084)	(0.0059)	(0.0061)	(0.0036)
VOLAT	-0.0527**	-0.1277**	-0.0385	-0.0483	-0.0346
	(0.0180)	(0.0415)	(0.0361)	(0.0280)	(0.0210)
RD	0.0080	0.0707*	-0.0188	0.0256	0.0089
	(0.0158)	(0.0355)	(0.0251)	(0.0216)	(0.0180)
NoA	-0.0004	0.0939	-0.3088	0.0314	0.1241
	(0.1603)	(0.2722)	(0.2564)	(0.2194)	(0.1543)
Industry FE	Yes	Yes	Yes	Yes	Yes
Observations	1,192	1,192	1,192	1,192	1,192
Adjusted R^2	0.208	0.152	0.206	0.193	0.134

Note. The table reports the estimation of equation (1) with the two aggregate ESG measures and with the MSCI environmental, social and governance pillar scores as dependent variables, with all control coefficients shown. ΔSHADOW is the December to December change in the shadow federal funds rate, in percentage points, so that a positive value denotes a tightening of the monetary policy stance. Industry fixed effects are included; year fixed

effects are not, since the policy variable takes a single value per year. Standard errors clustered at the firm level in parentheses. * p < 0.05, ** p < 0.01, *** p < 0.001. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.

Appendix 4. Board independence and the sustainability committee as moderators, full control coefficients

	(1) Independence, industry FE	(2) Independence, industry + year FE	(3) Committee, industry + year FE
ΔSHADOW	-0.0613	-	-
	(0.0460)	-	-
INDEP (centred)	0.0136	0.0121	-
	(0.0121)	(0.0124)	-
ΔSHADOW × INDEP	0.0029	0.0032	-
	(0.0050)	(0.0049)	-
COMMITTEE	-	-	0.6026*
	-	-	(0.2664)
ΔSHADOW × COMMITTEE	-	-	-0.0449
	-	-	(0.1188)
SIZE	0.2167*	0.2081*	0.1784
	(0.0992)	(0.1004)	(0.0996)
ROA	0.0350*	0.0338*	0.0355*
	(0.0150)	(0.0149)	(0.0149)
LEV	-0.0019	-0.0034	-0.0037
	(0.0077)	(0.0080)	(0.0078)
VOLAT	-0.0926**	-0.1070**	-0.1064**

	(1) Independence, industry FE	(2) Independence, industry + year FE	(3) Committee, industry + year FE
	(0.0352)	(0.0384)	(0.0384)
RD	0.0581*	0.0588*	0.0695*
	(0.0285)	(0.0288)	(0.0283)
NoA	-0.1205	-0.0735	-0.0587
	(0.2402)	(0.2463)	(0.2377)
Industry FE	Yes	Yes	Yes
Year FE	No	Yes	Yes
Observations	1,601	1,601	1,610
Firms	398	398	398
Adjusted R^2	0.114	0.123	0.116

Note. The dependent variable is the MSCI industry-adjusted aggregate ESG score throughout. Cells marked - indicate coefficients that are not estimated: the main effect of the policy variable is absorbed by the year fixed effects in columns 2 and 3, and each moderator appears only in the columns in which it is tested. Standard errors clustered at the firm level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Source: author's calculations on the estimation sample described in Section 4.1, using MSCI ESG ratings, Compustat, CRSP, I/B/E/S and Bloomberg data.