

ABSTRACT

Title of Dissertation: GREENWASHING, FIRM ESG STRATEGY, AND EMPLOYEE IMPACT

Nathan Barrymore, Doctor of Philosophy, 2022

Dissertation Directed By: Associate Professor Rachelle C. Sampson
Logistics, Business and Public Policy
Robert H. Smith School of Business
University of Maryland

This dissertation studies the causes and consequences of firms' environmental and social (ESG) actions, with a specific focus on employees. Essay 1 examines greenwashing: when firms present an overly positive view of their environmental and social outcomes. I ask how top managers and investors' ESG preferences influence companies' self-reported environmental and social policies, and their independently reported environmental and social outcomes. I find that managers' ESG preferences, as proxied using their language on earnings calls, correlate with both ESG policies and outcomes. However, investors' ESG preferences correlate with only policies and not outcomes, suggestive of greenwashing. I conclude that agency issues explain these divergent results.

Essays 2 and 3 ask how employees respond to firms' ESG outcomes and to firms' pay policies. Essay 2 explores the relationship between a firm's ESG outcomes and labor productivity. In two contexts, we find that ESG outcomes predict higher labor productivity, but only when there is sufficient information about firm behavior. In one study, the positive impact on labor productivity only exists for large firms. In another study, the positive relationship appears only after a government regulation requiring that firms disclose their carbon emissions. Essay 3 provides large scale evidence on the relationship between wages and employee attrition. We find that paying above median wages for a specific role decreases attrition rates, but only among low and middle wage workers in the US.

If stakeholder capitalism is to sustain and integrate into the US corporate system, the movement needs to be based on accurate assessments of environmental and social outcomes. These essays provide an advance in that direction, by using independently reported ESG data to examine how ESG issues impact firm strategy.

GREENWASHING, FIRM ESG STRATEGY, AND EMPLOYEE IMPACT

By

Nathan Barrymore

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
Of the requirements for the degree of
Doctor of Philosophy
2022

Advisory Committee:

Associate Professor Rachelle C. Sampson (Chair)
Associate Professor Cristian Dezsö
Professor Ginger Jin
Professor Lemma Senbet
Associate Professor David Waguespack

Acknowledgements

Thank you to my family for their support, and to my advisors for their patient guidance.

Contents

Introduction	1
Chapter 1: Green or Greenwashing?	4
Introduction.....	5
Firm Responses To Normative Pressure	10
Empirical Application	17
Results.....	27
Discussion.....	34
Tables and Figures.....	39
Chapter 2: ESG Performance and Labor Productivity.....	53
Introduction.....	54
Environmental, Social, and Financial Performance.....	56
ESG And Labor Productivity	58
Study 1: Distribution of ESG-Labor Productivity Relationship.....	61
Study 2: Impact of Revealed Environmental Performance on Labor Productivity.....	69
Implications & Conclusion.....	80
Tables and Figures.....	84
Chapter 3: Wages and Employee Attrition.....	96
Introduction.....	97
Background.....	99
Empirical Application	104
Results.....	110
Discussion.....	114
Tables and Figures.....	117
Appendix A: Non-Equilibrium Model Proof.....	123
Appendix B: Equilibrium Model Proof.....	130
Appendix C: Additional Tables from Chapter 1.....	136
Appendix E: Simulation Model of the Wage-Attrition Relationship.....	140
References.....	143

List of Tables

Table 1: Examples of ESG Related Articles from TruValue	39
Table 2: List of Exemplar Firms, by Green/Brown(washing) Category and Sector; Public US Firms.	40
Table 3: Exemplar Firms, Food and Household Sector	41
Table 4: Correlation Table	42
Table 5: Policy-Outcome Correlations by Sector-Year	42
Table 6: Sample	43
Table 7: Summary Statistics	44
Table 8: Correlation Table (4,032 Observations)	44
Table 9: ESG Stated Preferences, and ESG Policies & Outcomes: Using Word Count Measure	45
Table 10: Effect of Manager and Investor ESG Preferences on ESG Policies and Outcomes, Including Language Polarity and Subjectivity	46
Table 11: Managers' and Investors' ESG Stated Preferences, and ESG Policies & Outcomes, SASB Semantic Distance Measure.....	47
Table 12: ESG Stated Preferences, and ESG Policies & Outcomes: Using Investor Time Horizons	48
Table 13: Ratings Variance and Greenwashing	49
Table 14: Illustrative news articles from TruValue Labs dataset	88
Table 15: Sample Construction.....	89
Table 16: Industry Coverage	89
Table 17: Descriptive Statistics	90
Table 18: ESG and Labor Productivity	91
Table 19: UK Stufdy, Sample Stock Exchanges.....	92
Table 20: UK Study, Descriptive Statistics	93
Table 21: Labor productivity and excess emissions, by region-year.....	94
Table 22: Sector level regressions of labor productivity on excess emissions	94
Table 23: Difference in difference estimates of effect of UK regulation	95
Table 24: Example Rows of Data.....	117
Table 25: Sample Description	117
Table 26: Job Category Wage Distributions, July 2019	118
Table 27: Summary Statistics, January 2018 through June 2020	119
Table 28: Correlations	119
Table 29: Relationship between relative wages and yearly turnover percentage	120
Table 30: Wage Quintile Regressions, No Instrumentation.....	121
Table 31: Wage Quintile Regressions, Instrumented.....	121
Table 32: Job Market Size Quintiles, No Instrumentation	122
Table 33: Job Market Size Quintiles, Instrumented.....	122

List of Figures

Figure 1: Coefficient of ESG Policies On ESG Outcomes	50
Figure 2: Management’s ESG Language on Earnings Calls (SASB)	51
Figure 3: Shareholder’s ESG Orientation (Sustainable AUM Percentage)	51
Figure 4: Ratings Disagreement and Greenwashing.....	52
Figure 5: US Firm Estimates, by sector	84
Figure 6: US Firm Estimates, early sample period.....	84
Figure 7: US Firm Estimates, late sample period.....	85
Figure 8: Firm effects by firm size (assets and revenues)	85
Figure 9: Emissions intensity, high and low emitting firms, by region-year	86
Figure 10: Coefficient estimates of relationship between labor productivity and excess emissions, by region-year.....	86
Figure 11: Coefficient estimates of relationship between revenue and excess emissions, by region-year	87
Figure 12: Coefficient estimates of relationship between employees and excess emissions, by region-year	87
Figure 13: Earnings Call ESG Discussions by Sector	138
Figure 14: Earnings Call Environmental Discussions by Sector	138
Figure 15: Earnings Call Social Discussions by Sector.....	139
Figure 16: Synthetic Data Example of Simulation Model	142

Introduction

“Companies must benefit all of their stakeholders, including shareholders, employees, customers, and the communities in which they operate,” Larry Fink

“When a measure becomes a target, it ceases to be a good measure,” Marilyn Strathern

The three essays in this dissertation are linked around the question of how to evaluate stakeholder capitalism. A majority of the world’s population believes capitalism does more harm than good.¹ While capitalism is still viewed favorably in the United States, younger adults are losing faith in the system.² Employees are increasingly demanding that the companies they work for share their values, and take on a mission beyond maximizing profits.³ In response to these concerns, corporations in the United States and Europe, and organizations such as The World Economic Forum, have latched onto stakeholder capitalism as a guidepost.

This abstract stakeholder capitalism movement has found shared language within C-suites and boardrooms in the form of environmental, social and governance (ESG) issues. Firms communicate their ESG goals and actions as part of their corporate strategies. Institutional investors use quantified ESG metrics to allocate capital, much to the appeal of retail investors, who increasingly purchase investment products with ESG labels. In 2019 over one third of assets under management in the US claimed to consider ESG criteria in asset allocation.⁴

But has stakeholder capitalism resulted in substantive change? Almost all the largest public corporations in the US now employ corporate social responsibility (CSR) departments and produce annual CSR reports which outline the positive impact each corporation has on the environment and

¹ John M. 2020, January 20. Capitalism seen doing ‘more harm than good’ in global survey. Reuters.

² Manchester J. 2021, June 28. Majority of young adults in US hold negative view of capitalism: poll. The Hill.

³ Gaskell A. 2021. Employees Demand That We Become More Sustainable. Forbes.

⁴ Principles for Responsible Investment. 2019. Taking Stock: Sustainable Finance Policy Engagement and Policy Influence.

society. The role of CSR officer has been promoted to the C-Suite, and CEO and top executive remuneration is increasingly linked to ESG performance targets (Flammer, Hong, and Minor, 2019). Despite this growth, critics suggest that most ESG is greenwash – symbolic adherence to the movement without real change. Assessing this critique is a challenge, and research and practice around stakeholder capitalism suffers from both a hesitancy to recognize potential tradeoffs between financial and ESG goals (King and Pucker, 2021), and ESG data which is noisy and polluted by firm self-reporting (Berg, Kolbel, and Rigobon, 2020).

Across three essays, this dissertation analyzes the causes and consequences of firms' ESG strategies, while addresses these issues of greenwashing, ESG-financial tradeoffs, and a reliance on firm-reported data. Chapter 1 investigates greenwashing. I ask which actors drive substantive ESG strategies, and which actors drive symbolic strategies (greenwashing). I find that managers' ESG preferences predict substantive ESG adoption, while investors' ESG preferences drive symbolic ESG adoption, consistent with a principal-agent framework. In chapters 2 and 3, I ask how employees respond to firms' ESG outcomes and pay policies. I find that ESG outcomes predict employee-related financial outcomes, but only for certain types of firms and workers. Both chapters have implications for how managers can design ESG strategies in response to employee demands.

My findings suggest that ESG strategies depend on the quality of ESG information stakeholders receive. Chapter 1 outlines the ESG information asymmetry between managers and investors and concludes that firms greenwash for investors because of this asymmetry. Chapter 2 finds that mandatory reporting can create a link between ESG outcomes and employee productivity. Each of these chapters also addresses the problematic use of voluntary and unaudited firm reported ESG data. Chapters 1 and 2 use independent media reports of firms' ESG outcomes, while chapter 3 uses employee-reported information to examine worker attrition rates.

I also identify tradeoffs between ESG and financial outcomes. Chapter 2 finds that the correlation between ESG performance and labor productivity is positive for large firms, but negative for smaller firms in the US and Europe over the past decade. Chapter 3 finds that higher wages are correlated with reduced attrition, but not for all workers.

By examining components of the stakeholder capitalism movement, each essay that follows provides insight for managers around how to create ESG strategy. Chapter 1 suggests that managers who wish to differentiate their firms based on their ESG outcomes will have to find novel strategies that are prohibitively costly for ESG laggards to imitate. Chapter 2 provides evidence that managers need to better publicize ESG outcomes to employees if they intend to use ESG as a lever for employee productivity. Chapter 3 points to increasing wages as an effective tool to decrease employee attrition, but only for low and moderate wage workers. Each chapter also provides novel signals for ESG investors to use when evaluating companies for inclusion in ESG investment products. If stakeholder capitalism is to sustain and integrate into the US corporate system, the movement needs to be based on accurate assessments of environmental and social outcomes. These essays provide an advance in that direction.

Chapter 1: Green or Greenwashing?

How Manager and Investor Preferences Shape Firm Strategy.

Abstract

When faced with pressures from stakeholders, firms may respond by substantively changing their products and processes, or by symbolically responding without substantive change. In this paper, I theorize and empirically study how managers' and investors' preferences with regards to one such pressure – that for environmental and social (ESG) responsibility – causes firms to either make substantive changes that result in improved outcomes or to greenwash: adopt symbolic policies. Empirically, I find that managers' ESG preferences, as proxied using their language on earnings calls, correlate with both ESG policies and outcomes. However, investors' ESG preferences correlate with only policies and not outcomes, suggestive of greenwashing. These results are consistent with a principal-agent model in which managers greenwash due to information asymmetry. Finally, I show that my measure of greenwashing correlates with ESG ratings disagreement, providing practical insight for managers and investors. These results have strategic implications in situations in which investors must rely on third parties to verify firms' claims, as well as for the current effectiveness of investor-led stakeholder capitalism. Furthermore, this paper shows how symbolic management and decoupling can be caused by individual stakeholders' preferences.

Introduction

Firms are under increasing normative pressures from stakeholders, including pressures to improve their environmental and social outcomes (Marquis, Toffel, and Zhou, 2016). Firms can respond to these pressures by substantively enacting change within the organization, or by symbolically adopting policies that do not result in tangible outcomes (Durand, Hawn, and Ioannou, 2019). Because substantive change can be difficult and costly, firms sometimes symbolically adopt corporate governance structures, product portfolios, or ethical codes in response to pressures from customers, investors, regulatory bodies, social activists, and external evaluators (Fiss and Zajac, 2004; Kim and Jensen, 2011; Stevens *et al.*, 2005). Even CEO dismissals can be symbolic responses that aren't intended to substantively change firm outcomes (Gentry *et al.*, 2021).

Research has identified how institutional environments, activist pressure, firm capabilities, and issue salience affect a firm's decision to substantively or symbolically respond to normative pressures (Bundy, Shropshire, and Buchholtz, 2013; Durand *et al.*, 2019; Greenwood *et al.*, 2010; McDonnell and King, 2013). I extend this work by studying how preferences of specific actors within firms – top managers and investors – shape whether firms respond substantively or symbolically to normative pressures.

I study the question of symbolic responses in the context of firms' environmental and social actions. Customers, employees, retail investors and activists often pressure firms to improve their environmental and social (ESG) impacts (Burbano, 2016; Dyck *et al.*, 2019; Hartzmark and Sussman, 2019; McDonnell, King, and Soule, 2015). Managers and institutional investors have responded by stating their commitment to improve the ESG outcomes produced by the firms they manage and own (Chava, Du, and Malakar, 2021; Gond and Piani, 2013). However, assessing firms' ESG outcomes, and thus whether firms are making substantive changes or simply taking symbolic stances, remains a challenge because of the potential for firms to greenwash, that is, to take actions

that, intentionally or not, result in a deceptively positive impression of their environmental and social outcomes (Allen, 2018; Berg *et al.*, 2020; Chatterji *et al.*, 2016; Kim and Lyon, 2015; Lyon and Maxwell, 2011; Marquis and Qian, 2014; Terrachoice, 2010). This paper asks the question: when do firms greenwash? Specifically, I ask, how do managers' and investors' ESG preferences correlate with greenwashing, as measured by the gap between a firm's ESG policies and its ESG outcomes.

This investigation into greenwashing is motivated by the claims of stakeholder capitalism emanating from CEOs and investment managers. One third of assets under management in the U.S. are invested in funds which state that they consider ESG information in their investment decisions, a dramatic increase over the past decade (Nason, 2020; Principles for Responsible Investment, 2019). Meanwhile, several hundred CEOs of the largest U.S. corporations have publicly stated that their firms have responsibilities to non-shareholding stakeholders, and these same CEOs are increasingly taking public stances on social issues (Business Roundtable, 2019; Chatterji and Toffel, 2018).

While firms claim they are doing good by society, environmental activists, employees and customers have accused firms of greenwashing (Lyon and Montgomery, 2015). Unlike financial accounting, ESG reporting over the past decade has largely been voluntary, non-standardized and unaudited. Though ESG ratings agencies and other organizations aim to provide unbiased assessments of firms' overall ESG outcomes, their ratings are, out of necessity, largely based on these unaudited and self-reported data (Peirce, 2019). This fertile ground for greenwashing prohibits investors from efficiently allocating capital according to their ESG preferences, provides poor signals to managers as to their actual ESG outcomes, and limits firms' ability to choose exchange partners (e.g., suppliers) that match their social or environmental preferences (Dai, Liang, and Ng, 2020). Furthermore, greenwashing threatens to impact any link present between ESG outcomes and financial performance because it may prohibit firms from credibly signaling their ESG outcomes to

interested stakeholders. Mining firms with higher ESG outcomes have lower cost of capital (Henisz and McGlinch, 2019) and firms can reduce online worker wage requirements through socially conscious messaging (Burbano, 2016). However, in situations where greenwashing is prevalent and difficult to observe, creditors, employees, and other stakeholders will not be able to distinguish firms with positive outcomes from those which are simply greenwashing.

Thus, stakeholders are left with a series of questions around the impact of the ESG movement. Which firms are greenwashing? How do CEOs and top managers influence a firm's greenwashing decision? Do investment companies claiming to have sustainable investments invest in firms with better ESG outcomes? I address these questions by using recently developed data and techniques to measure greenwashing and ESG language, and through a correlational study grounded in formal theory.

My first key finding is that ESG policies, gathered from Thompson Reuters ESG ratings, and ESG outcomes, gathered through news media reporting, are largely uncorrelated between 2008 and 2019. This calls into question the signal value in ESG policies listed in firms' CSR reports and suggests greenwashing may be a common firm strategy. I then examine how investor ESG preferences, as captured through mutual funds labeled as sustainable, and managers ESG preferences, as captured through their language on earnings calls, correlate with greenwashing. In a sample of the 1,500 largest public US firms, CEO preferences correlate positively with both ESG policies and ESG outcomes across firms. This is consistent with the idea that that when CEOs have a preference for better ESG outcomes, their firms put in place ESG policies, and importantly follow up these policies with actions which produce better ESG outcomes. However, investors' preferences correlate with firms' ESG policies, but not with firms' ESG outcomes. This ESG policies-outcomes gap is suggestive of greenwashing, consistent with the idea that firms with ESG-oriented investors,

but not ESG oriented management, on average, greenwash by putting in place symbolic ESG policies that do not result in improved ESG outcomes.

These results are consistent with a principal-agent model in which investors (the principals) are unable to align the incentives of firm managers (the agents) with their own preferences. Using a principal-agent framework, I model greenwashing⁵ as the misalignment between a firm's ESG policies (what they say they do) and their ESG outcomes (what they actually do). The model assumes that managers have access to ESG outcome information but present potentially selective ESG information to investors in the form of ESG policies on the books. Managers are thus incentivized to greenwash when the firm has ESG-oriented investors, which is consistent with my empirical findings. There are limits to greenwashing, however. Greenwashing can be curtailed by environmental activists, which attack greenwashing firms, as well by rational investors, who discount unverifiable ESG policies. Both limitations are discussed in the theory and discussion sections.

Finally, I ask whether greenwashing may contribute to ESG ratings disagreement. A major challenge for incorporating ESG issues into the financial markets is that ESG ratings are noisy. Different ratings agencies assign radically different ratings to the same firms (Berg *et al.*, 2021, 2020). I find that my empirical measure of greenwashing correlates with ESG ratings disagreement. This finding both adds credence to the use of the ESG policies-outcomes gap as a measure of greenwashing and suggests that investors concerned about greenwashing may want to commit additional resources to scrutinize the ESG outcomes of firms with noisy ESG ratings before adding these firms to their ESG investments products.

⁵ I also model brownwashing, when firms present an unrepresentatively negative view of their ESG outcomes (Kim and Lyon, 2015). Firms may brownwash to, for example, avoid being perceived as claiming to be green when they are not (Carlos and Lewis, 2018) or because investors perceive ESG outcomes as substitutes for financial returns (Kim and Lyon, 2015).

Theoretically, these results add to the literature on symbolic management and decoupling (Meyer and Rowan, 1977) by identifying firm-specific sources of symbolic management: the normative preferences of managers and investors. Past research has identified how institutional and firm-level characteristics are correlated with firms' decisions to adopt symbolic policies. The domain of environmental and social performance highlights the possibility that the normative preferences of firm insiders such as managers may determine whether a firm substantively or symbolically responds to normative pressures from outside stakeholders.

This study also has implications for managers, investors, and civil society actors concerned with how stakeholder capitalism manifests in firm strategy. For investors, these results point to potential biases in ESG data used over the past decade. While investment has flowed into firms with high ESG ratings by third party ratings agencies, these ratings have not correlated with ESG outcomes reported in the news. This suggests that investors must both consider new independently reported data when assessing firms and be wary of the potential for firms to greenwash their ESG policies.

For managers, these results are consistent with the idea that greenwashing has been a common strategy over the past decade. This suggests that, at a minimum, the average firm genuinely committed to ESG outcomes has not been able to credibly differentiate themselves from their greenwashing peers. Thus, managers wishing to differentiate their firms based on their ESG outcomes will need to find novel mechanisms to do so that are costly for greenwashing firms to imitate.

If firms can greenwash and investors are unable distinguish the greenwashing firms from those with positive ESG outcomes, then pressure from investors' capital allocation decisions alone cannot produce positive ESG outcomes for society. This study suggests that managers may be able to produce positive social outcomes, but it is doubtful that stakeholders ought to rely solely on managers' consciences to produce these positive outcomes for society. Thus, the results in the paper

are consistent with the idea that alternative solutions, such as investors' use of voice through shareholder proposals and proxy votes to change management, or, more broadly, governmental action in the form of mandatory, standardized, and audited disclosures, or regulation, may be more effective levers of change (Fancy, 2021; Flammer, 2015).

This paper proceeds as follows: I first synthesize existing theory on the drivers of greenwashing and then theorize as to how managers and investors impact ESG performance. While verbal theory is presented in the body of the paper, a model in the appendix formalizes the predictions in the paper. I then outline my measures and show empirical results. The paper ends with a discussion.

Firm Responses To Normative Pressure

Altria's 2018 Corporate Social Responsibility Report claimed that "the most important issue we face is the epidemic rise in youth vaping" (Altria, 2019). The same year, Altria bought a 35% stake in JUUL (Gray and Badkar, 2018), the e-cigarette company selling flavors such as Mango, Fruit, and Cucumber, along with light-hearted decorative covers for their products with smiley faces and other high-school friendly designs that helped teenagers disguise their nicotine habits from the eyes of adults.

While the Altria example presents one approach, firms can greenwash in a variety of ways, each preventing investors and other stakeholders from accurately evaluating firms' ESG outcomes. I define greenwashing as actions taken by managers in organizations that result in a view of the organization as being more environmental or socially responsible than justified given the environmental or social outcomes the firm produces. This definition follows recent calls to expand the study of greenwashing beyond only disclosure decisions (Bowen, 2014; Delmas and Burbano, 2011; Lyon and Montgomery, 2015), and encompasses social as well as environmental outcomes – focusing more broadly than the term's origin in environmental actions. Furthermore, this definition of greenwashing can encompass intentional and unintentional actions by management. I consider

well-intentioned ESG policies that do not produce positive environmental or social outcomes to be greenwashing, albeit unintentional greenwashing. As such, greenwashing can be a strategic choice, a failure of execution, or simply the result of circumstances outside of management's control.

There are numerous examples of firms greenwashing, including implementing ESG policies which are not supported by or followed up with action. Delmas and Montes-Sancho (2010) find that late entrants to the U.S. Department of Energy's Climate Change Program were likely to symbolically cooperate in the program, but not reduce their carbon emissions. Ramus and Montiel (2005) find that companies adopt environmental policies due to isomorphic pressures, but, on average, firms selectively adopt policies to appear green, without substantively improving outcomes. Environmental organizations point out that most environmental claims on products are in some way misleading (Terrachoice, 2010).

Delmas and Burbano (2011) introduce a theoretical framework which argues that greenwashing is due to market, nonmarket, organizational, and individual psychological factors. Empirically, research has found evidence identifying market, nonmarket and organizational level factors. At the institutional and regulatory level, greenwashing thrives when there is weak stakeholder and regulatory oversight of firms (Delmas and Montes-Sancho, 2010; Kim and Lyon, 2011; Marquis *et al.*, 2016; Ramus and Montiel, 2005). Turning to firm-level characteristics, both large and growing firms are more likely to greenwash (Kim and Lyon, 2011, 2015; Ramus and Montiel, 2005). For example, Kim & Lyon (2015) find that growing investor-owned utilities greenwash to gain support from stakeholders necessary for expansion, while firms with poor financial performance brownwash, possibly to avoid giving investors the impression that they are sacrificing financial value for environmental outcomes.

In this paper, I turn attention to the preferences of actors within the firm – managers and investors – and ask how they influence greenwashing. To do so, I take a principal-agent framework

(Hölmstrom, 1979). Due to the difficulty and cost of assessing firms' ESG outcomes, there is naturally an information asymmetry between firm insiders (i.e., managers), and investors with regards to ESG outcomes. To theorize as to how a principal-agent dynamic develops with regards a firm's ESG outcomes and greenwashing, I next discuss how managers and investors have heterogeneous preferences for ESG outcomes, how these preferences influence the firm's decisions around investing in ESG outcomes, and whether these preferences may lead to greenwashing.

Managers

Following a principal-agent perspective, managers are actors with inside information and the decision-making power over which ESG policies are put in place and how many firm resources are committed to achieving ESG outcomes. As such, if managers have preferences for better ESG outcomes, they can, on average, achieve those outcomes. Related empirical work has found that managers' non-financial preferences and personality characteristics do correlate with firm outcomes beyond traditional financial measures. Chin, Hambrick, and Treviño (2013) find that firms with more liberal CEOs are more likely to receive higher ESG from third parties. Personality characteristics – narcissism and hubris – also predict a firm's ESG rating (Petrenko *et al.*, 2016; Tang, Mack, and Chen, 2018). However, ESG ratings are noisy and may include both ESG policies and ESG outcomes, and as such the question remains whether managers affect real outcomes, or simply greenwash (Berg *et al.*, 2020; Chatterji *et al.*, 2016). More closely related to measuring ESG preferences and outcomes, though in a more specific domain, Muller and Kolk (2010) find that, in a sample of auto part manufacturing firms in Mexico, managers' personal commitment to ethics is the main driver of their firms' social outcomes, even more so than economic pressures. I expect this effect of managers' commitment to ethics to extend broadly to the ESG domain. Thus, I expect managers with higher ESG preferences to commit organizational resources and energy toward achieving better ESG outcomes, leading to my first set of hypothesis.

Hypothesis 1: Firms with managers with higher ESG preferences will have more ESG policies.

Hypothesis 2: Firms with managers with higher ESG preferences will have better ESG outcomes.

Investors

Managers are ultimately accountable to investors (via the board of directors). Investors aim to employ executives and managers in the firm who will act according to investors' preferences, including their preferences for ESG outcomes. However, ESG outcomes are difficult to observe. To address this difficulty in measurement, a group of ESG ratings agencies have developed under the premise that these firms will measure ESG using a single set of fixed costs so that diffuse individual investors do not have to (Avetisyan and Hockerts, 2017). However, these ratings agencies have historically relied heavily on unaudited ESG information which firms self-report (Amel-Zadeh and Serafeim, 2018). Because investors use ESG ratings as a signal of ESG outcomes, but these ratings are influenced by unverifiable policies, managers can tailor their ESG policies to their investor audience, while they commit organizational time, energy, and resources toward ESG outcomes only if doing so aligns with the manager's own preferences.

While managers may have the ability to decouple policies and outcomes, the direction of this decoupling depends on whether managers or investors have a greater preference for ESG outcomes. That is, managers may greenwash or brownwash to please investors. We thus must consider whether the agency problem at hand is one of managers overinvesting in ESG or underinvesting in ESG outcomes, relative to the investment the firms' shareholders would prefer.⁶ Benabou & Tirole (2010) outline theoretical arguments in both directions. There is the potential that ESG investment are win-

⁶ I do not consider the important issue of social optimality in this paper. ESG performance may have positive externalities through, for example, the reduction in greenhouse gas emissions. If this is the case, it may be that all firms are underinvesting in ESG relative to the social optimum. However, if a firm's investors perceive a tradeoff between personal financial returns in stock price and the firm making the socially optimal reduction in carbon emissions reduction by the firm, they may prefer the former. Thus, when I discuss overinvesting or underinvesting in ESG, the reference point is always to investors' preferred level of ESG investment.

win scenarios which managers may not recognize unless pushed by investors (Flammer *et al.*, 2019) as well as the potential that managers invest in ESG programs to gain private nonpecuniary benefits, against the preferences of shareholders (Cheng, Hong, and Shue, 2014). Empirical research also identifies both effects. Supporting the argument that managers overinvest in CSR, Werbel and Carter (2002) find that a firm's charity foundation giving is consistent with nonpecuniary benefits for the CEO, and not to the firm. Cheng, Hong, and Shue (2013) find that increasing managerial ownership and financial alignment with the firm decreases firm CSR, suggesting managers overinvest in CSR in relation to what shareholders would prefer. Personality characteristics often characterized as negative also appear to contribute to ESG activities. Petrenko *et al.*, (2016) find that narcissistic CEOs invest more in CSR and Tang *et al.*, (2018) find that CEO narcissism and comparison to peer firms influences a CEO's investment in ESG. Both studies suggest that, on average, ESG outcomes provides nonpecuniary benefits for CEOs, causing overinvestment in ESG.

However, counterbalancing this potential overinvestment in ESG, investors have become much more receptive toward firms' ESG activities over the past decades. Shin, Lee, and Bansal (2021) find that in the 1980s and 1990s investors, on average, rewarded CEOs for poor ESG ratings, but that this relationship reversed in the 2000s. In the early 2000s, sin stocks, that is, companies in industries such as tobacco or oil, financially outperformed the market (Hong and Kacperczyk, 2009). These stocks are often excluded from ESG investments, suggesting that investors were willing to pay a price to not invest in negative ESG performers. This recent transition toward a preference for ESG outcomes is consistent with Hart and Zingales's (2017) argument that managers should maximize shareholder welfare, not wealth, by considering shareholders' preferences over environmental and social outcomes. The median shareholder appears to have adopted this pro-ESG stance. Hartzmark and Sussman (2019) find that when Morningstar made their fund ESG ratings more visible, investment flowed into high-sustainability funds and out of low-sustainability funds. In very recent

years ESG investing has become a mainstream investing strategy, with one third of assets under management claiming to take ESG information into account in portfolio allocation decisions (Principles for Responsible Investment, 2019).

Many investors, through boards of directors, now perceive too little ESG investment as an agency issue, as corporations have recently begun to contract around ESG outcomes. Including ESG related targets has become more popular in the past decades, with around a third of S&P 500 firms using such incentives by 2012 (Flammer *et al.*, 2019; Maas, 2018; Maas and Rosendaal, 2016). Flammer *et al.*, (2019) find that adding these CSR criteria in executive compensation increases firms' long-term orientation and value, consistent with the idea that CSR contracting focuses executives' attention in areas that provide value but are not executives' default focus. Furthermore, this suggests that firm financial value can be increased through incentivizing managerial attention to these issues.

Moving beyond average effects, different types of investors have heterogeneous preferences for ESG outcomes. Institutional investors, specifically European institutional investors, appear to value firms' ESG outcomes (Chen, Dong, and Lin, 2020; Dyck *et al.*, 2019). Furthermore, Chen *et al.*, 2020 find that institutional investors impact firms' ESG ratings through both choice (ownership) and voice (shareholder proposals and contact with management). Long term investors also own firms with higher average ESG ratings, and push firms to use integrated reporting, a type of reporting consistent with a commitment to ESG outcomes (Serafeim, 2015; Starks, Venkat, and Zhu, 2017). However, U.S. hedge funds appear to view ESG investments as a wasteful expenditure which sacrifices firm financial value. Hedge funds target firms with higher ESG scores, and then reduce the ESG ratings of firms they acquire (DesJardine and Durand, 2020; DesJardine, Marti, and Durand, 2020). Consequently, firms targeted by activist investors have increases in employee injury rates and sickness after they have been targeted (Xia and Meyer-Doyle, 2021).

The agency problem in ESG outcomes is thus not unidirectional. Investors are heterogeneous in their preferences for ESG outcomes, as are executives. Research has identified agency issues in managers overinvesting in ESG projects relative to investor preferences, as well as underinvesting. Thus, based on a firm's investor base, the firm's managers may make choices that cause firms to either greenwash or brownwash (Kim and Lyon, 2015). However, due to the unverifiable nature of the effectiveness of firms' ESG policies, managers will be able to greenwash or brownwash to an investor audience. Formally,

Hypothesis 3: Firms with investors with high ESG preferences will have more ESG policies on the books.

Hypothesis 4 (not formally tested): Investor ESG preferences will have a larger effect on ESG policies than ESG outcomes.

Constraints to Greenwashing

If a firm has ESG-oriented investors, why do the managers not, at low cost, put in place all possible ESG policies to placate their investors and simply run their business according to their (the managers') own preferences? Because, as firms increasingly greenwash, there is an increasing chance of being targeted by an environmental or social activists, as Baron and Diermeier (2007) establish theoretically and King and McDonnell (2012) find empirically. This activist-imposed cost to the firm may take the form of a boycott, leading to loss in sales or lower employee engagement. These activists can damage a firm's reputation and impact a firm's financial position (Kim and Lyon, 2011; McDonnell and King, 2013; McDonnell *et al.*, 2015). Thus, managers may choose to greenwash, but will be limited in how much they do so due to the threat from environmental activists.

Another potential limit to greenwashing comes from investors. A fully rational investor would understand that the ESG information they receive is substantially based on unaudited, and unverifiable, firm reported ESG policies. As such, a rational investor would discount that information. If ESG information were sufficiently unverifiable to produce no signal value, investors

would place no weight on ESG policies, and, as such, ESG policies would not correlate with investment trends. Two empirical facts stand in the way of this information-discounting hypothesis. First, ESG ratings have been largely based on self-reported data (Berg *et al.*, 2020), and second, investors do use these ESG ratings (Amel-Zadeh and Serafeim, 2018). While sophisticated investors may see through this problem, and point it out in popular press (Mackintosh, 2018), the fact remains that ESG investment has increased rapidly over the 2010-2020 timeframe (Principles for Responsible Investment, 2019). The sophistication of institutional investors does beg the question I return to in the discussion section: are firms greenwashing in response to investors, or are institutional investors allowing firms to greenwash, to the detriment of less sophisticated retail investors?

Empirical Application

To empirically test my theory, I gather data on large U.S. public firms from 2008 to 2019. To measure greenwashing, I source ESG data from TruValue (now owned by FactSet) and Thomson Reuters. Before analyzing how managers' and investors' preferences correlate with greenwashing, I describe my greenwashing measure and provide descriptive statistics. To introduce the measure, I return to the earlier example of Altria, who in 2018 claimed youth vaping was the most pressing issue they faced, while also buying a large stake in the company most responsible for the youth vaping epidemic.

The contradiction in Altria's stated claims around its environmental and social (ESG) impacts, and the consequences of the firm's actions can be seen more broadly by comparing their corporate policies meant to ameliorate negative (and improve positive) ESG impacts against the consequences of their corporate actions as judged by society. In 2008, Altria had 25 ESG policies on the books⁷, the industry median number. By 2018, they had a total of 58 ESG policies on the books, the 9th most

⁷ Using policies tracked by Thomson Reuters, a major provider of ESG ratings.

out of 215 companies in their industry. In fact, Altria was in the top 10% of firms in their industry in terms of ESG policies on the books for every year from 2011 through 2019.

News media coverage of Altria specifically covering ESG issues, a measure of how society assesses the company's ESG outcomes, paints a very different picture. For example, in 2019, the F.D.A. accused Altria and Juul from backing off of promises to address youth vaping – the very challenge that Altria identified in their 2018 CSR report as the most important they faced (Kaplan, 2019). Looking across a wide array of ESG issues and across hundreds of thousands of news sources, Altria received more negative coverage of their ESG outcomes than 75% of their industry peers every year from 2011 through 2019, often falling in the bottom 5% of firms.

I measure greenwashing empirically by comparing these two features: a firm's ESG policies and the firm's ESG outcomes as assessed by news media. To measure ESG policies, I aggregate individual datapoints from the Thomson Reuters ESG dataset (formerly ASSET4).⁸ There are 67 different policies firms can adopt, which include, for example, "Waster Water Reduction Program" and are rated as either 0 (no policy) or 1 (firm has a policy). I generate a total policies variable as the total count of policies.

I measure ESG outcomes from news media coverage gathered from TruValue Labs, an ESG data provider. TruValue Labs machine reads articles from tens of thousands of news media sources daily. Their algorithm then uses natural language processing to identify articles covering firms, categorizes these articles by ESG issue and scores the articles as being positive or negative ESG events on a 0-100 scale. These articles, by in large, report on actual ESG outcomes, rather than firm reputations or PR statements. For example, an article stating that Johnson and Johnson was sued for allegedly deceiving patients and doctors about the addictive nature of opioids, the article was

⁸ This dataset includes the individual metrics used in developing firms' overall ESG scores, which are unavailable from other ESG data providers. Hawn and Ioannou (2016) use this same micro-data to compare firms' internal facing ESG activities to their external-facing ESG activities.

classified as negative and categorized under “Product Quality and Safety.” When employees at Textron Inc, an industrial conglomerate, were injured in a plant explosion, the article was again negative, and categorized under “Critical Incident Risk Management.” Further examples articles can be seen in Table 1.

[Insert Table 1 Here]

While the TruValue data is backwards looking and contains firm-level measures as far back as 2008, it is important to note that at the time these articles were published, stakeholders, including investors, by in large did not have access to aggregated news media at this level. Only with the advent of natural language processing and the demand from the investment world did products such as TruValue arise. Thus, while stakeholders going forward may use this sort of data to make decisions, no investor in 2016 had access to this dataset, and only in 2021 did this type of data become regularly used by investors. Thus, while I expect the empirical patterns this paper uncovers to potentially be impacted by the increased visibility of aggregated news measures among investors, the patterns this paper reveals can appropriately inform us about the causes of greenwashing.

To measure the positive or negative ESG impacts of firms at a yearly level I use a weighted average of articles about the firm, weighing more recent articles more heavily.⁹ The final ESG Outcomes measure assesses firms on a 0 to 100 scale, with higher values of this measure indicate more positive ESG outcomes.¹⁰

I classify firms as green, greenwashing, brown or brownwashing based on their ESG policies on the books and their ESG outcomes as reported in news media. Green firms are firms that have many policies on the books and positive ESG outcomes. Greenwashing firms have many policies on

⁹ I use TruValue’s Material Insight Score, in which articles are weighted using an exponential weighted average with a half-life of 6 months. TruValue data is matched to Thomson Reuters data using the last day of the firm’s fiscal year, rounded to the nearest month end.

¹⁰ This 0-100 scale is an absolute scale. It is an overall measure of the average positivity or negativity of ESG events and is not normalized by industry or year. However, in regression analysis I use normalized versions of the variable.

the books, but poor ESG outcomes.¹¹ Brown firms are firms which have few policies on the books and poor ESG outcomes. Brownwashing firms are those which have few policies on the books but high ESG outcomes.

Table 2 lists the firms that exemplify the characteristics of each category from 2016 through 2019. For example, as a green firm, GAP has both many ESG policies on the books as well as positive media assessments of their ESG outcomes (all relative to their industry peers). As exemplars of a greenwashing firms in different industries, Exxon Mobil and Altria both have many ESG policies on the books, but very poor ESG outcomes reported by the news media.

[Insert Table 2 Here]

For a more detailed view of firms in Food and Household Goods Sector, Table 3 shows further details for firms that exemplify each category from 2016 through 2019. Pilgrim's Pride is classified as a brown firm because of the relatively few polices on the books (23) as well as the poor outcomes reported in news media. Pilgrim's Pride faced both significant recalls in 2016, as well as lawsuits from farmers alleging that chicken buyers, including Pilgrim's Pride, colluded to keep farmers' pay lower than it would be in a freely competitive marketplace.

[Insert Table 3 Here]

If ESG policies are a signal of better ESG outcomes, then there should be a positive correlation between policies and outcomes. However, if greenwashing is prevalent, there may be no, or even a negative correlation between policies and outcomes. To overview the prevalence of greenwashing, I outline the correlation between policies and outcomes, how this correlation varies by sector and year, the role of firm size, and discuss the potential for lags between policy implementation and outcomes.

¹¹ There may be a lag between the time firms put policies in place, implementation of those policies, and subsequent results. While the classification system here does not take time lags into account, lags are incorporate and addressed in the following section. These illustrative examples change very little if time lags are included.

In outlining greenwashing among large firms, it is important to note that environmental and social issues vary significantly across sectors and have become more salient for investors and managers over the past decade. To identify correlations between policies and outcomes across industries and over time, I normalize both policy and outcome variables within sector-year and use these normalized variables by default in regressions.¹²

I first turn to bivariate correlations between outcomes, policies, and firm size for 4,000 large US corporations between 2016 and 2019. Correlations are displayed in Table 4. Most notably, policies and outcomes show no consistent correlation. The raw count of policies has a small positive correlation with outcomes, and outcomes normed by industry-year. When policies are normed by industry-year, those correlations become small and negative. Large firms have more policies on the books, and the correlation is strong, ranging from 0.41 to 0.57 depending on the measure. However, large firms have slightly worse outcomes reported by news media, on average, with correlations between firm size and outcomes of around -0.03.

[Insert Table 4 Here]

These correlations vary by sector and over time, as shown in Table 5. Policies and outcomes are negatively correlated in the energy and technology sectors, positively correlated in the industrial and resource sectors. In other sectors, policies and outcomes have both positive and negative correlations across different years.

[Insert Table 5 Here]

The strong correlation between policies and firm size is consistent with the idea that the expectations for ESG reporting are higher for large firms. Furthermore, these firms may be able to commit to larger fixed cost budgets to their CSR teams which both preparing their CSR reports and

¹² For example, the normalized ESG policy measure is calculated as $Normalized\ Policies = \frac{Policies - \overline{Policies}_{jt}}{\sigma_{policies_{jt}}}$, where $\overline{Policies}_{jt}$ is the average number of policies in the focal sector (j) and year (t), and $\sigma_{policies_{jt}}$ is the standard deviation of the number of policies firms have in the focal sector-year.

may develop their policies on the books. For regression analysis, I control for firm size. The varying correlations between outcomes and policies is consistent with the idea that greenwashing is standard in some sectors, while in others, policies on the books provide a good signal of the firm's ESG impacts on society.

Firms may implement policies which result in improvements in outcomes only after months or years of implementation. By comparing concurrent outcomes and policies, I may be ignoring this time lag. To investigate whether there is a time lag, I regress ESG outcomes on ESG policies with varying time lags of 0, 1 or 2 years. All regressions are shown in Figure 1. Notably, the correlation between policies and outcomes is small and not statistically significant in any models and does not seem to be increasing as the lag between policy implementation and outcomes increases.

Over the 2016-2019 timeframe, reporting on ESG issues became more standardized and prevalent across large US corporations. Has the slight increase in standardization and greater reporting expectations increased, decreased, or had no effect on the signaling value of ESG policies? To examine this, I correlate policies and outcomes by sector-year. These results are also visible in Figure 1. The coefficient on policies in the regressions predicting outcomes is always small and never statistically significant. Policies have not become a better predictor of outcomes over time, suggesting that industry-led efforts around ESG reporting have not improved the signal in ESG policy adoption, and consistent with the idea that they have not decreased greenwashing.

[Insert Figure 1 Here]

These results showing that policies do not predict outcomes, suggest that greenwashing is widespread and is not decreasing over time. I now turn to the paper's main question do investor or CEO claims around ESG, as captured by funds labeled as "sustainable" and discussion of ESG during earnings calls, correlate with greenwashing (policies, not outcomes) or substantive improvements in ESG performance (policies and outcomes)?

Data and Sample

My sample consists of the 1,500 largest U.S. public firms from 2008 to 2019. In addition to the greenwashing data from TruValue and Thomson Reuters, I use data from MSCI, Sustainalytics and RobecoSAM, earnings call transcripts from Capital IQ, investors holdings from Morningstar and 13F filings accessed through WRDS, and firm financial and industry membership data from Compustat.

Dependent Variables

My outcome variables are firms' ESG policies and firms' ESG outcomes, as presented in the previous section.

Independent Variables

To find an empirical proxy for investors and managers' ESG preferences, I turn to their public ESG stances. For investors I use the descriptions of their investment fund offerings, and for managers I use their statements to investors in earnings calls.

Investors' Stated ESG Preferences

To measure a firm's investors' stated ESG preferences, I measure the proportion of their institutional shareholders' funds labeled as sustainable. I first develop an ESG measure at the institutional investor level. To do so, I gather all funds listed as sustainable according to their prospectus from Morningstar. I merge this data with 13f filings¹³ to calculate an institutional investor-year level measure of total assets under management (AUM) and sustainable AUM. I then calculate, at the institutional investor-year level, a measure of ESG orientation by dividing sustainable AUM by total AUM. Using 13f filings to identify each firm's investors, I then calculate, at the firm-year level, a weighted average of all the firm's investors' ESG orientations.

¹³ All institutional investors with over \$100M in assets must file 13f reports indicating how their assets are invested.

Previous research has found that long term investors have higher ESG preferences than short term investors (Oikonomou, Yin, and Zhao, 2020). As a robustness check, I calculate a measure of investor time horizon as an alternative measure of a firm's investors' ESG preferences. I follow past research and calculate a firm's investors' time horizon as the proportion of long term (dedicated) investors minus short term (transient) investors using the classification scheme from Bushee (2001).

Top Management's ESG Orientation

To develop an empirical proxy for the CEO and top management team's ESG preferences, I examine the language they use on earnings conference calls, following past work in finance (Brochet, Loumioni, and Serafeim, 2015; Loughran and McDonald, 2010; Yang, UY, and Huang, 2020) and emerging work around ESG issues (Smeuninx, De Clerck, and Aerts, 2020).

I measure the volume and sentiment of ESG-related language managers use on earnings calls and use this as a proxy for their ESG preferences. I identify ESG-related text in earnings calls using two methods: word counts, and semantic distance. I first use a word count measure, which comes with the advantage of being easily interpretable and used in past research. I use a list of words from Babcock *et al.*, (2022), which builds on a word list from Vracheva, Judge, and Madden, (2016), and extend the list of ESG related words based on a review of dozens of ESG reports. I then calculate, at the paragraph level, the proportion of ESG-words in the paragraph. This proportion, aggregated to the call level, is the "ESG volume" measure. I then measure the positivity and objectivity of the ESG language from top managers. To do so, I calculate the positivity and objectivity of each paragraph in the earnings call using TextBlob, a widely used natural language processing package. To calculate an ESG-positivity (and ESG-subjectivity) score, I multiply, at the paragraph level, the ESG volume measure and the positivity (objectivity) measures and then sum this score across all paragraphs spoken by executives. I then divide the final measure by the number of ESG words to generate a measure of positivity (and objectivity) per ESG word.

The disadvantage of using a word count measure is that it misses the context in which words are used in sentences. To account for this, I also introduce a novel measure of ESG language using recently developed natural language processing techniques. I measure the semantic distance between executives' language and the Sustainability Accounting Standards Board's (SASB) ESG standards. SASB is a nonprofit-run standards setting agency whose standards are now commonly used throughout the corporate and financial system. I calculate this semantic distance by converting earnings call paragraphs into vectors of real numbers using MPNet and Sentence Transformers (Song *et al.*, 2020), and then calculating the cosine similarity between each paragraph of earnings call text's vector representation and each of the 26 SASB standards' vector representations. This method identifies paragraphs that are, for example, close to the phrase "Communities in which we operate and impacts on core human rights and the treatment of indigenous peoples" or "Direct (Scope 1) greenhouse gas (GHG) emissions that a company generates." By using this large scale pre-trained model and sentence embeddings (vector representations), the algorithm takes advantage of the dense information in natural language, including both the semantic meaning of individual words as well as their position in sentences and paragraphs. This allows the algorithm to identify phrases and paragraphs with similar meanings, rather than finding only matching words or synonyms.

For all ESG language measures, I generate separate measures for the presentation and Q&A sections.

Control Variables

I control for firm and industry characteristics using financial variables from Compustat. I control for the number of employees, total revenue, and total assets in regression models. I also control for sector effects by using 2-digit SIC sector values. Finally, I use year fixed effects to control for the increasing salience of ESG issues over my sample timeframe.

Descriptive Statistics

I start with the sample of the largest 1,500 firms in the U.S. from 2008 to 2019. My sample is then restricted to firms covered in the TruValue dataset and firms for which I can gather earnings calls. Investor data from Morningstar and 13f filings does not significantly reduce my sample size. Merging in Thomson Reuters ESG data restricts the sample to 5,486 firm-year observations and skews the data toward recent years. Long term shareholding data, which I use in a robustness check, drops a few additional firm-year observations. Given this year skew of the sample, my results are thus largely driven by recent years (2015-2019). The sample is outlined in Table 6.

[Insert Table 6 here]

Summary statistics for my key variables are outlined in Table 7. Firms in the sample have between 0 and 67 ESG related policies, with an average of 21 and a median of 16 policies. My measure of outcomes, news sentiment from TruValue, has a mean value of 61.6 and a median value of 63.3, suggesting that the average ESG related news article in the sample has a slight positive sentiment (as a score of 50 represents neutral sentiment). Earnings call ESG focus measures have a slight positive skew (and are normed in the empirical models). Firms' investors have, on average, around 1% of their AUM labeled as sustainable.¹⁴ In unreported analysis, I find that both ESG mentions on earnings calls and investor ESG percentages are increasing greatly in later years of the sample. As a robustness check, I use shareholder time horizon. Long Term Shareholders is defined as the difference between the percentage of a firms shares held by dedicated shareholders and transient shareholders. On average, firms have more transient shareholders than dedicated shareholders.

[Insert Table 7 here]

¹⁴ While one third of assets under management in the United States claim to consider ESG information when making asset allocation decisions, funds explicitly labeled as sustainable in their prospectus make up a much smaller portion of total assets under management.

Half a percent of the words used in earnings calls relate to ESG issues, with a slightly positively skewed distribution. The median ESG word count is 0.463 percentage points, and the mean is 0.552 percentage points. Some earnings calls have no ESG related words, and the most ESG-focused earnings call is made up of 2.65 percent ESG words. The SASB measure is difficult to interpret because it is a measure of cosine similarity in a high dimensional space but shows a similar positive skew. As shown in Table 8, the two are moderately correlated (rho of 0.26).

Examples of how managers speak about ESG issues are in the appendix. Discussion of ESG issues experienced a short peak around the 2008 financial crisis, and after falling has steadily increased over the 2011-2020 period, with an especially pronounced increase since 2018. When, and how much management discusses ESG issues appears to be reflective of a firm's ESG reputation, at least among firms with strong ESG reputations (both positive and negative). Unilever, a firm with a very high ESG reputation, and Exxon Mobile, a firm with a very low ESG reputation, provide a key contrast. In their Q3 2020 earnings call, the Unilever CEO explicitly states their stakeholder-oriented framework in the opening paragraphs of the call. Alternatively, Exxon Mobile executives do not discuss ESG issues at all in their presentation, and only touch on renewable energy sources when asked by an analyst in the Q&A section in their Q3 2020 call. These differences are not just driven by industry. In BP's Q3 2020 earnings call, executives discuss stakeholders and renewable resources over multiple paragraphs of the presentation section of the call. For further summary statistics, see the appendix.

In all regression analysis I normalize all ESG variables by sector-year to account for change in the salience of ESG issues over time and the different expectations around ESG policies across sectors.

Results

Correlations between all key variables are presented in Table 8. Of note, the number of ESG policies a firm has correlates positively, but very weakly, with ESG outcomes in the news. The raw correlation between the two variables is 0.02. ESG language on earnings calls correlates with both ESG policies and ESG outcomes. Investor ESG orientation, on the other hand, correlates positively with ESG policies but have a very small and negative correlation with ESG outcomes. These correlations summarize the results told more precisely through regression analysis and visually through binned scatterplots in the remainder of the results section.

[Insert Table 8 here]

Regression results showing the correlation between managers' and investors' ESG preferences, controlling for financial characteristics, are found in Table 9 through Table 12. I first discuss the results related to managers, and then turn to investors.

Managers

The relationship (with no controls) between managers' ESG language and ESG policies and outcomes is shown visually in the binned scatterplot in Figure 2. From this raw correlation, we can see that firm-years with greater discussion of ESG issues on earnings calls have more ESG policies on the books, and better ESG outcomes as reported by news media.

[Insert Figure 2 here]

To examine whether these relationships are accounted for by alternative drivers such as firm financial characteristics or yearly trends, Table 9 regresses ESG policies and outcomes on both managers' ESG language and investors' ESG claims. Managers' ESG language is measured through word ESG counts. In models (1) and (2) of Table 9, the dependent variable is ESG policies on the books, while in models (3) and (4) the dependent variable is ESG outcomes. Models (1) and (3)

include sector and year fixed effects, while models (2) and (4) include firm and year fixed effects. The independent variables and control variables are lagged by one year in all models.¹⁵

The coefficient on Managers' ESG Language is positive in all models, but not always statistically significant. The coefficient on Managers' ESG Language has t-statistics of 3.64 and 5.09 in models (1) and (3), respectively. This means that across firms in the same industry, managers' ESG language is predictive of more ESG policies and better ESG outcomes. All key variables are normalized thus these results indicate that, across firms in the same sector, an increase of ESG language on earnings calls by 1 standard deviation is associated with a 0.04 standard deviation increase in policies on the books and a 0.08 standard deviation increase in ESG outcomes.

In firm fixed-effects models, results for managers' language are consistent directionally, but the point estimates in these models are around half to one third the size of the sector-fixed effects models and have confidence intervals which overlap with zero. Thus, these results suggest that while managers' ESG language is predictive of ESG outcomes in the news across firms within the same industry, they do not provide evidence that these effects are present within firms over time.

[Insert Table 9 Here]

I next examine whether the way in which managers speak about ESG issues predicts their firm's future ESG policies and outcomes. Table 10 includes the word count measure of managers' ESG language as well as the polarity (positivity) and subjectivity of the ESG language managers use. The coefficients on managers overall ESG language do not change significantly – they again predict more policies and better outcomes across firms in the same industry, but not within firms over time. The coefficient on the polarity of ESG language is 0.05 in model (1) and small and not statistically

¹⁵ In unreported results, I find that all results are substantively the same when using concurrent independent variables, or concurrent independent and control variables.

significant in model (2). This indicates that firms with managers who speak more positively about issues are likely to have more policies on the books than their industry peers. The coefficient on ESG language positivity is 0.05 and 0.06 in models (3) and (4), respectively. This indicates that firms with managers who speak more positively about ESG issues are likely to have better ESG outcomes than their industry peers, and improve their own ESG outcomes over time.

The coefficient on the subjectivity managers' ESG language is -0.049 in model (1) and -0.06 in model (3), and close to zero in the other models. This indicates that firms with managers who speak more subjectively about ESG issues are likely to have fewer ESG policies and worse ESG outcomes than their industry peers. However, the subjectivity of their language does not predict within firm changes over time.

[Insert Table 10 Here]

To examine the robustness of the method of measuring managers' language, I run the same analysis using the semantic distance measure in place of the word count measure for managers' ESG language. These results are displayed in Table 11. As in Table 9, the dependent variable in the first two columns is ESG polices, and in the last two columns is outcomes. Sector fixed effects are used in columns (1) and (3), and firm fixed effects are used in columns (2) and (4). The results from this model are consistent with the results using word count measure, with slightly smaller point estimates. In the sector fixed effects models, a one standard deviation increase in managers' use of ESG words is associated with a 0.03 standard deviation increase in ESG policies on the books and a 0.05 standard deviation increase in ESG outcomes the following year. In the firm fixed effects model, the coefficient on managers' ESG language is positive and not statistically significant.

[Insert Table 11 Here]

Shareholders

Turning to shareholders, these regression results show a clear distinction between (a) the correlation between shareholder ESG preferences and ESG policies, and (b) the correlation between shareholders ESG preferences and ESG outcomes. These relationships are presented visually in the binned scatterplot in Figure 3. The binned scatterplot shows that firms with investors who have more assets under management in sustainable-labeled fund have more ESG policies on the books. However, there does not appear to be a clear relationship between firms' investors and ESG outcomes.

[Insert Figure 3 here]

Turning to regression results in Table 9, the coefficient on Investors' ESG Orientation in columns (1) and (2) of are both positive, with t-statistics of 3.33 and 2.46, respectively. These coefficients indicate that, across firms in the same industry, a one standard deviation increase in a firms' investors' ESG orientation is associated with a 0.03 standard deviation increase in their ESG policies on the books. Within firms over time, a one standard deviation increase in investor ESG orientation is associated with a 0.02 standard deviation increase in ESG policies.

The correlation between investor ESG orientation and ESG outcomes is both smaller in economic terms, relative to the effect on policies, and not statistically different than zero. The coefficients on Investors' ESG Orientation in columns (3) and (4) have t-statistics of 1.2 and 0.22, respectively. These results provide evidence that increases in a firm's investors' ESG preference is correlated with an increase in ESG policies but does not provide evidence that ESG outcomes change with investors' ESG preference.

Previous research has found that long term investors prefer firms with high ESG ratings, and push firms to improve their ESG ratings. As a robustness check, I use investor time horizon as an alternative proxy for investors' ESG preferences. To calculate investor time horizon, I use the

proportion of a firm's shares held by dedicated investors minus the proportion of shares held by transient investors according to the classification system developed in Bushee (2001). Results are visible in Table 12. This alternate measure has the advantage of being established in the literature, but the disadvantage of being a less direct measure of ESG orientation. Table 12 is set up just as Table 11, but uses investor time horizon rather than investors' ESG funds as the key independent variable for investors. Column (1) shows that a one standard deviation increase in long term investors is correlated with a 0.036 standard deviation increase in ESG policies on the books in the sector fixed effects model. Within firms over time, long term shareholders are associated with lower ESG policies, but the coefficient has a 95% confidence interval overlapping with zero (t-stat of -1.69). With regards to outcomes, a one standard deviation increase in long term investors is correlated with a 0.0385 standard deviation *decrease* in ESG outcomes in the sector fixed effects model (t-stat of -2.55) and with a 0.0344 decrease in ESG outcomes in the firm fixed effects model (t-stat of -1.84).

[Insert Table 12 here]

In aggregate, these results suggest that the amount of time managers spend discussing ESG issues on earnings calls predicts both more ESG policies and better ESG outcomes, compared to firms in the same industry. The positivity of that ESG discussion predicts more policies and better outcomes both across firms and within the same firm over time.

Investor ESG orientation is associated with increased ESG policies both across firms and within firms over time but appears to have little or no association with ESG outcomes. In fact, firms with more long-term shareholders, who have been suggested to be more ESG-focused, have slightly worse ESG outcomes than peer firms.

Ratings Disagreement

A practical implication of greenwashing is that it increases the cost for investors and stakeholders to accurately assess firms' environmental and social outcomes, perhaps prohibitively so. This could be reflected in the fact that there is high disagreement between ESG ratings. Chatterji *et al.*, (2016) found that as of the mid 2010s, individual companies' ESG ratings are highly variable across ratings agencies, with inter-rater correlations of around 0.5, in comparison to financial ratings inter-rater correlations of around 0.97. Berg *et al.*, (2020) find the same pattern still exists through 2019, with the source of divergence split between scope disagreement (i.e., what behaviors and outcomes are ESG ratings meant to capture) and measurement error. Practitioners have also questioned our current ability to assess firms' ESG outcomes, citing both the disagreement over methodology and data (Mackintosh, 2018; Peirce, 2019). Researchers and practitioners are thus left with the question of what ESG information can be valuable, and what drives firms to potentially misrepresent their ESG outcomes.

As both a robustness check of the validity of using ESG outcomes reported in the news and ESG policies to assess greenwashing behavior, and as a practical insight for researchers and practitioners, I next examine the relationship between the ESG policy-outcome gap and ESG ratings disagreement. I define greenwashing empirically in this section as the number of ESG policies a firm has, normed at the sector-year level, minus the TruValue measure of ESG outcomes, again normed at the sector-year level for comparability. I calculate ratings disagreement as the variance between three large ESG ratings providers: MSCI, Sustainalytics and RobecoSAM.¹⁶ Due to data access, I have ESG ratings disagreement for only two years: 2018 and 2019, thus I present results only with sector fixed effects, as firm fixed effects would only capture one year of change. Results are presented visually in Figure 4. There is a positive relationship between greenwashing and ratings

¹⁶ I exclude Thomson Reuters because I gather my policies measure using Thomson Reuters data, but these results all hold if the Thomson Reuters ESG score is included in the ratings variance measure.

disagreement. Regression analysis is presented in Table 13, and confirms that this relationship is present across firms in the same industry.

[Insert Figure 4 and Table 13 here]

Discussion

This paper presents a theory of how a firm's stakeholders shape the firm's use of symbolic policies. Specifically, in the context of social and environmental performance, I examine how managers and investors shape greenwashing. Using a principal-agent framework, I hypothesize that firms greenwash in response to investor pressure around ESG issues when investors do not have sufficient information to verify managers ESG claims, and when managers are not ESG-oriented.

Empirically, I examine whether the ESG preferences of management and investors predict ESG outcomes and policies and find correlational evidence consistent with the principal-agent framework. Over the past decade, managers' ESG preferences, as proxied by their ESG language in earnings calls, correlates with ESG policies and ESG outcomes. However, a firm's investors' ESG preferences correlate strongly with ESG policies and the books, but not with outcomes, suggestive of greenwashing. These results are robust to alternative measures of managers and investor preferences, including investor time horizons and multiple methods of measuring managers' language.

This study has theoretical implications to multiple streams of research. First, this study suggests that individual and organizational preferences, especially those of top management, around ESG issues are driving factors in firms' ESG strategies. This promotes that idea that looking inside the firm can provide insight into how firms can benefit society. However, this study also provides a note of caution for ESG research. Significant ESG research has focused on the relationship between firms' ESG (or CSR) performance and financial performance. This study suggests that if the measures the research community uses to proxy for ESG performance are at best noisy and at worst

biased due to greenwashing, the true relationship between social and financial performance may be obscured when using historical ESG ratings.

More broadly, this paper contributes to sociological research on decoupling (DiMaggio and Powell, 1983; Meyer and Rowan, 1977). The empirical patterns observed here are consistent with organizations decoupling policy from practice to maintain external legitimacy while maintaining internal efficiency. Adding to past research, this paper shows that individual preferences, beyond industry and financial factors, may also affect the decoupling decision. Lastly, this paper suggests that agency problems around firms' ESG strategy could result in either overinvestment or underinvestment in ESG, relative to investors' preferences. Research has presented conflicting views of the directionality of the agency problem in ESG (Cheng *et al.*, 2014; Flammer *et al.*, 2019). This paper provides a model in which both are possible, depending on the relative ESG preferences of managers and investors.

This study contributes novel data measures and techniques to ESG research. I machine read earnings call transcripts and semantic distance and ESG word frequency to measure managers' ESG orientation. Semantic distance specifically can be used to generate insight across a broad array of strategic management questions, by providing an automated method for identifying what managers are talking about that can use the dense meaning of text across large corpuses. I also develop a new measure of greenwashing by comparing ESG policies to ESG outcomes reported in the news. This new method covers a larger set of firms and a wider set of ESG concerns than past research into greenwashing.

The sample, the correlational results, and the empirical proxies used in this study provide reasons to extend this work to identify causality and to consider the boundary conditions in interpreting this study. Greenwashing and the investor ESG movement are currently changing significantly in terms of scope and regulatory environment. Thus, it is important to note that this

study examines the behavior of firms in the United States between 2008 and 2019. The results presented here provide insight into this time and place and could change dramatically across jurisdictions and timeframes in which the reporting standards and culture around ESG issues vary. For example, ESG reporting in the United States is currently moving toward standardized frameworks, led by organizations such as SASB and the Global Reporting Initiative (GRI), and accounting for ESG reporting is moving toward assured, if not audited, data. These changes may help attenuate greenwashing.

The empirical results presented in this paper are correlational, and thus warrant additional investigation. While the results are consistent with a causal story and with the principal-agent theory, they do not provide causal evidence. Managers (and thus their ESG preferences) are not randomly assigned to firms, nor, likewise, are investors assigned randomly. Further research could examine exogenous changes to a firms' management and investor base or their preferences, and how greenwashing behavior changes as a result. Additionally, research could investigate whether greenwashing firms match with certain types of managers or investors.

The empirical proxies used in this study could introduce bias to the empirical analysis. ESG language may reflect managers' perceptions of demands from stakeholders, in addition to or rather than a reflection of their own preferences. Without dismissing this possibility, at a minimum this study's results indicate that this language is a signal that firms' ESG outcomes are judged to be more positive (and less harmful) by a broad set of societal stakeholders. Second, measuring ESG outcomes through media coverage may introduce bias. While individual media outlets undoubtedly come at reporting issues from a specific perspective, the breadth and specificity of the data sources gathered by TruValue, along with the elimination of press releases in their algorithms, provide some level of comfort that the events reported reflect actual outcomes. The use of these proxies is motivation for further research using alternative measures of ESG outcomes. Recent research looking

at only managers' language using alternative measures of ESG performance have found consistent results, suggesting that news media is in line with direct and more specific measures (Chava *et al.*, 2021).

Beyond theory, these results have significant implications for managers, investors, and regulators. The surface level takeaway is that over the past decade, greenwashing has been a feasible strategy for managers who are not committed to ESG issues. However, this result also indicates that for managers who want to pursue substantive ESG strategies, they will have to establish more robust methods to credibly signal their ESG commitments. Following the signaling perspective taken in this paper, this means that ESG-oriented managers must find practices that are sufficiently costly for low-ESG firms to imitate such that stakeholder will be able to differentiate high and low performing firms.

For investors, these results suggest that over the past decade, many firms have put in place symbolic ESG policies as a result of investor pressure, and these policies have not, on average, resulted in improved outcomes. Again, these results could have two interpretations. If institutional investors do not care about ESG outcomes, and only want to attract retail investors attracted to ESG-labeled products, these results indicate that institutional investors themselves can successfully greenwash their products. However, if institutional investors do want to push firms toward better ESG outcomes (as their retail clients appear to want based on their investment choices), they will need more reliable ESG signals and data from firms. This could mean demanding standardized, mandatory, and/or audited or assured ESG data. In terms of immediate practice, the finding that greenwashing correlates with ratings disagreement suggests that investors interested in measuring firms' ESG performance ought to consider using data reliant only on independently reported ESG outcomes in addition to traditional ESG ratings.

Governments and regulators both in the US and across the world are currently debating how to standardize ESG reporting. These results suggest that these steps will be necessary to turn the goals of the ESG and stakeholder capitalism movement into improved firm outcomes and impacts. If investors do not have reliable signals of ESG outcomes, firms can greenwash, and investor led stakeholder capitalism will not produce the substantive change it claims to pursue.

Tables and Figures

Table 1: Examples of ESG Related Articles from TruValue

MT / ArcelorMittal		
ESG Effect	Negative	Category Waste & Hazardous Materials Management
Article	ArcelorMittal has been accused of manipulating test results to conceal the amount of ammonia and cyanide that is being discharged from its plant in Indiana.	
Source	Steel Guru: steelguru.com/steel/arcelormittal-burns-harbor-accused-of-manipulating-test-results-by-idem	
AMZN / Amazon.com		
ESG Effect	Negative	Category Labor Practices
Article	Amazon has allegedly threatened to fire employees for criticizing the company's environmental policies.	
Source	Vanity Fair: vanityfair.com/news/2020/01/amazon-threatens-to-fire-employees-climate-change-environment	
JNJ / Johnson & Johnson		
ESG Effect	Negative	Category Product Quality & Safety
Article	Washington has sued Johnson & Johnson for allegedly deceiving patients and doctors about the addictive nature of the opioid pain killers it developed and marketed.	
Source	My Northwest: mynorthwest.com/1658163/ag-ferguson-johnson-and-johnson-opioid-lawsuit	
ORSTED / Orsted A/S		
ESG Effect	Negative	Category Systemic Risk Management
Article	Orsted will pay £4.5 million because its power station failed to remain connected after a lightning strike, causing a power outage in the UK that affected more than one million customers.	
Source	New York Times: nytimes.com/reuters/2020/01/03/business/03reuters-britain-powercut.html	
GE / General Electric Co		
ESG Effect	Positive	Category Product Design & Lifecycle Management
Article	GE said it won 1,215 MW of wind turbine orders or commitments from customers in China in 2019.	
Source	Renewables Now: renewablesnow.com/news/ge-bags-12-gw-of-wind-orders-in-china-in-2019-681584	
TXT / Textron Inc		
ESG Effect	Negative	Category Critical Incident Risk Management
Article	Several employees of Textron Inc were reported injured following an explosion at a plant in Wichita, Kansas.	
Source	News 8000: news8000.com/news/national-news/explosion-reported-at-kansas-aviation-manufacturing-plant	
ARCO / Arcos Dorados Holdings		
ESG Effect	Negative	Category Employee Health & Safety
Article	Peruvian authorities say that McDonald's Latin America franchisee Arcos Dorados violated the law with its working conditions after two employees were electrocuted and killed.	
Source	Haaretz: haaretz.com/world-news/wires/1.8321179	

Table 2: List of Exemplar Firms, by Green/Brown(washing) Category and Sector; Public US Firms.

	Green	Greenwashing	Brown	Brownwashing
Consumer Products	Gap Inc.	Fiat Chrysler Automobiles	Fox Factory Holding	IMAX Corp.
	PVH Corp.	Polaris Industries	Take-Two Interactive Software	Stoneridge Inc.
	Hanesbrands	Deckers Outdoor Corp.	Tile Shop Holdings	Interpublic Group
Energy	Vermilion Energy	Enbridge Inc.	TETRA TECHNOLOGIES	Pason Systems Inc.
	First Solar	Exxon Mobil	Ring Energy	Denison Mines Corp.
	Overseas Shipholding	Occidental Petroleum	Parsley Energy	Dorian LPG
Food & Household	Hormel Foods	Altria Group	Vector Group	Seaboard Corp
	Archer-Daniels-Midland	Philip Morris International	Cal-Maine Foods	Innophos Holdings
	Hershey Co.	Loblaw Companies Ltd.	Tootsie Roll Industries	Regis Corp.
Healthcare	PRA Health Sciences	Johnson & Johnson	Prestige Brands Holdings	Tivity Health
	Agilent Technologies	CVS Health	Tactile Systems Technology	National Healthcare Corp
	IQVIA HOLDINGS	Abbott Laboratories	Nevro	CytomX Therapeutics
Industrial	Steelcase	SNC-Lavalin Group	Omega Flex	Costamare
	CNH Industrial	W.W. Grainger	NL Industries	Kforce
	Xylem Inc.	Union Pacific Corp.	Universal Truckload Services	H&E Equipment Services
Resources	West Fraser Timber Co.	Southern Copper	Olin Corp.	American Vanguard
	Cabot Corp.	Sherritt International	Koppers Holdings	Hawkins, Inc.
	Newmont Mining	Barrick Gold	Myers Industries	Reliance Steel & Alum.
Technology	Workday Inc.	Alphabet Inc.	SINA Corp.	Ceva Inc.
	Virtusa	Cognizant Technology	Applied Optoelectronics	AppFolio Inc.
	Verisk Analytics	National Instruments	j2 Global	ORBCOMM
Utilities	Avangrid Inc.	Duke Energy	Atlantic Power Corp.	Middlesex Water Co.
	CMS Energy	Hydro One	Southwest Gas	SJW Corp.
	NextEra Energy	Dominion Energy	UGI Corp.	Pattern Energy Group

Top firms in each category are those closest to representing the extreme position suggested by the group. For example, the top greenwashing firms in each sector are those firms which have both the most policies on the books, but the worse outcomes reported by news media. Data is based on firm averages from 2016 through 2019.

Table 3: Exemplar Firms, Food and Household Sector

	Firm	Policy Count (0 to 67)	Outcomes (0 to 100)
Green	Hormel Foods	54	72
	Archer-Daniels-Midland	52	72
	Hershey Co.	55	69
Greenwashing	Altria Group	56	44
	Philip Morris International	56	46
	Loblaw Companies Ltd.	48	46
Brown	Pilgrim's Pride	28	42
	Core-Mark Holding Co.	11	63
	Empire Co. Ltd.	33	42
Brownwashing	Seaboard Corp	9	91
	Performance Food Group	14	70
	Illinois Tool Works	31	78

Policies measures a raw count of ESG policies on the books, and ranges from 0 to 65.

Outcomes is a sentiment measure of news media coverage of ESG outcomes, and ranges from 0 (most negative) to 100 (most positive), with 50 being neutral sentiment coverage.

Table 4: Correlation Table

	(1)	(2)	(3)	(4)	(5)	(6)
(1) Total Policies	1					
(2) Policies, normed	0.93	1				
(3) Outcomes	0.02	-0.04	1			
(4) Outcomes, normed	0.01	-0.01	0.94	1		
(5) Revenue, log	0.48	0.41	-0.04	-0.03	1	
(6) Employees, log	0.57	0.54	-0.03	-0.04	0.71	1

Policies are gathered from Thomson Reuters ESG dataset. Outcomes are measured as the material insight score from TruValue Labs, which uses news media coverage of firms' ESG performance. Policy and outcome variables in rows (2) and (4) are normed at the industry-year level. Sample is restricted to firm-year observations with data from Thomson Reuters, TruValue and Compustat.

Table 5: Policy-Outcome Correlations by Sector-Year

	2016	2017	2018	2019
Consumer Products	0.02	-0.01	0.04	-0.02
Energy	-0.08	-0.02	-0.01	-0.07
Food & Household	-0.02	0.03	0.02	0.04
Healthcare	-0.07	0.03	-0.02	-0.02
Industrial	0.05	0.08	0.06	0.05
Resources	0.03	0.05	0.04	0.04
Technology	-0.08	-0.06	-0.10	-0.06
Utilities	-0.04	-0.04	0.01	0.02

Each cell represents the bivariate correlation between ESG policies and ESG outcomes firms in the focal sector and year. Policies are gathered from Thomson Reuters ESG dataset. Outcomes are measured as the material insight score from TruValue Labs, which uses news media coverage of firms' ESG performance.

Table 6: Sample

Year	All Variables For Main Analysis	All Variables And Robustness Checks
2008	102	96
2009	125	118
2010	153	147
2011	179	172
2012	195	186
2013	219	211
2014	229	223
2015	471	461
2016	762	733
2017	905	842
2018	996	901
2019	1,150	931
Total	5,486	5,021

Counts refer to the number of unique firms in each year.

Table 7: Summary Statistics

Statistic	ESG Policies	ESG Outcomes	Earnings Call ESG (SASB)	Earnings Call ESG (Word Count)	Investor ESG Orientation	Long Term Shareholders	Employees, log	Revenue, log	Assets, log
N	4,990	4,990	4,990	4,990	4,990	4,990	4,986	4,990	4,990
Mean	21.3	61.7	0.002	0.552	0.009	-0.134	2.05	7.61	8.28
Std. Dev.	15.0	15.9	0.002	0.357	0.011	0.132	1.46	2.02	1.80
Median	16.0	63.3	0.001	0.463	0.006	-0.134	1.84	7.66	8.16
Minimum	0.0	1.0	0.000	0.000	0.000	-0.633	0.00	0.00	3.86
Maximum	67.0	99.8	0.032	2.650	0.100	0.782	7.70	13.20	14.80

Table 8: Correlation Table (4,032 Observations)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) ESG Policies	1.00							
(2) ESG Outcomes	0.02	1.00						
(3) Earnings Call ESG (SASB)	0.08	0.06	1.00					
(4) Earnings Call ESG (Word Count)	0.11	0.09	0.26	1.00				
(5) Investor ESG Orientation	0.12	-0.01	-0.00	0.03	1.00			
(6) Long Term Shareholders	0.15	-0.03	-0.01	0.00	-0.07	1.00		
(7) Employees, log	0.58	-0.07	0.03	0.06	0.11	0.10	1.00	
(8) Revenue, log	0.57	-0.10	0.04	0.12	0.11	0.09	0.81	1.00
(9) Assets, log	0.60	-0.12	0.01	0.06	0.15	0.14	0.68	0.85

All dependent and independent variables (1 through 7) are normalized at the sector-year level. Control variables (8, 9 and 10) are logged. All values above 0.03 (or below -0.03) have 95% confidence intervals that do not overlap with zero. All values above 0.04 have 99% confidence intervals that do not overlap with zero.

Table 9: ESG Stated Preferences, and ESG Policies & Outcomes: Using Word Count Measure

	(1) Policies	(2) Policies	(3) Outcomes	(4) Outcomes
Managers' ESG Language (Word Count)	0.0396*** [3.64]	0.0125 [0.94]	0.0781*** [5.09]	0.0299 [1.35]
Investors' ESG Orientation	0.0334*** [3.33]	0.0199* [2.46]	0.017 [1.2]	0.00296 [0.22]
Employees, log	0.317*** [19.75]	0.215*** [5.18]	0.0344 [1.52]	-0.13 [-1.88]
Assets, log	0.172*** [12.37]	0.0918** [2.95]	-0.0925*** [-4.72]	0.154** [2.97]
Revenue, log	0.0481*** [3.57]	-0.0224 [-1.08]	0.0245 [1.28]	-0.0573 [-1.66]
Constant	-1.802*** [-7.75]	-1.638*** [-8.53]	-0.633 [-1.93]	-0.486 [-1.52]
<i>Fixed Effects</i>				
Sector	x		x	
Firm		x		x
Year	x	x	x	x
Observations	4,606	4,606	4,606	4,606
R-squared	0.57	0.40	0.13	0.01

Managers' ESG Language is a measure of the proportion of ESG-related words managers use on earnings calls.

Investors' ESG Orientation is a measure of the proportion of a firms' investors' assets under management in funds labeled as sustainable in their prospectus.

Investor ESG Orientation, ESG Policies, and ESG Outcomes are all normed at the sector-year level.

All independent and control variables lagged by 1 year.

T-statistics in brackets. * p<0.05, ** p<0.01, *** p<0.001.

Table 10: Effect of Manager and Investor ESG Preferences on ESG Policies and Outcomes, Including Language Polarity and Subjectivity

	(1) Policies	(2) Policies	(3) Outcomes	(4) Outcomes
Managers' ESG Language (Word Count)	0.0284** [2.7]	0.00389 [0.33]	0.0813*** [5.47]	-0.00454 [-0.23]
Managers' ESG Language, Polarity	0.0510*** [4.01]	-0.00623 [-0.57]	0.0501** [2.79]	0.0599*** [3.31]
Managers' ESG Language, Subjectivity	-0.0491*** [-3.90]	0.00933 [0.85]	-0.0601*** [-3.38]	-0.00511 [-0.28]
Investors' ESG Orientation	0.0339*** [3.37]	0.0210* [2.58]	0.0166 [1.17]	-0.000126 [-0.01]
Employees, log	0.307*** [18.94]	0.143** [3.27]	0.0241 [1.05]	-0.126 [-1.72]
Assets, log	0.172*** [12.28]	0.111*** [3.48]	-0.0932*** [-4.71]	0.161** [3.04]
Revenue, log	0.0527*** [3.89]	-0.013 [-0.62]	0.0289 [1.51]	-0.0628 [-1.82]
Constant	-1.848*** [-7.92]	-1.716*** [-8.74]	-0.703* [-2.13]	-0.508 [-1.56]
Fixed Effects				
Sector	x		x	
Firm		x		x
Year	x	x	x	x
Observations	4,541	4,541	4,541	4,541
R-squared	0.57	0.39	0.13	0.01

Investor ESG Orientation, ESG Policies, and ESG Outcomes are all normed at the sector-year level. All independent and control variables lagged by 1 year.

T-statistics in brackets. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 11: Managers' and Investors' ESG Stated Preferences, and ESG Policies & Outcomes, SASB Semantic Distance Measure

	(1) Policies	(2) Policies	(3) Outcomes	(4) Outcomes
Managers' ESG Language (SASB Semantic Distance)	0.0295** [2.66]	0.0142 [1.48]	0.0474** [3.02]	0.0152 [0.95]
Investors' ESG Orientation	0.0345*** [3.45]	0.0201* [2.48]	0.0187 [1.32]	0.0034 [0.25]
Employees, log	0.315*** [19.62]	0.214*** [5.16]	0.0342 [1.51]	-0.131 [-1.88]
Assets, log	0.169*** [12.25]	0.0918** [2.95]	-0.0990*** [-5.07]	0.155** [2.99]
Revenue, log	0.0525*** [3.91]	-0.0219 [-1.06]	0.0334 [1.76]	-0.0571 [-1.65]
Constant	-1.852*** [-7.96]	-1.639*** [-8.54]	-0.723* [-2.20]	-0.494 [-1.54]
<hr/>				
<i>Fixed Effects</i> Sector	x		x	
Firm		x		x
Year	x	x	x	x
Observations	4,978	4,978	4,978	4,978
R-squared	0.56	0.40	0.12	0.01

Managers' ESG Language is a measure of the semantic distance between their language on earnings calls and sustainability standards (from SASB).

Investors' ESG Orientation is a measure of the proportion of a firms' investors' assets under management in funds labeled as sustainable in their prospectus.

Investor ESG Orientation, ESG Policies, and ESG Outcomes are all normed at the sector-year level.

All independent and control variables lagged by 1 year.

T-statistics in brackets. * p<0.05, ** p<0.01, *** p<0.001.

Table 12: ESG Stated Preferences, and ESG Policies & Outcomes: Using Investor Time Horizons

	(1)	(2)	(3)	(4)
	Policies	Policies	Outcomes	Outcomes
Managers' ESG Language (SASB Semantic Distance)	0.0300** [2.70]	0.0142 [1.48]	0.0470** [3.00]	0.015 [0.94]
Investors' Time Horizon (Long Term Orientation)	0.0356*** [3.33]	-0.0184 [-1.69]	-0.0385* [-2.55]	-0.0334 [-1.84]
Employees, log	0.312*** [19.35]	0.222*** [5.32]	0.0402 [1.77]	-0.122 [-1.76]
Assets, log	0.169*** [12.2]	0.0933** [2.99]	-0.0919*** [-4.70]	0.157** [3.02]
Revenue, log	0.0552*** [4.09]	-0.024 [-1.16]	0.0269 [1.41]	-0.06 [-1.74]
Constant	-1.895*** [-8.15]	-1.650*** [-8.59]	-0.738* [-2.25]	-0.51 [-1.59]
<i>Fixed Effects</i>				
Sector	x		x	
Firm		x		x
Year	x	x	x	x
Observations	4,610	4,610	4,610	4,610
R-squared	0.57	0.40	0.12	0.01

Managers' ESG Language is a measure of the semantic distance between their language on earnings calls and sustainability standards (from SASB).

Investors' ESG Orientation is a measure of the proportion of a firms' shares held by dedicated shareholders minus the proportion held by transient investors, using the Bushee (2001) classification.

Investor ESG Orientation, ESG Policies, and ESG Outcomes are all normed at the sector-year level.

All independent and control variables lagged by 1 year.

T-statistics in brackets. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 13: Ratings Variance and Greenwashing

	(1)	(2)
Greenwashing		0.0345* [2.13]
log Employees	-0.00543 [-0.19]	-0.00683 [-0.24]
log Assets	0.00885 [0.3]	0.00331 [0.11]
log Revenue	0.0695 [1.78]	0.0548 [1.39]
Constant	0.24 [0.8]	0.388 [1.27]
Observations	607	607

Greenwashing is defined as ESG policies (from Thomson Reuters) minus ESG outcomes (independently reported from TruValue, both variables are normed at the sector-year level before differenced). Ratings disagreement is the difference between (year normed) MSCI and Thomson Reuters ESG ratings.

Year and sector fixed effects are included in both models.

t statistics in brackets. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figures

Figure 1: Coefficient of ESG Policies On ESG Outcomes

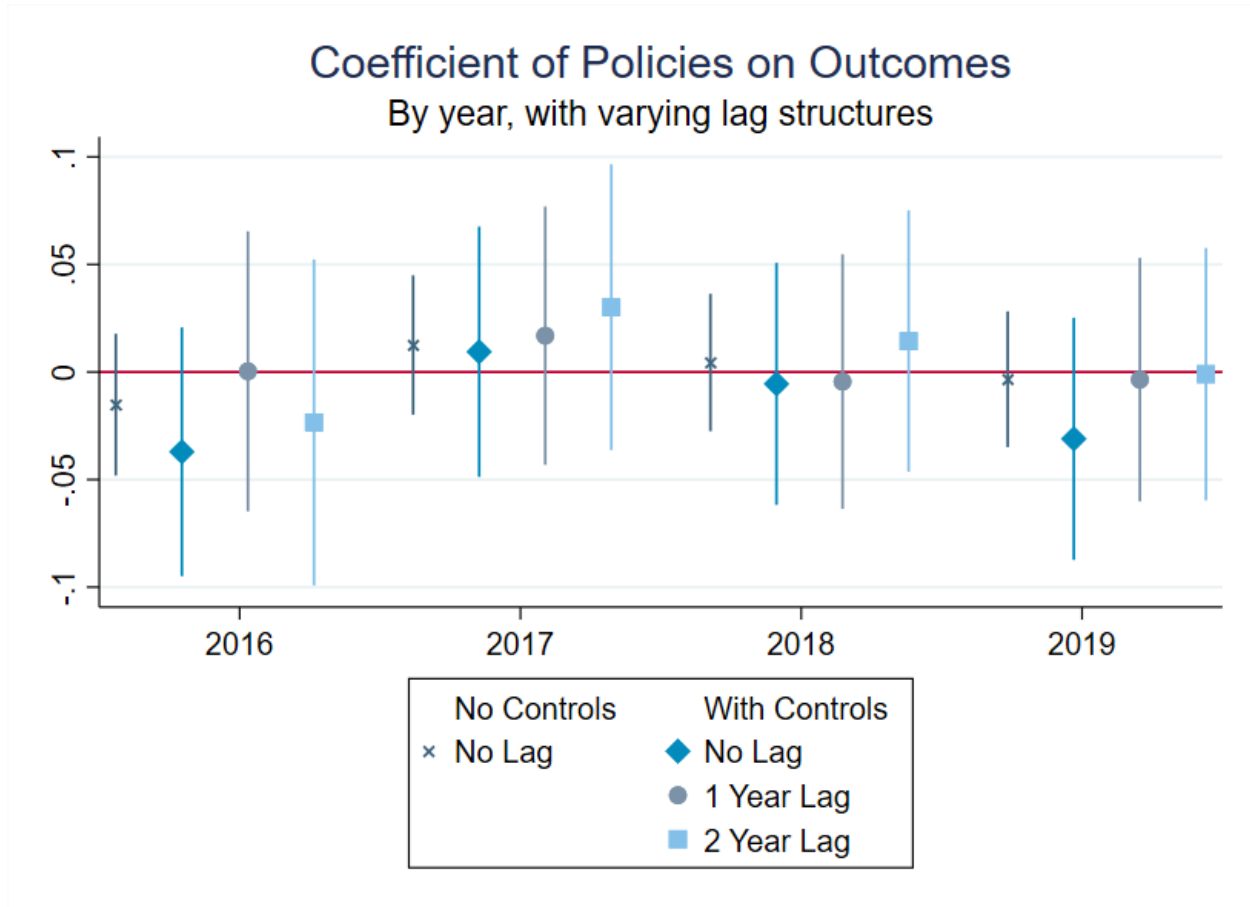


Figure 2: Management’s ESG Language on Earnings Calls (SASB)




Figure 3: Shareholder’s ESG Orientation (Sustainable AUM Percentage)

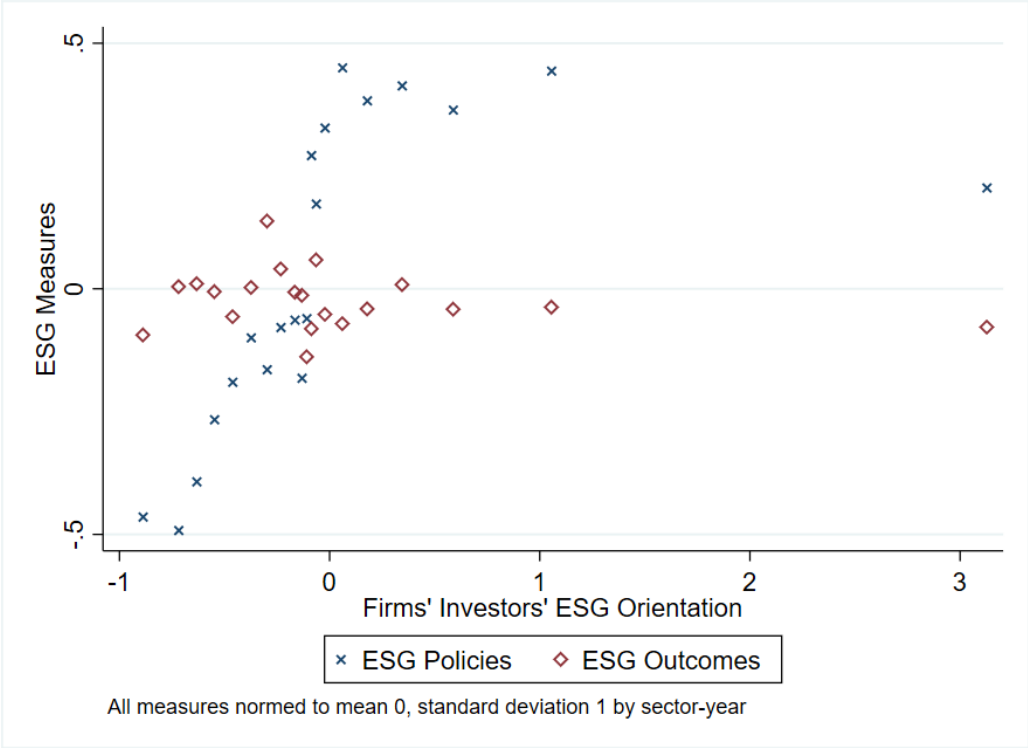



Figure 4: Ratings Disagreement and Greenwashing



Chapter 2: ESG Performance and Labor Productivity

Exploring whether and when ESG affects firm performance¹⁷

Abstract

We explore the relationship between a firm's environmental, social and governance performance (i.e., ESG) and labor productivity. We employ two approaches to examine this relationship: 1) generating firm specific estimates of the relationship between ESG and labor productivity to examine between and within industry variance; and 2) exploring a causal relationship via a quasi-experiment, where regulatory changes led to a shock in the comparability of greenhouse gas emissions data reported. Using random coefficient models, we estimate the distribution of the ESG-labor productivity relationship across a large sample of US firms. We find that the mean effect of ESG on labor productivity is statistically significant and positive in some sectors and significant and negative in others, while the variance of the relationship is always larger than the mean effect and statistically and economically significant. We find this variance is greater within than across sectors, suggesting significant firm-level variation in the returns to ESG. To examine the causal effect of revealing environmental performance, we examine the 2013 UK regulation mandating standardized greenhouse gas emissions reporting for UK firms. We find that treated firms with better (worse) than expected environmental performance have higher (lower) labor productivity post-regulation, relative to both pre-treatment trends and a control group. The combination of the correlational and quasi-experimental studies suggests that positive ESG performance positively impacts labor productivity. These results bear important implications for whether ESG is equally valuable for all firms as well as strong suggestive evidence of the impact of revealed environmental performance on labor productivity.

¹⁷ This chapter is coauthored with Rachelle C. Sampson

Introduction

As record amounts flow into ESG driven funds,¹⁸ questions have been raised about the validity of investing on the basis of ESG ratings. Critics argue that ESG (i.e., metrics that capture a firm’s environmental, social and governance profile) amounts to, “labelling based on incomplete information, public shaming, and shunning wrapped in moral rhetoric,”¹⁹ implying that currently reported ESG information is shoddy at best and bears little correlation with actual firm performance, even on the dimensions that such measures propose to capture (see, e.g., Chatterji, Levine and Toffel, 2009). Concerns have also been raised that the pursuit of ESG strategies by pension funds, for example, could undermine a money manager’s duty to act in the best interests of their clients by providing a return sufficient to fund retirement.²⁰

In contrast, a growing literature examines the impact of a firm’s ESG information on financial metrics such as stock returns and cost of capital, providing suggestive evidence of the positive impact of environmental and social initiatives on firm performance (e.g., Serafeim, 2018; Goss and Roberts, 2011; Deng, Kang and Low, 2013). Despite this burgeoning literature, we have little large-scale empirical evidence that helps define or capture the mechanisms by which environmental and social initiatives impact financial performance. That is, why does positive ESG performance correlate with better stock performance? What do investors expect to happen that will lead to increased value from ESG? Further, and perhaps more importantly, is ESG equally valuable for all firms? Answers to these questions may help guide firms towards more discerning investment in environmental and social initiatives by highlighting better when and what type of initiatives make most sense for firms. This is essential for making a more sophisticated business case for

¹⁸ Morningstar reports an increase of flows into socially responsible investment funds to \$17.67B in 2019 from \$2.83B in 2015 (<https://www.morningstar.com/articles/959379/10-sustainable-investing-stories-of-2019>). See also Białkowski and Starks (2016) and Hartzmark and Sussman (2019).

¹⁹ Hester Peirce, US SEC commissioner, as quoted in “ESG Investing: Poor Scores,” *The Economist*, 7 December 2019.

²⁰ See, for example, “SEC Questions Do-Good Funds,” *Wall Street Journal*, 17 December 2019.

environmental sustainability and social performance both within the firm and to outside stakeholders.

In this paper, we examine one such mechanism that may link firm environmental and social performance to financial performance: the relationship between ESG and labor productivity. Firms with positive ESG profiles may be more selective in hiring and better able to retain and engage employees, where employees share these firm values, leading to higher levels of effort and higher labor productivity. However, this relationship is not assured. If firms pursue a strategy of increasing environmental and social performance, they may shift productive resources such as capital, labor, or managerial effort toward the production of positive ESG outcomes, decreasing resources available for revenue producing activities (Becker 2011). Given these countervailing theoretical arguments and existing research suggesting a positive, but small and inconsistent mean effect between ESG performance and firm financial performance, we aim to provide evidence on when and for which firms ESG performance impacts labor productivity.

We employ two approaches to address this agenda. First, we examine the relationship between ESG performance and labor productivity in a large sample of international firms over the 2008-2018 time period. Second, we examine how an exogenous shock to the visibility and comparability of a firm's greenhouse gas emissions affects subsequent labor productivity.

Our work expands previous research by focusing on labor productivity, using independently reported ESG events, and examining the distribution of the ESG-financial performance relationships, rather than the mean effect. We find that while the mean relationship between ESG performance and labor productivity is statistically significant in some industries and not others, the variance of this relationship is statistically significant and much larger than the mean effect across our whole sample. This suggests that while on average, there is a small or insignificant correlation between ESG and labor productivity, for the median individual firm, this relationship is significant –

positive for some firms and negative for others. We also find that large firms have generally positive ESG-labor productivity relationships, while small firms have generally negative ESG-labor productivity relationships.

To go beyond this correlational study and identify potential causality, we study how a 2013 regulatory change in the United Kingdom mandating standardized greenhouse gas (GHG) emissions reporting impacted the labor productivity of high emitting and low emitting UK firms, relative to a control group of similar European firms. We find that when GHG emissions are costly to observe (i.e., no regulatory requirements for reporting or standardization), there is no relationship between ESG performance (emissions intensity) and labor productivity, but when standardized firm emissions become public information, in the year of and after the regulation, high-emitting UK firms show significantly lower labor productivity than both low emitting UK firms and a control group of European firms.

The results from both approaches suggest that over the 2008-2018 timeframe there was at most a very small mean relationship of environmental and social performance with labor productivity, but this relationship varies significantly with both firm characteristics and institutional characteristics, specifically the publicity of standardized and comparable ESG outcome data.

In the sections below, we briefly review research that examines how ESG affects firm financial performance and labor productivity. Descriptions of our datasets, empirical approaches and results follow. We conclude with a discussion of the implications of our results not only for the direction of future research and fruitful areas for exploration, but also for how firms might develop a more focused approach to environmental and social strategy.

Environmental, Social, and Financial Performance

The rising importance of ESG²¹ can be seen clearly through the significant increase in inflows to equity mutual funds with mandates to invest according to ESG performance, relative to those without such mandates (Bialkowski and Starks, 2016). In the US mutual fund market, Hartzmark and Sussman (2019) find that being categorized as a low sustainability fund led to net outflows of more than \$12B, while a high sustainability categorization yielded more than \$24B worth of inflows. These flows reflect the increasing demands by investors for accountability around environmental issues such as climate change and social issues such as fair and safe labor practices.²² Companies, in turn, are under increasing pressure from shareholders, customers and other stakeholders to report on ESG performance, in response to emergent environmental and social crises as well as the inherent risks to firm performance that such crises represent. Reflecting these trends, recent statements by the Business Roundtable that the social responsibility of business is to all stakeholders, not just shareholders,²³ suggests that firms are increasingly aware of the importance of business practices that underlie ESG metrics.

A growing body of recent research has investigated the financial performance implications of positive ESG performance, that is, being ranked above industry peers on environmental, social and governance metrics. In a review of 167 studies, Margolis, Elfenbein and Walsh (2007) find a modest positive correlation between corporate social (or ESG) performance and financial performance, though this correlation is smaller in studies with basic industry and firm controls. More recently,

²¹ We use ‘ESG’ to capture metrics that measure a firm’s environmental, social and governance profile, similar (though not identical) to what was more commonly called ‘CSR’ or corporate social responsibility in past years. While we use the acronym ESG in our text, we largely focus our empirical analysis below on environmental and social issues.

²² See, e.g., the ‘Joint Statement Between Institutional Investors on Behalf of Climate Action 100+ and Royal Dutch Shell plc (Shell),’ issued 3 December 2018, https://www.shell.com/media/news-and-media-releases/2018/joint-statement-between-institutional-investors-on-behalf-of-climate-action-and-shell/_jcr_content/par/textimage_d010.stream/1543782013771/767edfadfd44d01e425d68a17fd5a1fd24f4f32b/03-dec-2018-joint-statement-institutional-investors-and-shell.pdf

²³ See Business Roundtable’s August 2019 ‘Statement on the Purpose of a Corporation,’ <https://opportunity.businessroundtable.org/wp-content/uploads/2019/12/BRT-Statement-on-the-Purpose-of-a-Corporation-with-Signatures.pdf>

Khan, Serafeim and Yoon (2016) and Nagy, Kassam and Lee (2016) find that positive changes in ESG are correlated with positive abnormal stock returns.

However, as the modest correlation in the meta-analysis suggests, the ESG-financial performance relationship appears to be conditional. Many studies have found, for example, that better environmental performance does not motivate innovation, nor does it guarantee greater competitiveness (Jaffe, Peterson, Portney and Stavins, 1995; Ambec, Cohen, Elgie and Lanoie, 2013). In their extensive review of the literature, Kitzmueller and Shimshank (2012) argue that demand for corporate social responsibility is not universal. Cheng et al (2014) find evidence consistent with managers' personal interests driving ESG behavior, suggesting ESG can be an agency problem. This research on the ESG-financial performance relationship suggests there is significant conditionality in the relationship, so we now turn to our mechanism of interest linking the two: ESG's effect on labor productivity.

ESG And Labor Productivity

The role of human capital in firm productivity has changed dramatically over the past several decades, becoming relatively more important as more firms focus on innovation (e.g., Zingales, 2000). Bhaskarabhatla, Hedge and Peters (2019) find that highly talented inventors are attracted to firms with other talented inventors but weak firm-specific innovation capabilities, concluding that human capital, rather than firm capabilities, may be most important for firm innovation. This implies that attracting and retaining high skill workers is essential for performance of firms focused on innovation. Retention of workers is a key issue even for firms not focused on innovation, due to the costs associated with recruiting, selecting, and training new employees; Cascio (2006) reports that such costs often exceed 100% of the annual salary for the position being filled.²⁴ Therefore, to the

²⁴ Consistent with this, using store data from a retail chain, Ton and Huckman (2008) find that employee turnover is associated with lower profit margins. Employee turnover also has an important influence on organizational

extent that ESG affects a firm's ability to select, retain and motivate employees, it will affect labor productivity.

Recent work has directly analyzed the link between firm environmental practices, one component of ESG metrics, and labor productivity. Delmas and Pekovic (2012) find that firms that adopt environmental standards, such as ISO 14001, experience higher labor productivity. They posit that this higher productivity stems from greater employee commitment to the firm as well as increased training that may arise from adoption of environmental standards (Ibid.; Ambec and Lanoie, 2008). Similarly, Lannelongue, Gonzalez-Benito and Quiroz (2017) find that environmental management is positively correlated with labor productivity in firms, but only those with low capital intensity, arguing that such a positive relationship arises out of social identity with the organization and the attractiveness as an employer that results from good environmental practices. We expect that ESG metrics, which include but go beyond signals of environmental practices, may impact labor productivity through mechanisms that improve the ability for the firm to be more selective in hiring, reduce turnover and/or induce greater effort by employees.

Firms with good ESG performance may be perceived as 'better places to work' and therefore may be more selective in hiring. Turban and Greening (1996) find that corporate social responsibility ratings (the precursor to today's ESG metrics) are positively correlated with an organization's reputation and attractiveness as an employer.²⁵ Positive ESG performance may also contribute to perceptions of firm integrity among employees; integrity has been shown to be positively correlated with productivity and greater attractiveness of that firm to prospective job applicants (Guiso,

productivity, since turnover leads to a direct loss of firm-specific training and disruption of within-firm work groups (Brown & Medoff, 1978).

²⁵ Several studies point to the link between ESG (or CSR) performance and a firm's attractiveness as an employer, including Backhaus, Stone and Heiner (2002), Greening and Turban (2000) and Sen, Bhattacharya and Korschun (2006).

Sapienza and Zingales 2015). Consistent with these findings, Burbano (2016) reports that workers submit lower wage bids for jobs with organizations that are deemed socially responsible.

Good environmental and social performance may also lead to increased employee satisfaction, engagement and identification with the firm, thus, improving productivity, via retention and formal or informal on-the-job training that comes with tenure (Lichtenberg, 1981; Bartel, 1994)²⁶ or via inducing greater effort by employees. Peterson (2004) finds that employees report better work attitudes and greater commitment to the firm when that firm is perceived to be a good corporate citizen. Similarly, job embeddedness rises with perceptions of fit between an employee's values and corporate culture, a form of satisfaction, which can help improve retention (Mitchell, Holtom, Lee, Sablynski and Erez, 2001). In their meta-analysis of factors that influence turnover, Allen, Bryant and Vardaman (2010) find that two of the factors that bear the strongest negative correlations with turnover are the employee's job satisfaction and organizational commitment. In another meta-analysis of over seven thousand business units, employee satisfaction and engagement were found to be positively correlated with productivity, retention, and employee safety, among other variables (Harter, Hayes and Schmidt, 2002). Employee satisfaction has been linked with financial performance (Edmans, 2011, 2012); one potential mechanism to explain this improvement in financial performance is via an impact on labor productivity.

Finally, ESG metrics may also signal human resource practices within the firm (i.e., since many of the 'social' components capture elements relevant to labor), which have been shown to affect labor productivity. Past research indicates that human resource management practices are correlated with lower employee turnover (e.g., Huselid, 1995) and higher labor productivity (e.g., Ichniowski,

²⁶ Lichtenberg (1981) points out that length of service to a firm is an important determinant of labor productivity, usually as a result of firm specific on-the-job training. Training, which is generally linked to labor productivity growth (Bartel, 1994), is relatively more important for productivity growth in industries with higher human capital (Sepulveda, 2009), a point we take up in our empirical analysis below.

1990). To the extent that ESG factors also cover human resource practices, such as fair labor practices, compensation and benefits, accident and safety management, and efforts toward recruiting, development and retention, ESG performance on such labor related issues should also increase retention and, consequently, labor productivity.²⁷

However, just as in the ESG-overall financial performance domain, we expect these potential benefits to be accompanied by the costs of implementing ESG practices. We address this first by (i) using a random coefficient model to provide descriptive statistics on how this ESG-firm financial performance relationships differs across and within industries and (ii) examining how a regulatory shock changes the returns to ESG performance.

Study 1: Distribution of ESG-Labor Productivity Relationship

Our first study examines the relationship between ESG and labor productivity across a large sample of US firms from 2008 to 2018, inclusive. We examine correlations and firm heterogeneity using a variety of techniques described below. To construct our sample, we combine environmental and social performance data from TruValue Labs with financial and employee data for publicly listed US firms from Compustat. We require firms to have 5 years of data and a minimum coverage of firm ESG issues (more details below), providing us with a sample of 13,313 US firm-year observations. Our sample construction is outlined in Table 15.

[Insert Table 15 here.]

Our US sample ranges from 896 to 1,542 firms per year. We split our sample into industry sectors for our main specification. Around half of our sample firms are manufacturing firms; the sector breakdown of our firms is listed in Table 16.

²⁷ Huselid (1995:645) uses thirteen items to capture high performance work practices, consistent with prior research, which include practices in the areas of: personnel selection, performance appraisal, incentive compensation, job design, grievance procedures, information sharing, attitude assessment, labor-management participation, intensity of recruiting efforts, average number of training hours and promotion criteria. Several of these criteria map onto ESG metrics captured in our data below.

[Insert Table 16 here.]

Dependent Variable: Labor Productivity

Our dependent variable is labor productivity, defined as the log of revenues per employee,²⁸ a widely used measure across literatures. In the human resource management literature, Ichniowski (1990) uses logged sales per employee as the dependent variable to measure the effect of human resource practices on employee productivity. Koch and McGrath (1996) also use revenues divided by employees as their dependent variable to study the effect of modern human resource management strategies. In the environmental domain, Lannelongue et al (2017) use logged sales per employee to measure the effect of CO2 emissions on labor productivity. Revenues-based productivity measures are also commonly used in economics (Foster et al 2008).

Focal independent variable: ESG performance

Our main independent variable is ESG performance, which captures firm level environmental, social and governance performance. We use data from TruValue Labs (TVL), which measures ESG performance by analyzing and coding independent reporting of firm activities and events. While a relatively new data source, this ESG data has already been found to condition investor pricing of ESG ratings (Serafeim 2018) and affect credit prices in infrastructure projects (Henisz and McGlinch 2019). TruValue Labs gathers news media reports from thousands of news sources daily, identifies which articles cover relevant ESG issues in a firm, and codes articles as having positive or negative sentiment using natural language processing algorithms. For example, an article in Vanity Fair alleging Amazon threatened to fire employees for criticizing the company's environmental practices was coded as having negative sentiment and categorized under Labor Practices. In contrast, Duke Energy's announcement that it will start developing LEED certified buildings was coded as having

²⁸ We drop observations where revenues are negative or there are fewer than 5 employees in the firm. For all variables using log transformations we add one before taking the log.

positive sentiment and categorized under Energy Management. Additional illustrations of the types of articles that go into the TruValue Labs metrics are set out in Table 14.

[Table 14 here.]

A key advantage of this data is that it relies on information from outside the firm (e.g., news reports and information from government regulators, NGOs and analysts) as opposed to data reported by the firm itself. Given the lack of ESG standards and that such reporting is voluntary over our sample period, the quality of firm-provided ESG data has been questioned. Berg, Koelbel and Rigobon (2019) show that ESG ratings based on this self-reported ESG information diverge widely based on who is rating the firm's information. Highlighting the importance of gathering non-financial information from independent sources, Guiso et al (2015) find that firms' advertised values have no effect on performance, while employees' perceptions of firm values positively affect performance. The TVL data corresponds more directly to perceived values, since this data captures the sentiment of outsiders to firm events, whereas self-reported firm data on ESG corresponds more closely to the advertised values of the firm.

TVL aggregates individual article sentiment scores into an overall ESG 'Insight' score for each firm using an exponentially weighted moving average of all of a firm's articles' sentiment scores. (Articles have a half-life of six months.) We take this ESG score to be a measure of a firm's overall ESG performance and use this as our main independent variable in our empirical specification.

We also use disaggregated measures that capture environmental, social, and labor performance separately. Each article in the TVL database is classified according to one or more categories, producing ESG scores within each category, e.g., Labor Practices or Greenhouse Gas Emissions. We aggregate these individual category scores into omnibus environmental, social, and labor scores. For example, the environmental score is the average of 12 categories related to a firm's environmental performance, including greenhouse gas emissions, materials sourcing, and

water/wastewater management. While environmental and social omnibus scores are standard in the ESG community, we add a labor score as a subset of the social score due to our focus on labor productivity.

Different environmental and social issues are materially relevant to different industries. For example, biodiversity impacts are relevant to agricultural production, but not to e-commerce retailers. The Sustainability Accounting Standards Board (SASB) has developed a set of standards indicating which sets of environmental and social issues are material for which industries, including the previous examples. ‘Material’ data refers to firm information that would be deemed to be important in investment decisions by reasonable investors, according to the SEC. Using these SASB standards, Khan et al (2016) found that firms performing well on material ESG issues outperform poor performers on material ESG issues, but there is no equivalent effect from performance nonmaterial ESG issues. We use only news articles from material categories, as determined by SASB, when generating our measure of ESG performance.

Firms vary significantly in the volume of news media coverage they receive. While some firms receive no coverage in certain years of our sample, Alphabet, Microsoft, and Tesla all had over 2,000 articles covering ESG issues written about them in at least one year. Firms with very few articles in the dataset more likely have noisy measures of ESG performance; thus, we include only firms with a minimum of 3 articles per year to eliminate those with the lowest coverage and noisiest data.

ESG performance is measured on a 0-100 scale, with higher numbers representing positive ESG performance and lower numbers representing worse performance. A score of 50 is neutral performance. In our US sample, material ESG scores range from 1 to 99 with a mean of 63, indicating that the news media on average reports slightly positive news about firms’ ESG performance. Firms’ ESG scores vary by an average of 8 points over a 3-year period and 12 points over a 5-year period. These variances, our sample restrictions on minimum news coverage, and close

analysis of a subset of articles, give us confidence that variance in our ESG measure captures real news events around firms' ESG behavior and outcomes.

One potential concern is that changes in ESG scores reflect time trends characterizing industries or society rather than individual firms. In unreported results, we find that mean industry sector ESG scores are slightly increasing over time. While an interesting avenue for future study, we remove time trends and cross-industry differences in our main specification by running separate industry level regressions with year fixed effects.

Control variables

Capital intensity: Previous literature suggests that capital intensity influences the relationship between ESG performance and labor productivity. For example, using emissions data, Lannelongue, Gonzalez-Benito and Quiroz (2017) find that environmental performance is positively correlated with labor productivity in low capital intensity firms, and negatively correlated with labor productivity in high capital intensity firms. We define capital intensity as the log of assets divided by employees and include capital intensity in our main specification.

HHI: Labor productivity correlates with industry structure. Symeonidis (2008) finds that productivity is lower in highly concentrated industries in the UK. In our US sample we generate firm-year market shares within public firms as firm sales divided by four-digit SIC sales and create an industry-level Herfindahl Index by summing the square of all firm market shares in each industry, respectively. Descriptive statistics for our US sample are listed in Table 17.

[Insert Table 17 here.]

Estimation approach

Our main estimation strategy uses a random coefficient model (RCM) to estimate firm-specific effects of ESG on labor productivity. While common in economics, education, biostatistics and

other fields, RCMs are novel but growing in the strategy field.²⁹ We regress labor productivity on ESG performance, capital intensity, and industry concentration, and set all independent variables and our intercept to be free parameters.³⁰ This model estimates both a population mean coefficient for each variable, as in OLS, and a variance of the coefficient. The model then estimates firm-specific coefficients of the effect of, for example, ESG performance on labor productivity, that follow distributional assumptions generated from the data (Alcacer et al, 2018). We estimate a firm-specific ESG-labor productivity relationship for each firm in our sample. We do not lag any variables in our estimation. Our initial goal is examining the distribution of firm-specific correlations between ESG and labor productivity, and not to make a causal argument. In robustness checks not shown here, we find consistent results lagging ESG variables by one year.

Study 1: Results

Our first model examines the effects of ESG performance on labor productivity across sectors for US firms. Table 18 shows regression results for both fixed effects models (i.e., firm and time effects) and random coefficient models with sector effects using overall ESG performance and environmental/social/labor performance separately as independent variables.

In order to report summary results from our fixed effects and random coefficient models in a limited number of regressions, our initial models in Table 18 include sector-level controls rather than separate sector-level regressions. In both the fixed effects and random coefficient models in Table 18, there is no significant mean effect of ESG on labor productivity. However, there is significant variance in the relationship in the random coefficient models. This means that there is a significant distribution of both positive and negative ESG-labor productivity relationships across our sample of firms.

²⁹ See Alcacer, Chung, Hawk and Pacheco-de-Almeida (2018) for a detailed explanation of RCM econometrics and examples of recent papers in strategy that use RCMs.

³⁰ We also include sector and year fixed effects.

[Insert Table 18 here.]

To better capture expected differences in the distributions of relationships between ESG performance and labor productivity across sectors, we run separate random coefficient models for each sector. We run these sector level regressions for ESG, environmental, social, and labor performance, resulting in a total of 32 regressions (8 sectors, 4 ESG measures). To report these results concisely, we graph each model's mean and range of firm-specific coefficients. Figure 5 shows the distribution of firm-specific coefficients, which we refer to as 'firm estimates'. The diamond represents the mean coefficient of ESG events on labor productivity within the industry and the bar with 'x' symbols at the ends represents the range between the 5th and 95th percentile of firm specific estimates in that sector. That is, the bar is not a confidence interval, but represents the 5th to 95th percentile range of actual firm specific estimates generated using the random coefficient models. For all sector-level regressions where the mean coefficient on ESG performance was not significantly different from zero, we place the diamond representing the mean effect at zero. Thus, any mean estimate that differs from zero in the figures is statistically significant. For all sector-level regressions where the variance of the coefficient is not significant, we plot only the mean effect (i.e., the diamond shape only). Where variance is significant, we plot the range of estimates as a bar; there are four bars for each sector. These bars capture the range of firm relationships between labor productivity and: (1) the omnibus ESG measure (first bar for each sector in blue); (2) environmental categories only (second bar in green); (3) social categories only (third bar in orange); and (4) labor subset of social categories only (fourth bar in red).

[Insert Figure 5 Here]

The mean effect is statistically different than zero in our sample in the majority of regressions, sometimes positive and sometimes negative. The variance in the ESG-labor productivity coefficient is statistically significant, economically meaningful, and much larger than the mean effect.

The range of firm-effects varies within and across industries as well as between ESG categories (i.e., all ESG performance, or only environment, social or labor). Environmental events have more heterogeneous and larger correlations with labor productivity than the social, labor, or overall ESG measures. While in the case of the labor category, which has fewer data points and observations than the other measures, this may be due to sample size restrictions, comparison of estimates from the environmental versus the overall ESG metric is consistent with environmental issues having a more significant relationship (both positive and negative) with labor productivity. These results are also consistent with the signal to noise ratio in the environmental measure being higher than in the non-environmental measures. Manufacturing and mining firm estimates show the broadest distribution of the relationship between labor productivity and overall ESG and the environmental specific ESG metric (i.e., most significant within sector, between firm, differences in this relationship).

We next split our sample into an early period (2008-2013) and a late period (2013-2018), to better allow our firm specific estimates to vary over time. The importance of environmental and social factors to firms and their stakeholders has increased over time (at least in awareness) and we anticipate that the impact of a firm's ESG performance on labor productivity may similarly change over time. In the same fashion as Figure 5, we plot the mean and range of firm specific estimates by industry according to these two time periods in Figure 6 and Figure 7. Consistent with our expectation, we find that in the later period the variance of the relationship is larger, and the mean effect is more often significant. While not the subject of our analysis here, this observation is consistent with the rising focus of investors and other stakeholders on firm environmental and social performance.

[Insert Figure 6 and Figure 7 here.]

Observing the significant heterogeneity between firms, we perform a supplemental analysis to better understand why some firms have a positive versus negative relationship between ESG and productivity. We expect larger firms to face significantly different tradeoffs from ESG performance than small firms. For example, large firms may have more slack resources available, reducing the cost of ESG related investments and may also experience increased returns to reputation, increasing the payoffs of ESG performance. Conversely, small firms may face higher tradeoffs around ESG performance. We find that firm estimates (i.e., the firm-specific relationship between the omnibus ESG performance measure and labor productivity) correlate positively with two measures of firm size: $\log(\text{revenues})$ (significant at the 1% level) and $\log(\text{assets})$ (marginally significant, p-value of 0.07). Figure 8 shows these relationships using binned scatterplots.

[Figure 8 here.]

Study 2: Impact of Revealed Environmental Performance on Labor Productivity

Our large-scale analysis of the relationship between ESG performance and labor productivity above shows significant variance across and within industries. We find that there are firms for which high ESG performance is correlated with high labor productivity, and firms for which high ESG performance is correlated with low labor productivity. However, these patterns are correlations and do not allow us to make any claims about the causal impact of environmental or social performance on labor productivity.

One approach to test for the causal effect of ESG performance on labor productivity is by examining a shock which changes the information revealed to management and employees, while not changing actual ESG performance. This information path can provide insight into how individuals react to ESG performance. We identify one such regulatory change in the UK in 2013 which provides an exogenous change to the information available around one significant aspect of

ESG performance: greenhouse gas (GHG) emissions. Specifically, this regulation both mandates and standardizes the reporting requirements for GHG emissions. This standardization allows direct comparison of firms within and across sectors that was largely impossible prior to the 2013 regulation. The shock we use as a quasi-experimental study is the impact of revealed, comparative data on GHG emissions by firms that were already reporting GHG emissions, albeit in a non-standardized format. In our analyses below, we analyze how the relationship between excess GHG emissions and labor productivity changes as a result of the regulation and subsequent standardization of reporting.

UK Companies Act Regulation 2013

The UK Companies Act Regulation 2013³¹ required all publicly listed firms to report their carbon emissions, beginning in the 2013 fiscal year. This law provides an exogenous shock to GHG emissions reporting for UK firms that did not previously report standardized GHG emissions. This event allows us to use a difference-in-difference approach for making causal claims about how ESG information revelation affects labor productivity. With respect to this specific regulatory change, Jouevenot and Krueger (2020:12) note:

“What makes the U.K. GHG emissions disclosure regulation unique is that, in contrast to other disclosure regulations concerning non-financial information (e.g., the EU Directive on non-financial disclosures...), the U.K. regulation is prescriptive about (1) what kind of information needs to be disclosed, (2) how this information needs to be presented, and (3) where the information needs to be disclosed. In other words, the regulation is mandatory, prescriptive, and standardized.”

Initial research has found real changes in firm behavior as a result of this regulation. Jouvenot and Krueger (2020) find that UK firms subject to the regulation reduce their GHG emissions and emissions intensities by 14% and 17%, respectively, relative to other European firms post regulation. Furthermore, firms that report low emissions intensities see a rise in their stock price on the day they

³¹ The official legislation is the UK Companies Act 2006 (Strategic Report and Directors' Report) Regulations 2013. The 2013 regulation is an update to the 2006 law.

report emissions in 2013. Our analysis takes a different approach and analyzes responses in labor productivity to the improved information on a firm's carbon emissions that allows employees and other stakeholders to observe relative performance.

The UK Companies Act Regulations 2013 requires all firms listed on the London Stock Exchange to report their scope 1 and scope 2 greenhouse gas emissions, as well as their emissions intensity and methodology for calculating emissions. The law was planned over 2011-2012 and was passed and went into effect in 2013. A brief timeline of the law's development is listed below:

UK Companies Act 2006, Regulation 2013 Timeline

- | | |
|--------------|--|
| 2006 | The Companies Act of 2006 is passed. The law is the primary source of company law. No emissions reporting requirements are included. |
| 2008 | Climate Change Act of 2008 passed. The Act requires that, by 2012, the government write a law requiring directors of some companies to report on GHG emissions in some manner. Which companies must report and how and what these companies will report is left for future debate. |
| Aug 2011 | The UK government solicits feedback on potential versions of a GHG emissions law beginning in early 2011. No reporting, voluntary and mandatory reporting are all discussed. By August 2011, the government recommends the law mandating GHG emissions reporting by listed companies. At this point, firms listed on the London Stock Exchange can reasonably expect to be mandated to report carbon emissions beginning sometime in 2012 or 2013. |
| 2012 | Regulation is drafted, firms can expect to begin standardized reporting in 2013. |
| Jul 16, 2013 | Regulation is approved by the House of Commons. |
| Oct 1, 2013 | Regulation goes into effect; firms must report their GHG emissions by this date. |
-

Empirical Approach

Based on the development of the law over the 2011-2013 time period, we expect firms listed on the London Stock Exchange could predict that they would be required to report standardized GHG emissions for the 2013 fiscal year. While the law did not pass until mid-2013, the UK government recommended mandatory reporting for listed firms in mid-2011 and the regulation was drafted by

mid-2012. Accordingly, we expect firms to respond to the regulation in 2013, and we expect to see financial effects from the regulation beginning that same year. We focus on firms that were already reporting GHG emissions, but in a non-standardized format prior to 2013. Thus, the 2013 regulatory change is not leading the firms we focus on to report GHG emissions for the first time, but rather to standardize their reporting.

As noted above, the key aspect of this regulation was that it requires firms to report standardized and comparable GHG emissions data. Before the regulation, most firms listed on the London Stock Exchange reported some GHG emissions data in their CSR reporting or in their Directors' Reports. However, while comparable data was collected and reported to ESG data collectors such as the Carbon Disclosure Project, firms publicly reported selective and non-comparable GHG emissions data points. For example, in 2012, firms may have reported that they reduced CO₂ output by 4% from the previous year or reduced total emissions from a baseline many years back. Some firms chose to exclude new acquisitions from their emission calculations. As a consequence of the regulation, beginning in 2013 all firms had to report the absolute Scope 1 and Scope 2 emissions in their publicly available and easily accessible Directors' Reports. This change allowed stakeholders including employees to make comparison across firms at a much lower cost.

We are interested in the financial effects (i.e., revenues-based labor productivity) of this requirement to disclose comparable emissions data. We expect large firms, firms with high levels of production, and capital-intensive firms to have different levels of emissions than their smaller, less output-intensive or capital-intensive counterparts. We also expect different industries to have different levels of absolute emissions and emission intensities. As such, similar to King and Lenox (2000), we use an emissions production function to estimate a predicted level of emissions for each firm. We then generate a measure of excess emissions for each firm-year observation by subtracting this predicted emissions value from their actual reported emissions.

The 2013 regulation should make excess emissions more salient to stakeholders, including employees. Before the regulation, accessing comparable emissions data was costly or simply inaccessible for most stakeholders. Given stakeholders' lack of information, we expect any pre-regulation relationship between excess emissions and labor productivity to reflect operational efficiencies or managerial choices between GHG emissions and revenues, and not employee or consumer responses to firms' relative emissions levels. The mean relationship between ESG and labor productivity in our larger analysis was close to zero in many industries and, as such, we conservatively expect there to be no significant relationship between GHG emissions and labor productivity before the regulation goes into effect.

When the regulation goes into effect, we expect there to be a shock to the information available to employees, consumers, managers, and other stakeholders. Given this shock to the information available and the negative connotation of excess emissions, we expect the excess emissions residual to have a negative relationship with labor productivity for firms subject to mandatory GHG reporting (i.e., UK incorporated firms) beginning in 2013.

Data and Sample Description

Our treatment sample consists of UK incorporated firms listed on the London Stock Exchange. Our control sample consists of firms listed on other major European stock exchanges. We gather financial data from Thomson Reuters Worldscope. We gather greenhouse gas emissions data from Thomson Reuters ESG (formerly Refinitiv) and supplement this data with GHG emissions data from Bloomberg.

Our final sample consists of 152 treated (UK) firms and 319 control (European) firms. Our European firms come from 16 European stock exchanges, with the majority coming from Paris, Frankfurt, Milan, Madrid, Zurich, Stockholm, and Helsinki, in that order. The stock exchanges we use in our sample are listed in Table 19.

[Insert Table 19 here.]

We require firms to have GHG emissions data for all years between 2010 and 2015 inclusive to be included in our sample. Importantly, we are able to observe standardized GHG emissions data pre-2013, while employees and other stakeholders were not able to in those years. We are able to observe standardized GHG emissions pre-2013 for those firms reporting to the Carbon Disclosure Project (CDP). While such reporting was not mandatory, nor audited, the reporting to CDP was standardized. We can now observe those reports via Thomson Reuters data, but note that this data was not disclosed at the firm level by CDP to the general public (and still is not). While it is possible that firms would report inaccurate information to the CDP, given the private nature of the disclosure, the incentive to falsely report is reduced. Thus, all of the firms in our treatment sample collected GHG emissions information before 2013. However, these firms were not legally required to report these emissions publicly in such a standardized format. Employees (as well as consumers and other stakeholders) would not have been able to access this comparable GHG emissions data for the vast majority of British firms until the UK regulation made such reporting mandatory in 2013.

By including only firms with GHG information for all years between 2010 and 2015 we keep a consistent sample of firms pre- and post-regulation. This allows us to isolate the effects of the revelation of comparable ESG information to employees, removing any influence first time reporting firms would have on our sample. This addresses our concern that firms which report for the first time in 2012 or 2013 may have different labor productivity patterns than firms that consistently reported emissions before the regulation was drafted.

Our main measure of GHG emissions is the sum of scope 1 and scope 2 emissions, what we call total emissions or simply emissions for the remainder of the paper. Scope 1 emissions are emissions from sources that are owned or controlled by the company (e.g., boilers, furnaces or company-

owned vehicles). Scope 2 emissions are emissions from electricity purchased and consumed by the company (e.g., emissions produced by generating electricity purchased by the focal firm to power their buildings). We use the sum of scope 1 and scope 2 emissions because these are the two types of emissions firms are required to report according to the UK regulation.

To generate our focal independent variable, ‘Excess Emissions’, we first use an emissions production function to generate predicted emissions for each firm year. Industries vary significantly in their emissions intensities, so we run separate regressions at each industry-year level. The regressions estimate total emissions, y_{it} , as a function of assets, revenues and capital intensity. Formally, the production function is:

$$y_{it} = \beta_{1jt}x_{it} + \beta_{2jt}x_{it}^2 + \beta_{3jt}w_{it} + \beta_{4jt}w_{it}^2 + \beta_{5jt}z_{it} + \beta_{6jt}z_{it}^2 + \varepsilon_{it}$$

The variables x_{it} and x_{it}^2 are firm i 's assets and assets squared, respectively, in year t ; w_{it} and w_{it}^2 are firm i 's revenue and revenue squared, respectively; z_{it} and z_{it}^2 are firm i 's capital intensity and capital intensity squared, respectively. Capital intensity is defined as assets divided by employees. All coefficients are specific to their industry (j) and year (t). We choose these variables and specification to maximize the fit of the regression and predictive power to estimate expected emissions. While several of these regressors are clearly correlated with one another, our objective is to pull as much firm specific variance that can be explained by characteristics such as production levels, sales, and the capital to labor ratio (i.e., capital intensity) as possible out of the error term. As such, the relatedness of the regressors is less of an issue than would be typical.³²

Excess Emissions are defined as actual reported emissions (from Thomson Reuters or Bloomberg) minus predicted emissions ($E[y_{it}]$, estimated from the regression above). This excess

³² By estimating our emissions production function by industry sector by year, we generate over 200 sets of regression estimates. In the interests of brevity, we do not set these estimates out here, but note that the emissions production functions explain much of the emissions variance in our firms with an average r-squared of 0.86 and adjusted r-squared of 0.74. Full regression results are available upon request.

emissions variable is our main independent variable of interest. We are interested in how excess emissions predict labor productivity before and after the 2013 regulation for UK firms, and after regulation for UK firms relative to their European counterparts.

Descriptive statistics for emissions and financial variables are set out in Table 20. Firms in our control (European) sample are on average larger than firms in our treatment (UK) sample. This is because in order to be in our sample we require firms to have GHG emissions data in one of our data providers for all years between 2010 and 2015 inclusive. A higher proportion of UK listed firms reported GHG emissions over this time period than European firms (Jouvenot and Krueger, 2020). In order to adjust for these differences, we include various controls for firm size in our emissions production function and include firm fixed effects in some specifications. While absolute size and emissions vary between our treatment and control samples, our focal variables of interest, emissions intensity and labor productivity, are scaled by firm size and are much closer between the two samples.

[Insert Table 20 here.]

Study 2: Results

Our descriptive statistics show that while firms in our European sample are larger and have higher absolute emissions, our UK and European sample have similar emissions intensities. Figure 9 graphs emissions intensity over time, for UK and European firms split into above or below-median emissions intensity groups. UK high emitting firms decrease their emissions intensity post regulation, both in 2014 and then to a greater degree in 2017 and 2018. In contrast, high emissions intensity European firms see a slight increase post regulation in 2014. UK low emitters see a slight increase in emissions in 2015, while no change is observed in low emissions intensity European firms post regulation.

[Figure 9 here.]

Our main inquiry is how ESG performance affects labor productivity. Here we specifically examine the impact of environmental performance by regressing labor productivity on excess emissions in separate year-by-year regressions and examine the coefficient on excess emissions in the years before and after the UK regulation. We do not control for firm financial characteristics in these regressions because these financial characteristics are used in our earlier estimation used to create the excess emissions variable. Results for these regressions are set out in Table 21 and are plotted graphically in Figure 10.

The coefficient on excess emissions (or emissions residual) is statistically insignificant and close to zero for both the treatment (UK) and control (European) firms in 2010, 2011 and 2012. In 2013, the year the regulation went into effect, we observe a negative and statistically significant coefficient on excess emissions in the UK sample, and a positive and statistically significant coefficient in the European sample. Thus, excess emissions do not correlate with labor productivity for UK or European firms in the years leading up to the UK regulation. However, in 2013 and 2014, the year of and the year after the regulation was implemented, excess GHG emissions correlated with lower labor productivity for UK firms. The effect on European firms not subject to the regulation was the opposite – a positive correlation between excess emissions and labor productivity in 2013, and then a return to a non-significant relationship in 2014 and beyond.

[Insert Table 21 and Figure 10 here.]

To examine which firms are driving this relationship, we re-estimate these regressions for two relevant sector groups: industrial firms and consumer products firms. We choose these two sector groups in order to better distinguish between effects on our revenues-based labor productivity measure that may be driven more by consumer effects (i.e., changes in sales due to revealed GHG emissions being greater or lower than expected) versus employee effects (i.e., more revenues generated per employee because of increased labor productivity). The industrial firms group consists

of all firms in energy, chemicals, mining, resource, industrial, and transportation sectors. The consumer goods group consists of firms in consumer products and services, food, and household sectors. We expect that results for consumer facing firms are more likely to reflect changes in consumer behavior than results for industrial firms. Observed effects in industrial firms may be better attributed to changes in labor productivity. Regressions of labor productivity on excess emissions for these groups of firms are shown in Table 22.

[Insert Table 22 here.]

Results for industrial firms are qualitatively consistent with the results reported above. The coefficient on the emissions residual is statistically insignificant and close to zero until 2013, at which point the coefficient is significant and negative for UK firms and significant and positive for European firms. The 2013 coefficient on excess emissions for UK firms is around 50% larger in the industrials sub-sample than in the full sample estimate shown in Table 21. In our consumer goods sample, we see no statistically significant correlation between excess emissions and labor productivity. Thus, the negative correlation between excess emissions and labor productivity in Table 21 appears to be driven by industrial firms.

As a robustness test, we also perform a difference-in-differences regression on the pooled sample. We regress labor productivity on excess emissions over the 2010-2015 time period. We choose to end our sample in 2015 because of the Brexit vote in mid-2016 and to maintain a sample close to the treatment year. We include year dummies in all regressions to control for year-specific effects on labor productivity, and a region dummy in all regressions to control for consistent differences in labor productivity between firms in our sample in the UK and in Europe. We include a dummy called 'Regulation' which equals one for all firm-years subject to the regulation, that is, UK firms in 2013, 2014 and 2015. Any mean change in labor productivity for all UK firms in our sample over the 2013-2015 time period will be reflected in the coefficient on the regulation variable. While

we do not expect such a mean level change, we do expect that the regulation will differentially impact high and low emitting firms. To see this effect, we interact excess emissions with the regulation dummy. The coefficient on this interaction term captures the effect of the regulation across different levels of firm emissions.

We see a consistently negative coefficient on the interaction term in all three models (i.e., where included) and the coefficient is statistically significant at traditional levels in the baseline and sector fixed effects model. The coefficient is also negative with a p-value of 0.098 in the firm fixed effects model (column 6). This negative coefficient indicates that UK firms under the mandatory disclosure regime with high excess emissions had lower labor productivity than their European peers who were not subject to the regulation. Furthermore, this negative correlation between excess emissions and labor productivity was not present before the regulation was instituted.

[Insert Table 23 here.]

Our results indicate that in the year the UK regulation mandating standardized GHG emissions disclosure went into effect (and perhaps one to two years after) there was a significant negative correlation between excess GHG emissions and labor productivity for UK firms. Furthermore, this effect is driven completely by industrial firms, with no effect in consumer facing firms. Industrial firms have, on average, significantly higher absolute emissions and emissions intensity than the consumer facing firms in our sample. We also expect firms in our consumer category to have more interaction with final consumers. These two facts suggest to us that the regulation had a significant impact only on firms for which GHG emissions is a material ESG issue, and the effect of the regulation is driven at least in part by employees, or other stakeholders within the firm, rather than only consumers.

Our results appear to have an effect over a two-to-three-year period. This suggests that while the shock to the ESG information available to stakeholders has an immediate impact on the ESG-labor

productivity relationship, in our sample there appears to be a return to an equilibrium after a few years where the average relationship is close to zero.

Our results from this regulatory change complement the large-scale results correlating ESG and labor productivity in the first part of our paper. By analyzing an exogenous change in the regulatory regime, we are better able to assess the direct impact of ESG performance, and stakeholders' awareness of ESG performance, on labor productivity. Evidence from this regulation provides suggestive evidence that increasing information around ESG performance increases the correlation between ESG performance and labor productivity, at least in the short run.

Implications & Conclusion

In this paper, we examine the relationship between a firm's ESG performance and its labor productivity. Given the increasingly cited idea that people prefer to work for firms that share their values and have a positive impact on the world, it follows that ESG performance, as an indicator of both a firm's values and degree of positive impact, could drive how attractive the firm is as an employer. If a firm's environmental, social and governance profile is tied to its reputation and positive values lead to greater employee engagement and satisfaction, reducing turnover and potentially increasing effort, we would expect such a profile to affect labor productivity.

We focus on labor productivity as a measure of current firm performance, rather than a measure that also captures future expectations. Many past studies on the relationship between ESG and firm performance, while informative, have largely focused on stock returns, which captures not only past performance but market expectations around future returns. In this sense, we seek to link the impact of ESG metrics on a stakeholder of interest to the firm (i.e., employees) and on an outcome that ultimately bears a relationship with firm financial performance (i.e., labor productivity). Thus, we seek to better understand the feedback loops between ESG and elements of performance that may

not be immediately transparent to either investors or the firms themselves but are crucial to better understanding the (potentially conditional) importance of firm practices that underlie ESG.

We first provide correlational evidence on the distribution of the ESG-labor productivity relationship across a large set of US firms. Using a recent source of ESG data that relies on external observation of events (from news sources, government regulators, NGOs and other outside parties) rather than firm self-reported information that bears an unclear relationship with actual ESG performance, we find significant heterogeneity between and within industries in the relationship between ESG and labor productivity. Further, the variance of this relationship is in most cases more significant within industries than between them; however, we note that some industries, such as utilities, exhibit less variance on this relationship than others such as manufacturing and mining.

We also note that the estimates of the relationship between ESG and labor productivity vary according to which elements of ESG are included. For example, environmental elements of ESG show greater variance than social elements in the relationship with labor productivity in mining, whereas the opposite is true in the services sector. In some cases, these firm level relationships are positive, suggesting a benefit to firm labor productivity from ESG performance, while some relationships in the same industry are negative, suggesting that in those firms there may be a resource tradeoff from investing in ESG that diminishes labor productivity, along the lines suggested by Becker (2011) with respect to environmental management. We briefly explore the types of firms which have positive relationships between ESG performance and which firms have negative relationships. We find firm size correlates positively with the labor productivity-ESG relationship, suggesting large firms may see a benefit from ESG while small firms see a cost.

Via a quasi-experimental setting, we also show causal evidence of how a change in the visibility of ESG performance affects the ESG-labor productivity relationship. We find that the introduction of mandatory and standardized greenhouse gas emissions reporting in the UK in 2013 correlated

with a significant drop in the labor productivity of high emitting UK firms and an increase in the labor productivity of low emitting UK firms. Furthermore, this relationship held when comparing to a control group of European firms unaffected by the regulation.

We find the financial benefits of stronger ESG orientations vary according to firms, at least when considered through the lens of labor productivity. Thus, a firm's strategy around ESG should not be one size fits all, even within the same industry. While this seems somewhat obvious, from news reports and past research that generates mean effects across a sample, one may get the impression that investing in practices that drive ESG performance has no downside. From a financial performance perspective, our results imply that investors and pension funds (as well as the firms themselves) need to be sophisticated in their evaluation of what elements of ESG are most important to focus on, given a firm's competitive strategy. This is essential knowledge to most efficiently allocate resources towards solving our most pressing challenges, such as climate change.

Furthering the type of analysis contained in this paper can also help firms prioritize their ESG strategy. With the proliferation of reporting standards on different elements of ESG (e.g., TCFD, GRI, UNGP Reporting Framework, CDP, etc.), firm prioritization of ESG issues can be challenging. With better understanding of how elements of ESG link to firm performance that are specific to firm characteristics, firms can better devise an effective strategy for implementation of practices driving ESG metrics that both improves financial performance and has a meaningful impact on stakeholders most affected by that firm.

If firms are able to better understand what elements of ESG are important for different aspects of performance, they can better make the case for investments in such areas with boards and investors. This will also allow a more nuanced creation of ESG strategy that moves from a shotgun approach to one that is more tailored to the overall strategy and vision for the firm. Such information may also reveal the time horizons for payoffs from investments in practices that drive

ESG metrics and performance. Further exploration on how long it takes for positive ESG performance to translate into benefits to human resources – that is, better selectivity, retention, engagement and, potentially, effort – that drive labor productivity is also warranted. This type of analysis may also help identify how firms develop authenticity around ESG practices, which is thought to be an important element in translating these practices into greater customer and employee loyalty and lowering risk more generally with stakeholders.

Our study is limited by our sample, which is constrained to large, public firms. These are the firms where external reporting is available to provide a more informative picture of a firm's ESG performance. We find that the relationship between ESG performance and labor productivity is less positive for smaller firms in our sample. These results suggest that smaller firms may experience more tradeoffs when investing in practices that generate ESG performance. To the extent that ESG type investments detract from the overall performance and stability of the firm, such investments and subsequent ESG performance may make that firm a less attractive employer. A broader examination of these potential tradeoffs in smaller firms is a fruitful avenue for future research once more robust ESG data is available for these firms.

Alternatively, our causal study suggests that changing the visibility of ESG performance causes greater alignment between ESG performance and labor productivity, consistent with the firm size results from our correlational results. As ESG reporting frameworks continue to standardize and reporting becomes increasingly mandatory, these results suggest that the ESG-labor productivity link will become more positive over time, making ESG an increasingly financial material issue for firms.

Tables and Figures

Figure 5: US Firm Estimates, by sector




Figure 6: US Firm Estimates, early sample period

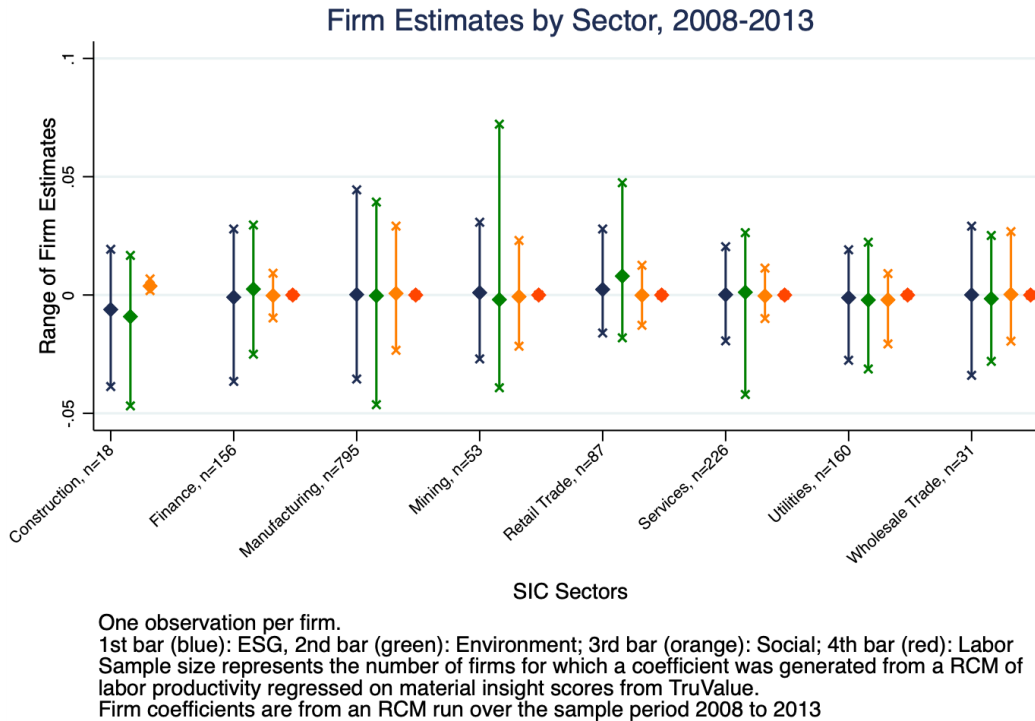
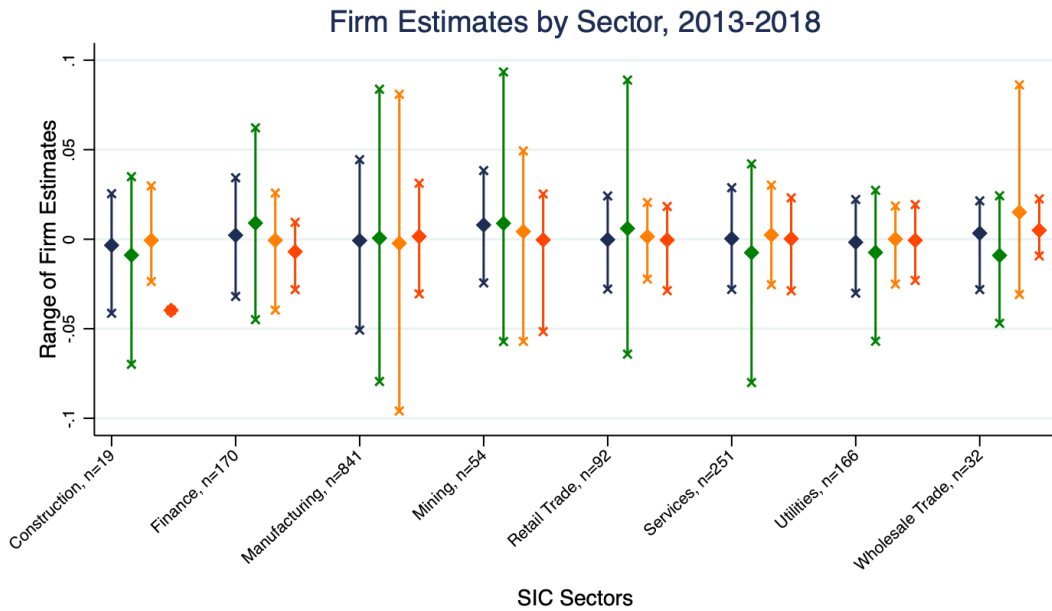


Figure 7: US Firm Estimates, late sample period



One observation per firm.
 1st bar (blue): ESG, 2nd bar (green): Environment; 3rd bar (orange): Social; 4th bar (red): Labor
 Sample size represents the number of firms for which a coefficient was generated from a RCM of labor productivity regressed on material insight scores from TruValue.
 Firm coefficients are from an RCM run over the sample period 2013 to 2018

Figure 8: Firm effects by firm size (assets and revenues)



Figure 9: Emissions intensity, high and low emitting firms, by region-year




Figure 10: Coefficient estimates of relationship between labor productivity and excess emissions, by region-year

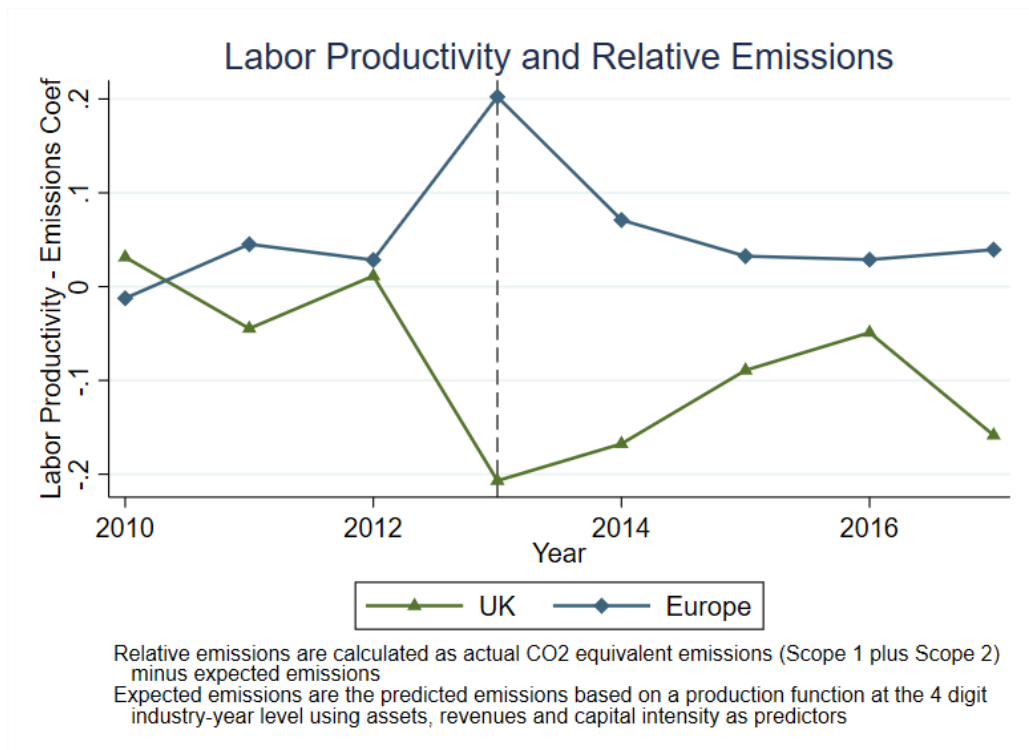


Figure 11: Coefficient estimates of relationship between revenue and excess emissions, by region-year




Figure 12: Coefficient estimates of relationship between employees and excess emissions, by region-year

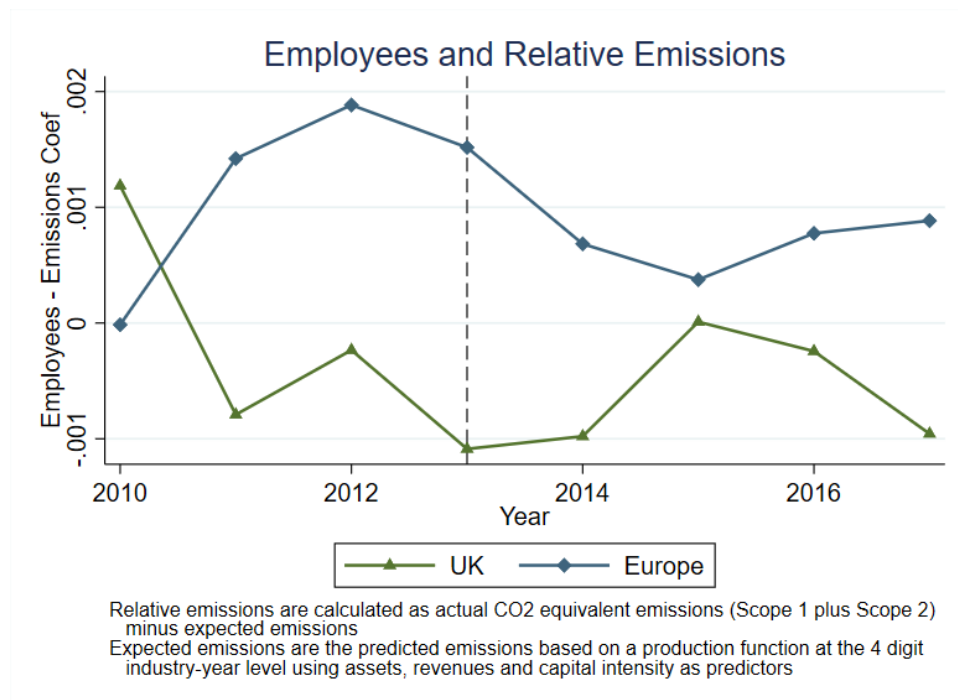


Table 14: Illustrative news articles from TruValue Labs dataset

MT / ArcelorMittal		
ESG Effect	Negative	Category Waste & Hazardous Materials Management
Article	ArcelorMittal has been accused of manipulating test results to conceal the amount of ammonia and cyanide that is being discharged from its plant in Indiana.	
Source	Steel Guru: steelguru.com/steel/arcelormittal-burns-harbor-accused-of-manipulating-test-results-by-idem	
AMZN / Amazon.com		
ESG Effect	Negative	Category Labor Practices
Article	Amazon has allegedly threatened to fire employees for criticizing the company's environmental policies.	
Source	Vanity Fair: vanityfair.com/news/2020/01/amazon-threatens-to-fire-employees-climate-change-environment	
JNJ / Johnson & Johnson		
ESG Effect	Negative	Category Product Quality & Safety
Article	Washington has sued Johnson & Johnson for allegedly deceiving patients and doctors about the addictive nature of the opioid pain killers it developed and marketed.	
Source	My Northwest: mynorthwest.com/1658163/ag-ferguson-johnson-and-johnson-opioid-lawsuit	
ORSTED / Orsted A/S		
ESG Effect	Negative	Category Systemic Risk Management
Article	Orsted will pay £4.5 million because its power station failed to remain connected after a lightning strike, causing a power outage in the UK that affected more than one million customers.	
Source	New York Times: nytimes.com/reuters/2020/01/03/business/03reuters-britain-powercut.html	
GE / General Electric Co		
ESG Effect	Positive	Category Product Design & Lifecycle Management
Article	GE said it won 1,215 MW of wind turbine orders or commitments from customers in China in 2019.	
Source	Renewables Now: renewablesnow.com/news/ge-bags-12-gw-of-wind-orders-in-china-in-2019-681584	
TXT / Textron Inc		
ESG Effect	Negative	Category Critical Incident Risk Management
Article	Several employees of Textron Inc were reported injured following an explosion at a plant in Wichita, Kansas.	
Source	News 8000: news8000.com/news/national-news/explosion-reported-at-kansas-aviation-manufacturing-plant	
ARCO / Arcos Dorados Holdings		
ESG Effect	Negative	Category Employee Health & Safety
Article	Peruvian authorities say that McDonald's Latin America franchisee Arcos Dorados violated the law with its working conditions after two employees were electrocuted and killed.	
Source	Haaretz: haaretz.com/world-news/wires/1.8321179	
DRE / Duke Realty		
ESG Effect	Positive	Category Energy Management
Article	Duke Realty has announced that, starting this year, it will develop only LEED-certified buildings.	
Source	MarketWatch: marketwatch.com/press-release/all-duke-realty-new-development-projects-will-be-leedr-certified-2020-01-07	

Table 15: Sample Construction

	Year											Total
	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	
TruValue Data	1,693	2,111	3,208	3,711	4,113	4,483	4,885	5,199	5,285	5,258	5,046	44,992
Compustat coverage	1,533	1,902	2,622	2,972	3,251	3,532	3,805	4,005	3,987	3,910	3,712	35,231
5+ Years Data	1,100	1,387	1,937	2,231	2,468	2,701	2,902	2,821	2,716	2,599	2,407	25,269
Min ESG Article Volume	873	1,052	1,204	1,338	1,444	1,563	1,667	1,631	1,570	1,494	1,375	15,211
Sample firms	850	1,023	1,175	1,306	1,408	1,524	1,620	1,585	1,527	1,451	1,335	14,804

Table 16: Industry Coverage

	SIC Sector								Total
	Construction	Finance	Manufacturing	Mining	Retail Trade	Services	Utilities	Wholesale Trade	
US Firms	19	173	855	54	93	262	173	32	1,661

Table 17: Descriptive Statistics

	Mean	Median	SD	Min	Max	N
Construction						
Labor Productivity	13.3	13.4	0.9	12.0	16.6	164
ESG Performance	67.8	70.1	15.5	8.4	99.1	164
Capital Intensity	13.2	13.2	1.2	11.2	16.4	164
Finance						
Labor Productivity	13.2	13.2	0.9	10.0	17.1	1,407
ESG Performance	58.3	58.3	14.5	3.0	94.3	1,407
Capital Intensity	15.3	15.5	1.4	11.3	20.0	1,407
Manufacturing						
Labor Productivity	12.6	12.7	1.0	4.7	17.7	6,975
ESG Performance	65.6	67.9	14.7	3.4	99.4	6,975
Capital Intensity	13.1	13.0	0.9	9.5	18.4	6,975
Mining						
Labor Productivity	13.5	13.4	1.2	6.7	15.9	420
ESG Performance	52.5	52.3	16.3	2.6	97.9	420
Capital Intensity	14.3	14.3	1.4	10.7	17.2	420
Retail Trade						
Labor Productivity	11.9	11.9	0.9	10.1	15.2	754
ESG Performance	59.3	59.6	14.3	5.9	99.3	754
Capital Intensity	11.4	11.3	1.0	9.6	15.0	754
Services						
Labor Productivity	12.2	12.3	0.9	7.3	16.7	1,893
ESG Performance	59.9	61.4	15.3	1.0	98.9	1,893
Capital Intensity	12.6	12.7	1.2	7.7	19.0	1,893
Utilities						
Labor Productivity	13.1	13.1	0.8	10.3	15.5	1,441
ESG Performance	60.6	61.9	14.8	6.5	96.9	1,441
Capital Intensity	13.9	13.9	1.1	10.8	17.6	1,441
Wholesale Trade						
Labor Productivity	13.9	13.8	0.9	11.8	17.1	259
ESG Performance	66.3	70.3	16.4	10.7	92.4	259
Capital Intensity	13.1	13.2	0.9	10.8	15.5	259
Total						
Labor Productivity	12.7	12.7	1.0	4.7	17.7	13,313
ESG Performance	62.8	64.7	15.3	1.0	99.4	13,313
Capital Intensity	13.3	13.1	1.4	7.7	20.0	13,313

Labor productivity and capital intensity are log transformed.

Table 18: ESG and Labor Productivity

DV = Labor Productivity												
	RCM			FE			RCM			FE		
	FE	beta	sd	FE	beta	sd	FE	beta	sd	FE	beta	sd
	(1)	(2a)	(2b)	(3)	(4a)	(4b)	(3)	(4a)	(4b)	(3)	(4a)	(4b)
ESG												
All Categories	-0.002	-0.002	0.078									
ESG	(0.377)	(0.601)	(0.000)									
Environmental Performance				-0.0002	-0.001	0.090						
				(0.951)	(0.808)	(0.000)						
Social Performance							0.003	0.001	0.078			
							(0.574)	(0.821)	(0.000)			
Labor Performance										-0.003	-0.002	0.000
										(0.328)	(0.363)	(0.000)
Financials												
Assets (logged)	0.056	0.062	0.190	0.068	0.062	0.196	0.051	0.079	0.191	-0.034	-0.029	0.246
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.026)	(0.303)	(0.000)
Industry Concentration	-0.010	-0.010	0.000	-0.014	-0.013	0.000	-0.011	-0.005	0.000	0.036	0.007	0.028
	(0.191)	(0.127)	(0.000)	(0.101)	(0.073)	(0.000)	(0.529)	(0.742)	(0.000)	(0.054)	(0.706)	(0.000)
Observations	8,149	8,149		6,729	6,729		4,091	4,091		1,202	1,202	
Groups	983	983		803	803		502	502		141	141	
R-squared	0.022			0.026			0.018			0.019		
Log likelihood		14174			11862			5996			2548	
LR p-value		0.000			0.000			0.000			0.000	
Chi-Square		215.139			179.165			142.305			35.848	

p-values in parenthesis; FE = Fixed effects models; RCM = Random coefficient models. LR p-value refers to the p-value of the LR test vs. a linear regression. All models include time period fixed effects; fixed effects model includes firm fixed effects; random coefficient models include industry sector fixed effects

Table 19: UK Stufdy, Sample Stock Exchanges

Stock Exchange	Number of Firms
London Stock Exchange	162
Euronext Paris	63
Deutsche Boerse AG (XETRA)	52
Borsa Italiana S.P.A.	28
Bolsas de Madrid	27
Six Swiss Exchange	27
Nasdaq Stockholm	26
Nasdaq Helsinki	22
London, Foreign Listing	21
Euronext Amsterdam	18
Oslo Bors ASA	16
Nasdaq Copenhagen	15
Euronext Brussels	11
Wiener Boerse AG	10
Euronext Lisbon	6
Irish Stock Exchange	1

In order to be included in the counts above, a firm must have emissions data present for 2010 through 2015, inclusive. Foreign listed firms on the London Stock Exchange are split into their own category and excluded from empirical analysis because the greenhouse gas emissions reporting regulation did not apply to these firms.

Table 20: UK Study, Descriptive Statistics

	UK	Europe		UK	Europe
Total Emissions			Revenue (\$M)		
<i>N</i>	1,863	3,765	<i>N</i>	1,863	3,765
<i>mean</i>	1,737,648	5,652,852	<i>mean</i>	\$9,980	\$19,300
<i>median</i>	76,968	235,345	<i>median</i>	\$2,180	\$6,610
<i>std dev</i>	7,113,538	20,800,000	<i>std dev</i>	\$28,600	\$36,600
<i>p5</i>	1,750	4,957	<i>p5</i>	\$288	\$844
<i>p95</i>	9,200,000	25,300,000	<i>p95</i>	\$45,500	\$84,500
Scope 1 Emissions			Employees		
<i>N</i>	1,468	2,799	<i>N</i>	1,863	3,765
<i>mean</i>	1,487,807	822,816	<i>mean</i>	24,604	46,293
<i>median</i>	35,423	136,000	<i>median</i>	7,908	16,702
<i>std dev</i>	6,204,848	2,812,882	<i>std dev</i>	58,745	73,902
<i>p5</i>	310	2,682	<i>p5</i>	630	1,581
<i>p95</i>	10,200,000	3,796,000	<i>p95</i>	92,000	182,865
Scope 2 Emissions			Assets (\$M)		
<i>N</i>	1,459	2,799	<i>N</i>	1,744	3,483
<i>mean</i>	562,965	822,816	<i>mean</i>	\$3,590	\$10,000
<i>median</i>	37,425	136,000	<i>median</i>	\$986	\$3,300
<i>std dev</i>	2,334,409	2,812,882	<i>std dev</i>	\$8,760	\$21,300
<i>p5</i>	827	2,682	<i>p5</i>	\$114	\$445
<i>p95</i>	1,974,997	3,796,000	<i>p95</i>	\$17,400	\$41,200
Emissions Residual			R&D (\$M)		
<i>N</i>	1,350	2,657	<i>N</i>	332	1,305
<i>mean</i>	-13,520	6,869	<i>mean</i>	\$644	\$940
<i>median</i>	-2,393	-85	<i>median</i>	\$118	\$190
<i>std dev</i>	2,242,962	2,195,541	<i>std dev</i>	\$1,380	\$1,780
<i>p5</i>	-2,166,473	-2,523,072	<i>p5</i>	\$0	\$9
<i>p95</i>	1,351,210	2,285,197	<i>p95</i>	\$4,940	\$5,130
Labor Productivity			Advertising (\$M)		
<i>N</i>	1,848	3,763	<i>N</i>	38	479
<i>mean</i>	\$958,271	\$1,746,533	<i>mean</i>	\$1,970	\$641
<i>median</i>	\$245,022	\$678,371	<i>median</i>	\$81	\$95
<i>std dev</i>	\$2,559,226	\$3,100,140	<i>std dev</i>	\$3,610	\$1,730
<i>p5</i>	\$47,229	\$111,282	<i>p5</i>	\$3	\$5
<i>p95</i>	\$4,334,466	\$7,172,797	<i>p95</i>	\$9,650	\$3,130

Table 21: Labor productivity and excess emissions, by region-year

UK	2011	2012	2013	2014	2015	2016
Excess Emissions	-0.0447	0.0114	-0.207**	-0.168*	-0.0892	-0.049
	[0.553]	[0.862]	[0.002]	[0.048]	[0.208]	[0.255]
N	135	135	135	135	135	135
r-squared	0.003	0.000	0.069	0.029	0.012	0.01
Europe	2011	2012	2013	2014	2015	2016
Excess Emissions	0.0452	0.0284	0.202*	0.071	0.0324	0.0287
	[0.545]	[0.763]	[0.035]	[0.377]	[0.558]	[0.602]
N	269	269	269	269	268	268
r-squared	0.001	0.000	0.017	0.003	0.001	0.001

Results from yearly regressions of labor productivity on an emissions residual, defined as the residual of a regression of emissions on assets, revenue, capital intensity, and all three squared.

p-values in brackets, * p<0.05, ** p<0.01, *** p<0.001

Table 22: Sector level regressions of labor productivity on excess emissions

Industrial Firms

UK	2011	2012	2013	2014	2015	2016
Excess Emissions	0.0042	0.045	-0.307**	-0.197	-0.0489	-0.0626
	[0.97]	[0.618]	[0.002]	[0.069]	[0.602]	[0.211]
N	60	60	60	60	60	60
r-squared	0	0.004	0.159	0.056	0.005	0.027
Europe	2011	2012	2013	2014	2015	2016
Excess Emissions	0.0214	-0.0178	0.254*	0.0992	0.0178	0.0417
	[0.825]	[0.88]	[0.019]	[0.328]	[0.773]	[0.414]
N	130	130	130	130	130	130
r-squared	0.000	0.000	0.042	0.007	0.001	0.005

Consumer Goods Firms

UK	2011	2012	2013	2014	2015	2016
Excess Emissions	-0.072	-0.13	-0.12	-0.0315	-0.0291	-0.0179
	[0.83]	[0.744]	[0.751]	[0.932]	[0.92]	[0.936]
N	54	54	54	54	54	54
r-squared	0.001	0.002	0.002	0.000	0.000	0.000
Europe	2011	2012	2013	2014	2015	2016
Excess Emissions	0.218	0.345	0.346	0.271	0.36	0.33
	[0.646]	[0.577]	[0.603]	[0.694]	[0.62]	[0.629]
N	73	73	73	73	72	72
r-squared	0.003	0.004	0.004	0.002	0.004	0.003

Table 23: Difference in difference estimates of effect of UK regulation

Labor Productivity and Excess Emissions						
	(1)	(2)	(3)	(4)	(5)	(6)
Excess Emissions	0.00616 [0.835]	0.0597 [0.06]	0.00124 [0.962]	0.0406 [0.141]	-0.00584 [0.505]	-0.00202 [0.823]
Regulation	66945.4 [0.806]	89229.2 [0.743]	61569.3 [0.794]	77928.7 [0.741]	65060.6 [0.213]	66769.3 [0.201]
Regulation * Excess Emissions		-0.392*** [<0.001]		-0.287*** [<0.001]		-0.0339 [0.099]
Year, Region F.E.	x	x	x	x	x	x
Sector Fixed Effects			x	x		
Firm Fixed Effects					x	x
N	2,423	2,423	2,423	2,423	2,423	2,423
Adjusted r-squared	0.01	0.02	0.27	0.28	0.05	0.05

Dependent variable: Labor productivity. Excess emissions are calculated as reported emissions minus predicted emissions, generated from an emissions production function with revenues, assets, and capital intensity as predictors. Regulation is a dummy that takes the value of 1 beginning in 2013 for all firms subject to the UK GHG reporting requirement.

p-values in brackets, * p<0.05, ** p<0.01, *** p<0.001

Chapter 3: Wages and Employee Attrition³³

Abstract

This paper provides large scale evidence on the relationship between wages, relative to comparable jobs in the same geographic area, and employee attrition in the US between January 2018 and June 2020. Confirming past studies, we find that among workers earning below \$62,000 per year, but not among those earning more, the correlation between relative pay and attrition is negative. We also find that the negative relationship between relative wages and attrition is strongest at moderate levels of local job market size, and near zero when the local job market is large. This may suggest that in highly competitive markets, firms may have to turn to incentives other than wages to reduce attrition. However, when we instrument for wages using competitor firm wages in other cities, we are unable to identify these negative relationships. Our results support the view that paying employees higher wages is a strategic choice firms can use to reduce attrition, while revealing various contexts in which this strategy appears to fail.

³³ This chapter is coauthored with Rachelle C. Sampson

Introduction

Employees in the United States move between jobs increasingly often, and as a result firms are under pressure to retain their workforce by reducing employee attrition. In 2021, more than 40% of employees were thinking of leaving their jobs, and from April through November a record number left their companies (Sull, Sull, and Zweig, 2022). Changing demographics increasingly contribute to these trends, as younger workers in the US display less attachment to their employers than past generations and move between employers more often (hiQ Labs, 2021). This employee attrition is a financial threat to firms, as attrition results in a loss of firm-specific human capital and increased training and recruitment costs (Dess and Shaw, 2001). Furthermore, this loss of employees can reduce unit-level efficiency and service quality, both resulting in worse firm financial performance (Glebbeek and Bax, 2004; Hausknecht, Trevor, and Howard, 2009; Kacmar *et al.*, 2006).

While firms can reduce employee attrition through company culture and by offering greater benefits (Dale-Olsen, 2006; Sull *et al.*, 2022), a key lever to reduce attrition is through pay policy. To reduce attrition, firms can pay above market wages with the goal of inducing employee loyalty and attachment to the firm (Shapiro and Stiglitz, 1984). Large scale correlations suggest this strategy is effective. On average, higher pay is associated with lower levels of attrition (Bassier, Dube, and Naidu, 2021; Duhautois, Gilles, and Petit, 2016; Heavey, Holwerda, and Hausknecht, 2013). However, employees and firms match endogenously, suggesting the pay-attrition relationship may be the result of unobserved but correlated worker characteristics or firm characteristics. For example, higher paying jobs provide, on average, greater job benefits (Sockin, 2021). Thus, a key question in the literature is: when does increasing pay cause lower turnover? Furthermore, increased pay may result in lower attrition only under certain conditions. For example, compensation may be a more effective lever to retain employees when they make below a certain level of earnings (Duhautois *et al.*, 2016). High levels of compensation can also be associated with increased job anxiety, which

could lead to greater attrition rates (Bryson, Barth, and Dale-Olsen, 2012). Thus, managers may reduce attrition most effectively by focusing on the wages of certain groups of employees.

This study addresses these questions of causality and conditionality in the wage-turnover relationship. First, we ask: what is the causal effect of wages on employee attrition? Second, if there is a relationship, for which firms and under what condition are wages an effective tool to reduce attrition?

We address these questions using large-scale data generated by aggregating all job postings in the United States alongside online job profiles and government labor data. As a result, our data provides a view of job flow dynamics across the United States at a highly localized level. We are able to identify the number of employees and employee flows in and out of jobs every month at the company-MSA-job category-seniority level. This data granularity provides several advantages. First, we are able to generate highly specific job comparison groups to establish workers' outside wage options. This allows us to assess whether firms offer wage premiums relative to the jobs their workers might otherwise be qualified for and considering for employment. This granularity also allows us to separate attrition due to firm characteristics from attrition due to local labor market conditions or attrition trends within specific job categories.

Identifying whether and when higher wages decrease employee attrition in the US provides several contributions to the literature and practice. First, we provide a novel theory to explain the U-shaped relationship we find between worker outside options and the wage-attrition relationship. This theory provides evidence as to why researchers may only identify this negative relationship among certain populations of workers. Specifically, it explains the lack of a wage-attrition relationship when markets are very thick, which the monopsony literature does not address. Our findings also lend support to the idea that changes in wages matter more for lower wage workers,

consistent with the literature on the relationship between well-being and wages (Kahneman and Deaton, 2010).

Our results also provide material guidance for managers attempting to respond to the great resignation of recent years. Companies and managers may find it valuable to focus on remuneration and retention strategies more focused on non-pecuniary benefits, such as better company culture. Even if firms suspect that higher wages decrease turnover, the obvious cost to this policy is increased salary bill. Our findings provide evidence that managers weighing the costs and benefits of increasing wages may find the greatest benefits when employees have moderate levels of outside options and low pay. Pay is only one lever to reduce attrition, and in situations where we find pay to be less effective, managers can turn to other measures such as benefits or culture.

One contribution we intend, but do not yet make, to the literature is identifying the causal relationship between relative wages and attrition rates. Recent literature has begun to move beyond correlational results, and we push this literature further using identification strategies from the economics literature applied to our questions of attrition (Azar, Berry, and Marinescu, 2019). However, the strategy we take in our paper does not currently provide us with a strong instrument to establish causality. We expect the strategy to be effective in identifying any causal effects when we receive updated data from our data provider, as a similar strategy has produced strong first stage results (Azar et al 2019).

To situate our study, we next outline related research and describe how this project extends past work. We then introduce the data and our measures. Results follow, and the paper ends with a discussion.

Background

According to the efficiency wage model, an equilibrium can exist where some firms pay above market wages, while others pay below market wages. Low wage firms can be profitable due to lower

wage bills, while high wage firms are compensated for their increase wage bill by through more productive workers and decreased employee attrition, resulting in lower recruitment and training costs. Consistent with this model, significant empirical evidence has found that high wage employees have lower turnover rates than low wage employees (Heavey *et al.*, 2013).

Empirical research estimating the wage-attrition relationship often faces the tradeoff between examining large swatches of the economy, and precisely identifying comparison groups to establish whether a firm pays their workers above market wages. On the specific context side, Siebert and Zubanov (2009) study a single set of workers in a single retail firm, thus allowing precise comparisons for wages relative to the local markets at different store locations. They find that sales associates have higher turnover when their wages are below the local county average wage. Similarly, Batt, Colvin, and Keefe, (2002) find that the relative pay of telecommunications workers is negatively associated with quit rates.

Large, matched employer-employee datasets have allowed researchers to study trends across whole economies. For example, Duhautois *et al.*, (2016) find that in France in the early 2000s, establishment level wage premiums are associated with lower churn rates within the establishments. Ilmakunnas and Maliranta (2005) study plants in Finland in the 1990s and find that high wage plants have lower turnover than low wage plants. Dale-Olsen (2006) study Norwegian workers in 1996 and 1997 and find higher wage firms have lower attrition. These studies establish large scale trends but come with the difficulty of accounting for differences in job types. If attrition rates vary systematically by job type, then these job type differences may influence the relationship between pay and attrition.

The first aim of this paper is to capture the causal effect of wages on firms' attrition rate. Firms which pay higher wages may be, on average, better places to work. If so, higher wages may be a signal and not a cause of lower employment attrition. Recent work in economics has begun to apply

causal inference tools to similar questions. For example, Bassier *et al.*, (2021) use a matched employer-employee dataset from Oregon between 2000 and 2017 and rely on employees who change firms to estimate the causal effect of wages on the elasticity of labor at the firm level. This paper approaches a similar question but focuses on firm-location-job category attrition rates rather than labor supply elasticities.³⁴ Our highly specific comparison groups also allow us to better capture the effect of wages relative to employees' outside options. We also take a different causal inference approach. We instrument for wages directly using two stylized facts. First, employers tend to prefer to set flat wage scales across different geographies, often out of concerns over equity within the firm (Azar *et al.*, 2019). For example, Costco has set national minimum wages for its own workers over the past decade (Tyko, 2021). Similarly, remote workers for large tech firms were often paid the same amount whether they worked in the Bay Area or rural Colorado. As a result, workers for these national firms in rural areas may often have higher salaries, relative to local competitors, due to their firm's footprint in high-cost cities. Second, firms respond to local economic conditions when setting their local wages. Our instrumental variables approach relies on the idea that if a firm has establishments in two locations and competitor firms in the first location increase their average wages, the firm is likely to increase their own employees' wages in both the first and second locations. This increase in the competitor firm's wages in the first location is unlikely to be correlated with the focal firm's attrition rates in the second location, except through this increase in wages. For example, consider a worker in Dayton Ohio in 2012 who was choosing between two similarly paying jobs, one with a firm headquartered in Detroit, and one with a firm headquartered in Austin. Over the past decade Austin has experienced much greater wage growth than Detroit. Thus,

³⁴ While related, we take the perspective that the labor supply elasticity is a separate concept than the relationship between wages and attrition rates within firms. Some work in economics measures labor supply elasticities by either measuring the total number of workers a firm employs, or assuming that the elasticity is double the attrition rate. We aim to explicitly study the differences in attrition, or churn, among their employee base, without making any assumptions about how attrition relates to labor supply elasticities.

our approach would expect the worker to have received higher salary growth in their Dayton, Ohio job if they chose to work for the firm headquartered in Austin.

Conditionality

In addition to identifying a causal relationship between relative wages and employee attrition rates, we investigate when and under what conditions paying higher wages is an effective strategy to reduce employee separations.

We predict higher wages will be more salient for low wage workers due to the decreasing marginal utility of money. Emotional well-being correlates strongly with income, but the relationship becomes weaker once workers earn around \$75,000 (Jebb *et al.*, 2018; Kahneman and Deaton, 2010), though there is debate as to this weakening relationship (Killingsworth, 2021). Thus, we expect higher wages to be more negatively correlated with attrition for lower wage workers. This finding would be in line with Duhautois, Gilles, and Petit (2016), who, though not their primary investigation, find that the negative relationship between firm wage premia and employee attrition reverses among high wage and high skill workers in their sample. We aim to extend their findings by incorporating our causal identification strategy to address this question.

Hypothesis 1: Higher wages will be most effective at decreasing employee attrition among lower wage workers.

Next, we turn to the outside options employees have when looking for work. A significant literature in economics investigates how worker wages respond to monopsony: the condition when there are very few firms in a labor market. Under monopsony, firms can be price setters for wages, and exercise great control over all aspects of employment due to workers' lack of outside options. As a result, firms pay workers well below their marginal productivity and not face the consequences of increased employee attrition (Azar *et al.*, 2019). Thus, the elasticity of labor supply is lowest when there is significant monopsony power in the market. This work in economics examines firms' labor

supply and find that it does not change significantly when their wages change. Unlike this literature, however, we are empirically interested in what reduce attrition rates for firms within specific geographic and job-category groups, even when firms rehire workers for the same role.

We expect that wages will be most effective in reducing attrition in markets where workers have an intermediate level of outside options. If an employee has few options to find work outside of their employer, when they search for a new job, they will be unlikely to secure one. In these cases, their current salary will have little effect on their decision to leave their employer because they cannot reach the point where they can compare the characteristics of their current job to those of a potential new job. As a worker's outside options increase, employees have the ability to apply for and gain offers for alternative employment. When deciding whether to accept the new job (and, thus, generate attrition for their current employer), they compare the characteristics of the new job, including wage, against those of their current job. In these cases, salary will be highly relevant to the worker, and changes in their current salary will have a large effect on their decision to leave the firm, and, thus, the firm's attrition rate.

However, as employees' outside options become very plentiful, it is increasingly difficult for a firm to offer better pay than all of their employees' outside options. At the limit, an employee will continuously be presented with the opportunity for better paying jobs. At this point, a firm will have to be more reliant on credence goods which are more difficult for employees to judge the quality of for firms at which they are not employed. Under these very high levels of employee outside options, we thus expect the relationship between wages and attrition rates to be lower than when employees have moderate levels of outside options.

These ideas, which are formalized in a simulation model included in the Appendix, lead to our next hypothesis:

Hypothesis 2: Higher wages will be most effective at decreasing employee attrition when workers have moderate levels of outside options.

Empirical Application

Data and Sample

We use a dataset of 15,000 public and private companies' workforce dynamics in the United States for each month between January 2018 and June 2020. The dataset contains data for all public firms which are among the 3,000 largest public firm by employment at any point between 2008 and 2020, and for all private firms which appear among the 3,000 largest during that time frame.

We source our data on granular employee wages and turnover from Revelio labs. Each observation in our data provides the number of employees, employee wages, and employee inflow and outflow information at the firm-month-MSA-job category-seniority level. Inflow and outflow information indicate when employees enter or leave a focal role in a company. Thus, our current data do not differentiate between employees which leave the company, and those who are promoted within the firm, or whose job is reclassified within the firm.³⁵ For example, one observation in the data can indicate that firm X employs 12 junior sales managers in the Charlotte, North Carolina MSA in June 2019, these sales managers are paid an average salary of \$41,512 per year, the firm gained two additional junior sales managers in Charlotte that month, and one left the role that month.

Revelio uses two main types of data sources to generate their data and supplements these data with government sources. To estimate the number of employees and the inflow and outflow of employees from jobs, Revelio uses information from online job profiles (e.g., from LinkedIn or Jobcase). To estimate job salaries, Revelio gathers data from all online job postings made, both on

³⁵ We expect to receive data which allow us to observe whether this inflow our outflow stems from within-company movement, or between company movement in late April.

job posting websites (e.g., Monster) as well as company websites. Of course, not all workers create online job profiles, and to adjust for the probability of a worker creating a profile, Revelio takes several steps to correct for unobserved workers using governmental data, outlined in more detail below.

The significant challenge in combining job postings and job profile data is that job titles and descriptions in each are not perfectly matched across a standard classification system. To address this issue, Revelio uses a large-scale natural language processing algorithm to classify all jobs into a single geometric space. This process can be thought of as dimension reduction on the domain of natural language in job postings and profiles. This space is then split into 150 job categories, with each job posting and job profile falling into one or more of these categories. Revelio also classifies all jobs and employees across four levels of seniority, from 1 (most junior) to 4 (most senior).

For a sample of our data format, see Table 24.

[Insert Table 24 Here]

Combining large scale data from online job profiles and all online job postings has significant advantages over past non-governmental data sources on company wage distributions. This data allows us to perform large analysis both across firms and within firms over time. Significantly, calculating within firm wage distributions has often had to rely on broad, firm-reported wage data (e.g., firm-reported median wages, which are subject to significant biases and limited generalizability), or single-firm case studies. For example, studies on within firm wage disparities have often relied on either single company studies, or professional sports teams (Shaw, 2014). While this data is advantageous due to data availability, a larger scale study such as ours can address questions of generalizability. Our data is much closer to census or other government administrative data and comes with corresponding advantages and disadvantages. See Duhautois *et al.* (2016) for a summary of recent work in the labor elasticity literature using detailed administrative data. As one

example, Mueller, Ouimet, and Simintzi (2017) use granular employee pay data from the UK to study the effect of within firm overall pay inequality on firm value and operating performance.

Our data is particularly advantageous because it is generated using publicly available data sources, and it is thus reproducible. The data is also generated through employee, rather than firm, reporting, it is highly geographically precise, and it uses a standard classification scheme to match different types of workers across firms and industries. Both the standard job classification system and the geographic granularity in our data allows us to address new questions in this research stream. For example, using the job category system with 150 different types of jobs and 4 levels of seniority, comparable across firms, we can ask questions that involve comparing worker pay in one firm, relative to similar workers in a local competitor firm, without relying on firms' self-reported employee classifications.

The main limitations of this data, relative to government administrative data, is completeness and precision. Not every worker creates an online job profile, and the workers who do choose to create a job profile are not a random sample. For example, one might expect low-wage workers to be less likely to create job profiles than high-wage workers. Some of these biases can be identified and largely corrected empirically. Revelio compares the total count of observed job profiles in a specific job-category and MSA and compares that count to BLS data on employment by job type. Revelio then weighs observed job profiles accordingly to generate an estimated actual number of employees in a specific firm-MSA-job category. Further, Revelio uses data sources that capture job profiles not just for office workers, but for non-office workers (e.g., Jobcase), such as retail, transportation, and other service type jobs.

The salaries in our data are based on job postings from the firms themselves. This allows us to capture average differences in wages across firms but does not capture post-offer negotiations between employees and employers. Thus, we cannot analyze wage differentials within firm, job

category and seniority with confidence and instead focus on differences: (a) within job category and across firms in the same MSA; and (b) differences in turnover rates within the same firm-job category-MSA over time.

Key Variables

Our main independent and dependent variables are relative wages and the annualized attrition rate, respectively. We calculate these variables at the firm-month-MSA-job category-seniority level as follows:

Attrition. We calculate attrition as the number of employees who separate from a role in a month, divided by the total number of employees in that role. Both are measured at the firm-month-MSA-job category-seniority. To ease interpretation, for regression models we use the annualized attrition rate, scaled such that a value of 25 indicates that the attrition rate would result in one quarter of the workers leaving that role each year. Attrition is calculated as:

$$\text{Annualized Attrition Rate} = \frac{\text{firm employee separations}}{\text{employee count}} * 100 * 12$$

Relative Wages. We calculate an employee group's relative wages as the percentage employees are paid above or below the median wage for employees in the same job category, seniority and MSA. Formally, relative wage is calculated as follows, where median wage is the median wage at the MSA-month-job category-seniority.

$$\text{Relative Wage} = \frac{\text{wage} - \text{median_wage}}{\text{median_wage}}$$

Controls

The granularity of our data allows us to accommodate a variety of possible sources of variation with fixed effects. To control for time-invariant industry or firm level characteristics, we use industry or firm fixed effects. To capture seasonal and national macroeconomic conditions, we use year-month fixed effects. To capture location effects and time-varying local economic conditions, we use

MSA-year fixed effects. To control for expected variation in turnover across job types, we use job category fixed effects. To adjust for different turnover rates among small and large firms, and small and large local offices, we include controls for both employee counts at the most granular level (firm-MSA-month-job category-seniority) as well as the employee count of the firm.

Sample and Summary Statistics

Our sample covers over 7,000 US firms, including both public and private firms, over the 30 months from January 2018 through June 2020. Our sample includes only US-based workers. A public (private) firm appears in our dataset if, at any point between 2008 and 2020 they are among the 3,000 largest public (private) firms according to Revelio's estimate of their employee count.

Our observations are at the firm-month-MSA-job category-seniority level. We restrict our sample to observations with at least 10 employees, which greatly reduces both the number of observations and firms in our sample.

A summary of the sample, with average statistics by year, is presented in Table 25. We observe between 5,000 and 7,000 firms each month, with more firms in the sample in later months. We observe workers 415 different Metropolitan statistical areas (MSAs), and our data classifies employees across 150 job categories and 4 levels of seniority. This level of granularity theoretically allows for over 1 billion possible firm-MSA-job category-seniority combinations, but of course firms do not hire employees of every job category or seniority or MSA. We observe between 10,000 and 18,000 different firm-MSA combinations in our data. Months in our data average between 30,000 and 70,000 observations. Over the full sample, the data contains 78,220 firm-MSA-job category-seniority groups, 61% of which appear in all 30 months.³⁶ We observe between 1 million and 3.7

³⁶ Not reported in tables. The remaining 39% of the groups appear in anywhere from one month (2.2% of groups) to 29 months (0.9% of groups) in the data.

million employees in any given month, and we observe average attrition rates between 11% and 26%.³⁷

[Insert Table 25 Here]

We provide an example of the most common job categories we observe, along with summary statistics and the wage distribution we observe for those jobs in June 2019, in Table 26. The top job categories we observe are manager, customer service representative, and operator. Customer service representatives have lower salaries than many of the jobs we observe, with half of their salaries falling between \$57,016 and \$72,360. Many high paying jobs make the list of most commonly observed jobs, with “Vice President,” “Software Engineer,” and “Doctor” all among the top 30 jobs. As we draw inferences from our analysis, we note that our sample is representative of largely, but not exclusively, white collar jobs.

[Insert Table 26 Here]

Table 27 displays summary statistics for our main variables of interest. We observe on average 47.2 employees per observation (that is, 47.2 employees within a single firm/MSA/job category/seniority observation). The distribution is highly skewed, ranging from 10 (our researcher imposed minimum) to 15,961 employees. The median observation contains 19.3 employees. The median wage in our sample is \$81,068, above the median household income in the US in 2019 of \$68,703, according to the Census Bureau (Semega *et al.*, 2021). Our sample is biased toward including more high wage workers relative to low wage workers. Likely this reflects the underlying data sources that are more complete for higher wage workers. While Revelio uses machine learning

³⁷ The lower average attrition rate in 2020 in our data is likely due to our data collection period ending in June 2020 and employees who recently exited roles not immediately reporting these separations on their online job profiles. Thus, in the last 6 months of our data we are likely undercounting attrition rates. While we do not expect a strong correlation between which employees are late to report attrition and our other variables of interest, and, thus, we do not have strong reason to expect biased estimates from this missing data, this data issue will be corrected when we receive updated data from our data provider.

and data from other sources, such as Bureau of Labor Statistics, to correct for this over representation of workers within certain (higher paying) job categories, this suggests that the correction is imperfect.

The distribution of yearly attrition rates is highly skewed. For many of our observations, there is no observed attrition, as would be expected given that many observations have less than 20 employees and our observations track employee groups monthly. For some observations, the annualized turnover rate is greater than 100%, meaning more than 1/12 of the employees turned over in a month. The average yearly turnover rate is 25.6%, which translates into 2.1% of employees separating from firms each month, lower than the rate reported by the Bureau of Labor Statistics.³⁸ This indicates that our data could be underestimating separation rates in our empirical analysis, or that our sample of slightly higher paid employees has lower turnover than the average worker in the US economy. As a result, we may be underestimating any effects of wages on turnover in our empirical models.

We use relative wage as our key independent variable. It has, by construction, a mean and median equal or close to zero. Relative wages vary from -0.71, meaning employees are paid 71% below the median wage of their job category-seniority-MSA peers, to 2.49, meaning employees are paid two and half times their median peer. To give meaning to coefficients on relative wages, we generally examine what an increase of 0.1 in the relative wage variable (a 10% increase in relative wages) would imply for attrition rates.

[Insert Table 27 Here]

Results

³⁸ For most of 2018 and 2019, the BLS reported a total separation rate of between 3.5% and 3.8% per month. Data is accessible at <https://www.bls.gov/jlt/data.htm>.

Correlations between our variables of interest are presented in Table 28. Wages are negatively correlated with attrition (rho of -0.02), as are relative wages, though the correlation is quite small.

[Insert Table 28 Here]

To test whether this correlation is accounted for by firm, location, and/or job factors, we regress the annualized attrition rate on relative wages with varying controls in Table 29. We include fixed effects for year-month to account for time varying macroeconomic conditions, MSA-year to account for local economic conditions, job category and seniority to account for different turnover rates by job type, and we also include industry or firm fixed effects depending on the model. Time-varying controls for the numbers of employees in the observation, and the total employees at the firm are included in the regressions and omitted from the table. Column (1) includes only year-month fixed effects, while column (5) includes all controls except for industry, as they wash out due to the inclusion of firm dummies.

[Insert Table 29]

The coefficient on relative wages is negative in all models, ranging from -0.65 to -2.6, with p-values all below 0.001. A coefficient of -2 means that a 10% increase in wages, relative to the median comparable job, is associated with a 0.2 percentage point decrease in yearly attrition rates. At the mean attrition rate of 13.6%, this 0.2 percentage point decrease is equivalent to a 1.5% decrease in attrition rates. Thus, across our whole sample, our correlational results are consistent directionally with past work, as they show a modest negative relationship between wages and attrition rates.

We next examine how the relationship between wages and attrition varies across wage bands by splitting our sample across wage quintiles. For this quintile split, we use the specification in column (4) which includes month, MSA-year, job-seniority, and industry fixed effects. Results are displayed in Table 30. Our sample is split by wage quintile within each month. The lowest quintile of wages

includes workers who earn below around \$62,000 (slightly less in January 2018, and slightly more in December 2021).

[Insert Table 30 Here]

Counter to our expectations, the negative relationship between wages and attrition we observed in the full sample is driven by workers in the middle of the wage distribution, not low-wage workers. Among the lowest quintile of wage earners in our sample (those earning below \$62,000), there is a positive correlation between wages and attrition rates. The coefficient for lower wage workers, displayed in model (1), is 3.1, indicating that a 10% increase in salary, relative to the median for comparable jobs, is associated with a 0.3 percentage point increase in annual attrition rate. For the 2nd through 4th quintile of wages, the coefficient on relative wage ranges from -2.1 to -5.9.

We next run the same wage quintile regressions but using our instrument to estimate the causal effect of wages on attrition. We instrument relative wages using the wages paid to workers of competitor firms in MSAs in which the focal firm has workers. To calculate the instrument, we exclude all workers in the same MSA or job category as the focal observation. These instrumented results are displayed in Table 31. The instrument is strong. The first-stage F-statistic in our regressions ranges from 1,476 to 5,519. Our instrumented results match the correlational results in sign but all show significantly larger coefficients. For lower wage workers (column 1), the coefficient on relative wages is 326. This coefficient indicates that a 10 percent increase in wages, relative to the median similar job, causes a 32.6 percentage point increase in annualized attrition rates. This is equal to more than doubling the annual average attrition rate in our sample. For the middle quintile of workers in our sample by wages, a 10 percent increase in wages, relative to the median is associated with a 9.8 percentage point decrease in annual attrition rates.

Overall, we do not find support for our first hypothesis. Among workers earning less than \$62,000, we find a significant positive correlation between wages and attrition, controlling for a wide

variety of industry, firm, job type, seniority, and geographic factors. For all other workers we find a negative relationship. Our causal identification strategy does produces results that are consistent directionally, but which are much stronger.

[Insert Table 31 Here]

To test our second hypothesis regarding employees outside options, we split our sample by the size of the local labor pool. We hypothesized that workers will be most responsive to wages at moderate levels of labor pool size, defined as the number of employees in the same MSA working in the same job category. Within each month, we split the sample into quintiles by the number of workers in each MSA job category group and run separate regressions for each quintile. The non-instrumented results are shown in Table 32. Like the previous wage quintile regressions, we use the model with month, MSA-year, job category-seniority and industry fixed effects. Columns (1) through (5) display the results of the first through fifth quintiles, respectively, of the number of workers within the MSA-job category. In the non-instrumented models, the relationship between relative wages and attrition is negative across all market sizes, except for the largest markets. The 1st, 3rd, and 4th quintiles of market size have coefficients on relative wage of around -1.5, which indicates a 10% increase in wages, relative to the median comparable job, is associated with a 0.15 percentage point decrease in annualized turnover.

[Insert Table 32 Here]

We next use our instrumental variables design along these job market size quintiles. The results are shown in Table 33. In these quintiles, our first stage is again strong, with first stage F-statistics always greater than 2,800. In the instrumented regression, all coefficients are negative, and they are larger by an order of magnitude compared to the non-instrumented models. In the smallest quintile of markets (column 1), the coefficient on relative wages is -52. This indicates that a 10% increase in wages relative to the median comparable job causes a 5.2 percentage point decrease in annualized

attrition rates (equal to a 38% drop in attrition rates relative to the mean). The coefficients for the remaining quintiles are smaller but still economically large and statistically significant, ranging from -10.5 to -29.4.

In summary, across labor market size quintiles, the non-instrumented regressions indicate small and negative correlations between wages and attrition in all but the largest job markets. The instrumented regressions give causal evidence of large and significant negative relationships between wages and attrition. While the overall negative coefficients match our expectations, these size of these coefficients across job market size quintiles do not match our expectations. We do not find evidence that there is a U-shaped relationship between the wage-attrition relationship and the size of the local labor market.

[Insert Table 33 Here]

Discussion

Workers are increasingly unattached from the firms they work for. The average tenure with a company is decreasing over time, and younger workers are showing less professional attachment to their employers. As a result, managers must determine how to retain their employees. In this paper, we examine whether and when pay decreases attrition rates across the US between 2018 and 2021. We specifically examine pay relative to similar workers in the same geographic area, providing a precise comparison of workers' outside wage options. Consistent with past research, we find that increasing pay is associated with moderate decreases in attrition rates. However, our results indicate that these relationships vary significantly based on wage levels and the size of the local labor market. Surprisingly, we find that the relationship between relative wages and attrition rates is positive among low-wage workers, with some causal evidence of this relationship. We also present causal evidence that the relationship between wages and attrition rates is most negative among the smallest local labor markets.

Our work has significant implications for managers attempting to avoid the negative impacts the great resignation may have on their companies. Managers want to keep workers to retain firm-specific human capital. Losing and having to hire new workers can result in increased recruitment and training costs, and decreased efficiency and product quality. Our results suggest that managers ought to turn to increased pay in many circumstances. However, increasing pay may not be effective among the highest and lowest paid workers. In these cases, managers may have to turn to other tools available to reduce attrition, such as increased benefits, better working hours, or by creating attachment to the firm through firm culture.

Significant additional work can be done to examine the causes of the results presented in this paper, and the surprising nature of wage-attrition relationship among low-wage workers specifically is cause for greater exploration into what is driving that relationship. Given the rich data available, future work can examine a wide possible range of explanations. Our causal inference strategy is driven by multi-location firms, and as such they may have an outsized impact on our results. Future work could separate large and small firms in the sample. Our lag structure of a single month may also be insufficient to capture the full effect of wages on attrition, as previous work has noted that workers take several months to respond to compensation changes (Sandvik *et al.*, 2021). We also are interested in whether firm-specific effects may be driving significant variation in both attrition rates and wage setting policies, and future iterations could include firm, rather than industry, fixed effects in the quintile regression models. Further firm-specific characteristics such as growth or contraction may also impact attrition rates. Including additional controls for these factors could reduce the noise in our models, and thus better identify the specific effect of wages on attrition.

Our paper also opens the door for research into the heterogeneity in the wage-attrition relationship by industry, firm, and worker type. Future research could ask for which industries the wage-attrition relationship is strongest. For example, do service industry workers respond more

strongly to increased wages than manufacturing workers, perhaps due to the specialized skills different industries demand? Turning to firms, research could ask whether the wage-attrition relationship varies based on whether firms are growing or shrinking, new or old, or high or low tech. Researchers could also turn to workers and ask whether younger or older workers respond more strongly to wages, or whether other demographic factors predict whether increased wages reduce attrition. Worker characteristics such as political affiliation or performance could also predict the strength of the wage-attrition relationship. These conditionalities could inform managers as to when wages are an effective lever to retain their firm-specific human capital and competitive advantage, and when they should turn to alternate methods to entice workers to stay with the firm.

Tables and Figures

Table 24: Example Rows of Data

Company	MSA	Month	Job Category	Seniority	Employees	Inflow, External	Inflow, Internal	Outflow, External	Outflow, Internal	Salary
Apple	Charlotte, NC	June 2018	Sales Manager	1 (Junior)	12.3	2.1	1.1	1.1	0.5	\$43,543
John Deere	Charlotte, NC	June 2018	Sales Manager	1 (Junior)	32.1	0.1	0.1	0.1	0.1	\$39,988
Sprint	Charlotte, NC	June 2018	Sales Manager	1 (Junior)	44.3	2.6	1.5	3.1	2.1	\$32,500
				...						
Apple	Charlotte, NC	June 2018	Sales Manager	2	8.1	2.2	0.1	0.3	2.5	\$46,402
				...						

Variable values are simulated in this table due to preserve confidentiality of the disaggregated data. Table is shown to demonstrate the granularity and format of the data. Employee counts, inflows and outflows are real numbers, not integers, because they are estimated algorithmically based on observed employee counts, and then adjusted for the fact that not all employees create job profiles.

Table 25: Sample Description

Year	Firms	Firm-MSA Groups	Firm-MSA-Job- Seniority Groups	Total Employee Count	Attrition Percentage
2017	8,958	95,920	261,246	7,988,089	12%
2018	8,985	99,439	274,915	8,475,315	13%
2019	9,000	103,595	292,173	9,008,968	13%
2020	8,969	105,567	300,397	9,189,923	13%
2021	8,904	108,508	317,394	9,643,250	16%

Note: All statistics are averaged across all months in each year

Table 26: Job Category Wage Distributions, July 2019

Job Category	Employees	Firms	MSAs	Annualized Attrition	Wage			
					Average	p25	Median	p75
Cashier	1,000,434	2,819	401	14%	\$53,186	\$45,928	\$47,566	\$51,627
Crew Member	994,031	2,908	400	15%	\$46,699	\$40,872	\$42,822	\$48,550
Distribution Specialist	749,607	3,255	400	8%	\$56,036	\$50,038	\$52,392	\$56,197
Software Engineer	642,943	2,348	285	14%	\$97,238	\$87,749	\$92,810	\$103,360
Financial Advisor	615,345	1,896	396	9%	\$76,133	\$68,107	\$72,716	\$77,075
Medical Rep	560,077	1,660	392	11%	\$69,196	\$59,957	\$65,640	\$77,317
Machine Operator	470,565	3,370	399	6%	\$64,819	\$53,951	\$60,699	\$74,818
Mechanical Engineer	369,185	2,382	387	8%	\$88,269	\$81,628	\$83,946	\$94,066
Receptionist	301,346	1,891	373	11%	\$67,572	\$52,666	\$64,070	\$76,505
Sales Associate	283,032	3,159	372	15%	\$88,316	\$77,949	\$82,383	\$96,807
Customer Service	208,183	2,376	370	12%	\$66,201	\$58,503	\$65,642	\$75,395
Accountant	190,896	3,153	326	15%	\$86,095	\$74,482	\$79,844	\$94,196
Coordinator	186,569	1,583	330	19%	\$63,973	\$51,502	\$57,666	\$75,575
Sales Representative	178,416	1,991	364	18%	\$64,611	\$46,455	\$53,458	\$77,943
Investment Specialist	176,428	1,100	215	15%	\$111,480	\$81,002	\$114,725	\$138,763
Retail Sales	156,427	1,547	377	13%	\$72,774	\$61,368	\$69,640	\$79,782
Legal	144,457	1,188	214	13%	\$102,612	\$78,421	\$96,858	\$114,937
IT Specialist	141,199	2,332	302	12%	\$84,320	\$81,058	\$82,784	\$84,836
Production Operator	138,446	1,690	369	10%	\$84,415	\$76,506	\$81,074	\$92,070
Operations Admin.	126,846	3,111	356	9%	\$61,532	\$53,564	\$59,222	\$68,623
Claims Specialist	108,178	965	297	9%	\$75,904	\$72,410	\$75,525	\$77,316
Producer	94,148	1,055	292	21%	\$75,203	\$62,518	\$68,758	\$81,069
Application Engineer	89,031	592	168	9%	\$100,244	\$87,737	\$94,027	\$106,754
Merchandiser	72,792	464	272	19%	\$62,520	\$54,090	\$57,186	\$71,935
Operations Manager	72,528	3,174	352	3%	\$121,613	\$95,898	\$123,045	\$142,651
Human Resources Specialist	71,157	2,331	297	15%	\$80,879	\$73,173	\$77,483	\$86,614

Notes:

Table 27: Summary Statistics, January 2017 through December 2021

	Mean	Std Dev	Min	p25	Median	p75	Max
Employee Count	30.6	133.4	5.0	6.8	9.3	19.8	28,360.6
Wage	\$76,419	\$25,744	\$14,507	\$57,212	\$74,342	\$88,795	\$164,069
Annualized Attrition	13.60%	76.35%	0%	0.02%	0.13%	1.38%	1200%
Firm Employment	7,989.65	19,469.37	5.0	588.3	1,827.6	5,956.0	268,091.8
Relative Wage	0.01	0.15	-0.8	0.0	0.0	0.1	5.8

Observations are of groups of workers at the Firm-MSA-K150 Job Category-Seniority level.

Sample is restricted to observations with at least 10 employees. There are a total of 17,353,500 observations. All variables are fully populated in the sample.

Wage is the average yearly salary paid to employees and is estimated by Revelio using job posting data.

Wage Relative to Median is defined as the focal firm's wage minus the median wage, all divided by the median wage. For example, a value of 0.05 means that the focal firm pays wages 5% above the median wages of employees of the focal K150 job category and seniority working at other firms in the same MSA.

Annualized attrition can be above 100% if a position turns over multiple times during the year.

The number of employees (count) is based on Revelio's algorithms for estimating the actual number of employees at the Firm-MSA-K150-Seniority level. This is estimated based on job profile data.

Table 28: Correlations

	(1)	(2)	(3)	(4)	(5)	(6)
(1) Employee Count	1.00					
(2) Wage	-0.04	1.00				
(3) Wage Relative to Median	-0.01	0.31	1.00			
(4) Attrition Pct	-0.00	-0.02	-0.00	1.00		
(5) Firm Employee Count	0.22	-0.01	0.01	0.01	1.00	
(6) Attrition Count	0.63	-0.04	-0.01	0.32	0.15	1.00

There are 17,353,500 observations. All correlations are statistically significant at the 1% level.

Table 29: Relationship between relative wages and annualized attrition

January 2018 through June 2020					
	(1)	(2)	(3)	(4)	(5)
Relative Wage	-2.037*** (-16.18)	-2.615*** (-20.38)	-2.348*** (-18.25)	-0.970*** (-7.42)	-0.652*** (-4.48)
Fixed Effects					
Month	Yes	Yes	Yes	Yes	Yes
MSA Year	No	No	Yes	Yes	Yes
Job-Seniority	No	Yes	Yes	Yes	Yes
Industry	No	No	No	Yes	No
Firm	No	No	No	No	Yes
Observations	17,353,500	17,353,500	17,353,500	17,353,500	17,353,500

T-statistics in brackets. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A coefficient of -2 means that a 10% increase in wages, relative to the median wage for similar jobs, is associated with a 0.2 percentage point decrease in annualized attrition rates.

Dependent variable is the annualized employee attrition percentage, equal to the number of employee-employer separations in a month, divided by the number of employees, times twelve, times one hundred.

Relative Wages is calculated as employee wages minus median wages of similar employees of the same job category and seniority within the same MSA in the same month, divided by the same median salary.

Observations are at the firm, MSA, k150 job category, seniority level

Sample restricted to observations with a minimum of 10 employees.

Time-varying controls included and not displayed: number of employees at the observation level, total number of employees in the firm

Table 30: Wage Quintile Regressions, No Instrumentation

	(1)	(2)	(3)	(4)	(5)
	Low Wage	2nd Quintile	Mid Wage	4th Quintile	High Wage
Relative Wage	3.108*** (7.83)	-3.282*** (-6.65)	2.090*** (3.79)	-5.864*** (-8.91)	-1.892*** (-7.22)
Obs	3,471,037	3,470,568	3,470,616	3,472,148	3,469,131

Dependent variable is the annualized attrition rate, where a value of 20 indicates that, on a yearly basis, 20 percent of the workers in the focal job left the company.

Column (0) includes all observations in our sample. Columns (1) through (5) each include one-fifth of the observations in our sample, split by the level of wages. Column (1) includes all observations in the lowest quintile of wages, column (2) the second quintile of wages, etc.

Controls for the month, MSA-year, job category-seniority, and industry are included in all regressions.

T-statistics in parenthesis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 31: Wage Quintile Regressions, Instrumented

	(1)	(2)	(3)	(4)	(5)
	Low Wage	2nd Quintile	Mid Wage	4th Quintile	High Wage
Relative Wage	326.1*** (18.26)	-29.34 (-1.67)	-98.26*** (-4.79)	-98.28*** (-5.20)	1.813 (0.14)
Obs	3,144,696	2,929,829	2,947,656	2,836,699	2,606,916
R-Squared	0.0000	0.0116	0.0082	0.0130	0.0093
First Stage F-Stat	2,120	3,384	3,360	5,519	1,476

Dependent variable is the annualized attrition rate, where a value of 20 indicates that, on a yearly basis, 20 percent of the workers in the focal job left the company.

Columns (1) through (5) each include one-fifth of the observations in our sample, split by the level of wages. Column (1) includes all observations in the lowest quintile of wages, column (2) the second quintile of wages, etc.

Controls for the month, MSA-year, job category-seniority, and industry are included in all regressions.

T-statistics in parenthesis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 32: Job Market Size Quintiles, No Instrumentation

	(1) Small Market	(2) 2nd Quintile	(3) Mid-Sized Market	(4) 4th Quintile	(5) Large Market
Relative Wage	-1.354*** (-4.32)	-0.679* (-2.28)	-1.565*** (-5.17)	-1.411*** (-4.76)	0.918*** (3.51)
Obs	3,471,453	3,471,861	3,472,800	3,474,404	3,462,982
R-Squared	0.0097	0.0111	0.0115	0.0130	0.0144

Dependent variable is the annualized attrition rate, where a value of 20 indicates that, on a yearly basis, 20 percent of the workers in the focal job left the company.

Column (0) includes all observations in our sample. Columns (1) through (5) each include one-fifth of the observations in our sample, split by the number of employees in the same MSA and the same job category as the focal observation. Column (1) includes all observations with the fewest workers in the focal MSA job category, column (2) the second quintile, etc.

Controls for the month, MSA-year, job category-seniority, and industry are included in all regressions.

T-statistics in parenthesis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 33: Job Market Size Quintiles, Instrumented

	(1) Small Market	(2) 2nd Quintile	(3) Mid-Sized Market	(4) 4th Quintile	(5) Large Market
Relative Wage	-52.17*** (-4.99)	-10.48 (-1.83)	-21.97*** (-5.39)	-26.65*** (-7.39)	-29.42*** (-7.15)
Obs	1,783,524	3,063,814	3,209,412	3,222,659	3,186,387
R-Squared	0.0031	0.0111	0.0106	0.0114	0.0111
First Stage F	2,863	9,705	19,573	23,246	14,150

Dependent variable is the annualized attrition rate, where a value of 20 indicates that, on a yearly basis, 20 percent of the workers in the focal job left the company.

Column (0) includes all observations in our sample. Columns (1) through (5) each include one-fifth of the observations in our sample, split by the number of employees in the same MSA and the same job category as the focal observation. Column (1) includes all observations with the fewest workers in the focal MSA job category, column (2) the second quintile, etc.

Controls for the month, MSA-year, job category-seniority, and industry are included in all regressions.

T-statistics in parenthesis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix A: Non-Equilibrium Model Proof

Consider a firm’s decision to self-report their ESG policies and invest in costly ESG outcomes (which may be difficult or impossible to verify). There are three actors: a manager, a shareholder, and an environmental activist³⁹ (e.g., The Sierra Club). The manager invests $c(x)$ into ESG programs to achieve ESG outcomes of level x .⁴⁰ The manager self-reports the firm’s ESG policies at level r , which, for now, we will assume is also set by the manager (the agent for the firm). The manager’s preference for ESG is defined as α , and could be determined by the manager’s personal preferences for ESG investment (e.g., Chin *et al.*, 2013) or by financial incentives to achieve ESG goals (e.g., Flammer *et al.*, 2019). The shareholder’s preference for ESG is defined as β , with, for example, European Institutional investors (Dyck *et al.*, 2019) having a greater preference for ESG outcomes than U.S. hedge funds (DesJardine *et al.*, 2020). Higher values of α and β indicate a greater preference for ESG outcomes. The shareholder influences the manager’s choice of ESG policies and investment in outcomes through the use of voice (i.e., calling the manager), choice (capital allocation) and contracting around ESG issues in CEO compensation packages (Amel-Zadeh and Serafeim, 2018; Flammer *et al.*, 2019). The environmental activist punishes the firm for self-reporting high levels of ESG policies while having low ESG outcomes, indicated by the cost function $f(r - x)$ (Baron and Diermeier, 2007; King and McDonnell, 2012). This activist penalty increases with the concentration of activist oversight of the firm, indicated by δ .

The manager determines the firm’s ESG policies and outcomes considering their own preferences as well as the preferences of the shareholder and activist. The manager maximizes the following expression, which I call the manager’s utility (u^F : $u^F = \alpha x + \beta r - c(x) - \delta f(r - x)$)

The terms in the utility function are defined below, along with a column listing an example application of a utility firm.

Term	Definition	Investor-Owned Utilities Example
α	Manager’s preference for ESG	The top management team’s perception that lower carbon emissions will lead to higher profits via, for example, lower regulatory burden or higher consumer loyalty, and/or their perceptions that the firm has a social responsibility to reduce their carbon emissions.
x	ESG outcomes	Firm’s overall carbon emissions.
$c(x)$	Cost of ESG outcomes	The cost of new solar power plant or wind farm installations and/or the costs of modifying existing plants to reduce or capture emissions.
β	Shareholder’s preference for ESG	Investors’ demand for lowering carbon emissions. This could be based on the belief that reducing carbon emissions provides financial benefits through reducing regulatory costs, or from a belief that the firm has a moral obligation to reduce carbon emissions to lessen climate change.
r	ESG policies (self-reported)	A firm’s reported carbon reductions (which could be from selected projects, and not representative of the overall firm’s emissions), or their self-reported carbon-reduction policies.

³⁹ This could be an environmental or social activist. I use “environmental activist,” or “activist” for brevity.

⁴⁰ Greenwashing may also be unintentional due to an ESG policy producing a worse than expected outcome. For simplicity, I assume ESG investments produce known ESG outcome x , rather than noisy ESG outcome $x + \epsilon$.

δ	Level of activist oversight	Oversight of the utility's ESG activities by environmental activists. This could be, for example, the membership and activity level of the local Sierra Club.
$f(r - x)$	Penalty activist imposes	The expected cost of an activist campaign to the firm, and by extension the manager. For example, the Sierra Club could organize a campaign against the utility for exaggerating their carbon reductions.

Results when Investors Trust Self-Reported ESG Policies

I will always assume that the manager selects the firm's ESG investment, and thus determines the firm's ESG outcomes, x . For now, suppose the shareholder is trusting such that the manager selects the firm's ESG policies, and the shareholder believes the firm's level of ESG outcomes are equal to their level of self-reported ESG policies. Specifically, they believe (potentially incorrectly) that $x = r$. The game proceeds as follows.

1. The manager, F , observes her type α (the shareholder's type, β , is public knowledge). The manager then selects a level of ESG investment which produces ESG outcomes x , and selects a level of self-reported ESG policies, r .
2. The firm's utility is determined based on $\{x, r\}$.

Under this initial assumption the manager sets the value of both x and r , the manager is the only actor, and all given values and functions are common knowledge: $\{\alpha, \beta, \delta, c(\cdot), f(\cdot)\}$. I assume the cost and penalty functions are both increasing and convex in their arguments. The equilibrium under these assumptions is simply the solution to the following maximization problem: $\max_{x,r} u^F = \max_{x,r} [\alpha x + \beta r - c(x) - \delta f(r - x)]$

The complete proofs for all equilibria (along with all function and parameter assumptions) are in the Appendix. For now, I simply present an overview of the results of the model. The manager's optimal choice of ESG policies, r^* , and ESG outcomes, x^* , are as follows:

$$x^* = c_x^{-1}(\alpha + \beta)$$

$$r^* = x^* + f_g^{-1}(\beta)$$

The manager faces incentives to increase ESG outcomes up to the point where the marginal benefits (from both the αx term and a reduction in the penalty function) are equal to the marginal costs (from the cost function). They also have the incentive to set ESG policies at the point where the marginal benefits from ESG policies (from the βr term) are equal to the marginal costs of ESG policies (from the increased activist penalty). Depending on the relative importance of the manager's own preferences, the shareholder's preferences, and the activist's penalty, the manager will want to understate, accurately report, or overstate the firm's ESG.⁴¹ If the penalty from activists is high relative to the benefits of self-reporting a high level of ESG policies, a firm will brownwash:

⁴¹ Here, I use overstate to mean they report higher policies than they achieve outcomes, that is choose $r|r > x$.

underreport their ESG (see Carlos and Lewis, 2018 or Kim and Lyon, 2015 for examples). Alternatively, if benefits to self-reporting a high level of ESG policies are high relative to the penalty from activists, the firm will greenwash.

Under this trusting shareholder assumption, as the manager's and investor's ESG preferences change, the firm's ESG policies and ESG outcomes change in the following ways:

Increasing Manager's ESG Preferences ...	Increasing Shareholder's ESG Preferences ...
$\frac{\partial x^*}{\partial \alpha} = \frac{\partial [c_x^{-1}(x_i)]}{\partial \alpha} > 0$ $\frac{\partial r^*}{\partial \alpha} = \frac{\partial x^*}{\partial \alpha} > 0$	$\frac{\partial x^*}{\partial \beta} = \frac{\partial [c_x^{-1}(x)]}{\partial \beta} = \frac{\partial [c_x^{-1}(x)]}{\partial \alpha} = \frac{\partial x^*}{\partial \alpha} > 0$ $\frac{\partial r^*}{\partial \beta} = \frac{\partial x^*}{\partial \beta} + \frac{\partial [f_\beta^{-1}(\beta)]}{\beta} > \frac{\partial x^*}{\partial \beta} > 0$

These results mean that the higher the manager's preferences for ESG (α), the higher the firm's ESG policies (r) and ESG outcomes (x). In fact, as the manager's ESG preferences increase, the firm's ESG policies and ESG outcomes will increase at the same rate. As the shareholder's preferences (β) for ESG increases, again both ESG policies and ESG outcomes increase. However, unlike with manager preferences, when investor ESG preferences increase, the firm increases their ESG policies (r) more than they increase their ESG outcomes (x), leading to greenwashing.

The preceding solution assumes that shareholders respond to a firm's self-reported ESG policies as if they reflect the firm's ESG outcomes. This could be because investors trust firms' self-reported ESG policies, or because they face institutional pressures to do something, lack more credible information, and think their eventual customers won't notice the difference. This latter point would suggest a different foundational source of greenwashing.

Taking an equilibrium approach would require the model to account for shareholders rationally discounting firms' self-reported ESG policies. This equilibrium model would provide an alternative set of relationships between stakeholder preferences and ESG outcomes and self-reporting. Thus, this model framework provides two possible sets of predictions: one in which shareholders trust firms ESG policies, and another in which they rationally discount ESG policies as cheap talk.

Do investors use firms' self-disclosed ESG policies, or do they rationally discount self-disclosed information as cheap talk? Recent academic literature and articles in the popular press suggest that investors are forced to use ratings based significantly on self-disclosed data due to high demand for such products and the lack of better options. While investors are rationally concerned about greenwashing (Ramanna, 2021), the rise in the quantification of ESG metrics has resulted in an information environment where investors are left with little option other than relying on largely voluntary, non-standardized and unaudited ESG disclosures (Peirce, 2019). Investors rely on ESG ratings in making investment decisions (Amel-Zadeh and Serafeim, 2018), and yet they say their top concern about this data they use is greenwashing (Carlson, 2021). Furthermore, there are known issues with ESG ratings such as divergence (Berg *et al.*, 2020; Chatterji *et al.*, 2016) and the impact of firm resources on ESG scores (Drempetic, Klein, and Zwergel, 2020).

The non-equilibrium predictions presented above, in which shareholders trust ESG policies, are more consistent with this empirical evidence and with testimonials in the popular press. Additionally, they are consistent with the results I present below in the empirical section. As such, I present this trusting shareholder version of the model as my inference to the best explanation and relegate the equilibrium solution to the model to the appendix.

Given this inference to the best explanation, the propositions from the trusting outside stakeholder version of the model are formally stated as follows:

Proposition 1: As managers' ESG preferences increase, firms' ESG policies will increase at the same rate as their ESG outcomes.

Proposition 2: As shareholders' ESG preferences increase, firms' ESG policies will increase more quickly than their ESG outcomes.

Full Proof of non-equilibrium model

Assume there is a firm F with type $\alpha \in (\underline{\alpha}, \bar{\alpha})$. The firm selects a level of costly performance, x with cost $c(\cdot)$, and reported performance, r . The only cost from reported performance coming from a penalty function, $\delta f(\cdot)$, increasing in $(r - x)$. The one-sided equilibrium game proceeds as follows:

1. The firm, F , observes its type α . It then selects a level of performance x , and a reported level of performance, r .
2. Utilities are determined based on $\{x, r\}$.

Under this initial assumption that the informed agent sets the value of both x and r . The informed agent is the only actor, all given values and functions are common knowledge: $\{\alpha, \beta, \delta, c(\cdot), f(\cdot)\}$. The cost and penalty functions are both increasing and convex in their arguments. The equilibrium under these assumptions is simply the following maximization problem.

$$\max_{x,r} u^F = \max_{x,r} [\alpha x + \beta r - c(x) - \delta f(r - x)]$$

Benefits to actual performance are αx , where $\alpha > 0$. Benefits to reported performance are βr , where $\beta > 0$. Define the costs to actual performance and associated partial derivatives (indicated using subscripts) as $c(x)$, $\frac{\partial c(x)}{\partial x} \equiv c_x(x)$ and $\frac{\partial^2 c(x)}{\partial x^2} \equiv c_{xx}(x)$. Further, assume that costs as a function of actual performance are continuous, twice differentiable and convex: $c_x(0) = 0$, $c_x(x) \geq 0$ and $c_{xx}(x) > 0$. These assumptions ensure the minimum a firm will spend on performance is greater than 0.

There are no direct costs of broadcasting performance (as in cheap talk models, i.e. Crawford & Sobel, 1982). However, I assume firms face a penalty for misrepresenting their performance. Call this gap between reported and actual performance $g \equiv r - x$ and the cost of misrepresenting performance $f(g)$. Assume the cost to firms for misrepresenting performance is continuous, twice differentiable, and convex: using subscripts to denote partial derivatives, $f_g > 0$ and $f_{gg} > 0$. This formulation captures a variety of possible cost structures for misrepresentation. Firms may face a stochastic possibility of being caught misrepresenting performance, such as through investigative reporting or regulatory oversight, in which case $f(g)$ represents the expected value of the cost. Alternatively, firms may face a consistent penalty from a reported/actual performance gap. For example, a set of discriminating consumers may be able to identify firms with gaps between reported and actual performance and choose not to purchase goods from those firms.

Because firms know they will face this broadcast decision when investing in performance, firms will make their investment and reporting decisions simultaneously. Formally, the firm's problem is to maximize the following utility function:

$$u_i = \alpha x + \beta r - c(x) - f(g)$$

$$\max_{x,r} \alpha x + \beta r - c(x) - f(g)$$

Where $c_x(x) \geq 0$, $c_{xx}(x) > 0$, $g \equiv r - x$, $f_g > 0$ and $f_{gg} > 0$

To differentiate between indicate inverse functions and the multiplicative inverse, I use notation as follows:

Inverse functions are indicated by a superscript. If $f(g) = \beta$, then $g = f^{-1}(\beta)$.

Multiplicative inverses are indicated by a fraction. If $f(g) = \frac{1}{\beta}$, then $\beta = \frac{1}{f(g)}$

I know solve for the equilibrium:

$$\max_{x,r} u^F = \max_{x,r} \alpha x + \beta r - c(x) - f(g), \text{ such that } g = r - x$$

The partial derivatives of the utility with respect to the choice variables must be equal to zero.

$$\frac{\partial u}{\partial x} = \alpha - c_x(x) - f_x(g) = \alpha - c_x(x) - f_g(g)g_x = \alpha - c_x(x) + f_g(g) = 0$$

$$\frac{\partial u}{\partial r} = \beta - f_g(g)g_r = \beta - f_g(g) = 0$$

Solving for x^* and r^* using the systems of equations:

$$\alpha - c_x(x) + f_g(g) = \beta - f_g(g)$$

$$2f_g(g) = c_x(x) + \beta - \alpha$$

$$f_g(g) = \frac{c_x(x) + \beta - \alpha}{2}$$

Plugging back into the equation for $\frac{\partial u}{\partial x}$

$$\alpha - c_x(x) + f_g(g) = 0$$

$$\alpha - c_x(x) + \frac{c_x(x) + \beta - \alpha}{2} = 0$$

$$2\alpha - 2c_x(x) + c_x(x) + \beta - \alpha = 0$$

$$c_x(x) = \alpha + \beta$$

$$x^* = c_x^{-1}(\alpha + \beta)$$

Plugging back into the equation for $\frac{\partial u}{\partial r}$

$$\beta - f_g(g) = 0$$

$$f_g(r - x^*) = \beta$$

$$\frac{\partial x^*}{\partial \alpha} = \frac{\partial [c_x^{-1}(x_i)]}{\partial \alpha} > 0$$

$$\frac{\partial r^*}{\partial \alpha} = \frac{\partial x^*}{\partial \alpha} > 0$$

This optimal behavior suggests that increasing α , the pressure for real performance, increases actual performance (through higher investment in performance by the firm) and higher broadcast performance. For example, if solar power becomes more efficient and provides increased operational benefits (α_i increases), then the firm will invest more in solar panels, increasing actual environmental performance (x_i) and will increase broadcast behavior (r_i) by the same amount.

The partial derivatives with respect to β_i are as follows:

$$\frac{\partial x^*}{\partial \beta} = \frac{\partial [c_x^{-1}(x)]}{\partial \beta} = \frac{\partial [c_x^{-1}(x)]}{\partial \alpha} = \frac{\partial x^*}{\partial \alpha} > 0$$

$$\frac{\partial r^*}{\partial \beta} = \frac{\partial x^*}{\partial \beta} + \frac{\partial [f_\beta^{-1}(\beta)]}{\beta} > \frac{\partial x^*}{\partial \beta} > 0$$

These partial derivatives show that increasing β , pressures for broadcast performance, also increases both actual and broadcast performance. However, broadcast performance increases more than actual performance.

For example, if institutional investors demand higher reported environmental performance from a firm, the firm faces two countervailing forces. The firm is rewarded for increasing their reported performance but increasing their reported performance without increases their actual performance increases the penalty they pay for misrepresenting performance. As a result, the firm will increase its actual performance, but not as much as they increase their reported performance due to the convexity of real performance costs. Thus, the gap between the reported and actual performance increases.

Appendix B: Equilibrium Model Proof

Under the full equilibrium model, outside stakeholders rationally discount self-reported ESG)

Suppose now that there is an audience which rationally sets the firm's ESG rating, r , with the goal of setting the rating as close to the actual performance as possible. Specifically, the audience's rule is to set $E[x|r] = r$. The only communication between the inside stakeholder and audience is a message sent by the inside stakeholder to the audience, which can be thought of as a noisy signal of the firm's actual ESG performance, x . Under this assumption, the model follows the structure of the cheap talk model introduced by Crawford and Sobel (1982). The formal solutions are presented below, but first I present a verbal summary of the equilibrium:

The solution concept for the model is a partition equilibrium. In a partition equilibrium, firms are sorted into one or more groups corresponding to ranges of performance, and the firm's rating, r , is equal to the average performance of all firm types in the range under which their performance falls. As the average (absolute value of the) bias of firms increases, the outside stakeholder can less accurately sort firms. As the bias becomes large, eventually there is a point at which the only solution is a 1-partition equilibrium, in which all firm types would be given a single rating equal to the average performance of all firm types.

For example, say the parameters of the model are such that at most a two-partition equilibrium is possible (a one-partition equilibrium is always possible, but higher partition equilibria always produce greater utility when they are possible). In a two-partition equilibrium, two ratings are possible: r' and r , such that $r' > r$. If the firm is type $\alpha > \bar{\alpha}$, then the firm will choose a performance level and send a message such that they receive the rating r' . If the firm is of the type $\alpha < \bar{\alpha}$, they will choose a performance level and send a message such that they receive the rating r . For such an equilibrium to hold, the average performance of all firms with type $\alpha > \bar{\alpha}$ must be r' , given that these firms know they will receive the rating r' . The equivalent must hold for all firms of type $\alpha < \bar{\alpha}$. As a result, firms are accurately, but imprecisely, sorted into "high" or "low" performers.

The intuitive tension in the model is that if the firm has a bias to have a rating that is significantly different than their actual performance, the audience will discount any message they send. The audience treats the message as cheap talk, and will assign a conservative (and, on average, less accurate) rating to the firm. As the audience's expectation of the firm's bias moves closer to zero, the audience will be better able to distinguish between different firms based on the messages they send, and average ratings accuracy will increase. The interesting component of the model is that the firm's bias can be negative (underreporting ESG performance, $r < x$), or positive (greenwashing, $r > x$), and that this bias is a function of the inside and outside stakeholders' preferences and activist presence. If outside stakeholders have low preferences for ESG performance relative to inside stakeholders' preferences and the level of activist presence, firms will underreport (i.e., not report) their ESG performance so as to not draw activist attention. As a result, third party assessments of ESG performance reliant on firms' self-reported data would be imprecise and inconsistent across raters. The lack of attention to ESG issues from investors in fossil fuels relative to the presence of environmental activists in the 1990s would fit this scenario.

Alternatively, if outside stakeholders have a high demand for ESG performance, while inside stakeholders do not, firms will greenwash their ESG behavior, also leading to imprecise and inconsistent third party ESG assessments. This may presently be the case in some firms facing high investor demands for ESG performance, but with top management doubtful of climate science or

believing there is responsibility for business to reduce their harmful externalities. I expect this greenwashing effect to dominate, given the recent rapid rise in ESG investing.

From this two-sided equilibrium, I develop the following proposition:

As the outside stakeholder's preference for ESG performance increases (i.e. investors' increased commitment to ESG ratings over the past decade) relative to inside stakeholder's ESG preferences, the accuracy of these ratings will increase up to a point, and then decrease. The disagreement between ratings will be directly inversely correlated with ratings accuracy, so stated using ratings disagreement, this formulation of the model leads to the following propositions:

Proposition 3: As outside stakeholder preferences for ESG performance (β) increases relative to that of the inside stakeholder (α), ratings disagreement will fall and then rise, following a U shape.

Complete Proof:

Begin with the setup of the one-sided model, but now assume that the uninformed stakeholder rationally sets a firm's reported performance, r , as the expected value of their performance given that they receive the rating.

The game proceeds in three stages:

1. The firm, F , observes its type α , selects a level of performance x , and sends a message, n , to the audience, A . This signal can be random and can be viewed as a noisy estimate of the firm's actual performance x .
2. The audience sets a rating for the firm, r .
3. Utilities are established.

The parameters of the model are defined as follows.

Parameters

Assume there is a distribution of firms defined by $\alpha \in (\underline{\alpha}, \bar{\alpha})$ with a cumulative distribution function $G(\alpha)$ with density $g(\alpha)$ such that $G(\underline{\alpha}) = 0$ and $G(\bar{\alpha}) = 1$.

The benefits to reported performance are defined by the parameter β (as in the one-sided model). The rating is assigned by the audience (unlike in the one-sided model, where the rating is set by the firm) and is defined as $r \in \mathbb{R}$.

Utility Functions

The utility of the firm is defined as $u^F(x, r, \beta)$, a continuous and twice differentiable function. The utility of the audience is defined as $u^A(r, x)$, a continuous and twice differentiable function.

Partial Derivative Assumptions

I use the notation $\frac{\partial u^F(x, r, \beta)}{\partial x} \equiv u_x^F(x, r, \beta) \equiv u_1^F(x, r, \beta)$ and $\frac{\partial^2 u(x, r, \beta)}{\partial x^2} \equiv u_{xx}^F(x, r, \beta) \equiv u_{11}^F(x, r, \beta)$. Further, I use $u^F(\cdot)$ to indicate the utility function with the arguments in any order, so a statement about $u_1^F(\cdot)$ is a statement about the partial derivative with respect to any of the arguments.

Assume, for $i \in \{F, A\}$, that for some r , $u_{rr}^i(r, x) < 0$, meaning that u^i has a unique maximum r for a given x . Further, assume that $u_r^F(r, x, \beta) = 0$ for some r and $u_{rr}^F(r, x, \beta) < 0$,

so u^F has a unique maximum r for a given (x, r, β) . Finally, assume that $u^i(r, x) > 0$. This last assumption is a sorting requirement meaning that the best value of r for a fully informed agent is increasing in x .

Equilibrium

I use a Bayesian Nash Equilibrium to define the solution. While the firm's choice of performance and message precede the audience's ratings assignment in time, the audience sees only the message, and therefore all these actions are strategically simultaneous.

The equilibrium consists of a signaling rule for the firm defined as, $q(n|x)$, and an action rule for the audience defined as $r(n)$, such that

- 1) The firm's signaling rule maximizes expected utility for a given type, defined by α , taking the audience's ratings rule as a given. Specifically,

for each $\alpha \in (\underline{\alpha}, \bar{\alpha})$, $\int_N q(n|x)dn = 1$, where the Borel set N is the set of feasible signals, and if n^* is in the support of $q(n|x)$, then n^* solves $\max_{n \in N} u^F(r(n), x, \beta)$.
- 2) The audience responds optimally to each signal, using Bayes' rule to update their prior, taking into account $G(\alpha)$, the firm's signaling strategy, and the signal received. Specifically, for each n , $y(n)$ solves $\max_r \int_{\underline{\alpha}}^{\bar{\alpha}} u^A(r, x)p(x|\alpha)d\alpha$, where $p(\alpha|n) \equiv \frac{q(n|\alpha)g(\alpha)}{\int_{\underline{\alpha}}^{\bar{\alpha}} q(n,\alpha)f(\alpha)d\alpha}$

Modifications from the classical Crawford and Sobel (1982) model.

Assume the firm's utility takes the functional form from the one-sided model:

$$u^F = \alpha x + \beta r - c(x) - f(r - x)$$

Where $x \geq 0$, $c(x)$ is defined over the domain $x \in [0, \infty)$, $c(0) = 0$, $c_x(x) \geq 0$, $c_{xx}(x) > 0$, $g_i \equiv r - x$, $f_g > 0$ and $f_{gg} > 0$.


The one-sided equilibrium is solved in the main paper. Here I restrict my attention to the two-sided equilibrium. In this equilibrium, the firm takes r_i as a given and maximizes their utility by choosing x_i :

$$\begin{aligned} u_i &= \alpha_i x_i + \beta_i r_i - c(x_i) - f(r_i - x_i) \\ \frac{\partial u_i}{\partial x_i} &= \alpha_i - c_x(x_i) + f_x(r_i - x_i) = 0 \\ x_i^* &= \{x_i | \alpha_i - c_x(x_i) + f_x(r_i - x_i) = 0\} \end{aligned} \tag{1}$$

We are interested in the point where $\alpha_i + f_x(r_i - x_i) = c_x(x_i)$

Lemma 1: The point x^* is unique for a given (α, r) .

Proof: The function $c_x(x)$ is monotonically increasing and convex, and the function $\alpha + f_x(r - x)$ is convex and monotonically decreasing in x_i . Therefore, the conditions for x_i^* satisfy the single crossing property, and there will be only one point x^* at which $\alpha_i + f_x(r - x) = c_x(x)$. Therefore, $x^*|r$ exists and is unique for each $\alpha > 0$.



Lemma 2: x^* is strictly increasing in α .


Proof: From Equation (1), $x^* = \{x | \alpha - c_x(x) + f_x(r - x) = 0\}$. By assumption, $\alpha_i > 0$ and $c_x(x) > 0$. Since $\alpha - c_x(x) + f_x(r - x) = 0$, we know $c_x(x) > f_x(r - x)$. If we increment α , then $c_x(x) - f_x(r - x)$ increases. Suppose x decreased (or remained constant) as a result of incrementing α . Then, $c_x(x)$ would decrease, at a greater rate than $f_x(r - x)$ would increase (or both would remain constant if α remained constant). Then $c_x(x) - f_x(r - x)$ would decrease (remain constant), and we have a contradiction.

Lemma 3: If there are two possible ratings to select into, $r' > r$, there is a specific point $\bar{\alpha}$ at which a firm will be indifferent between rating r and rating r' . Furthermore, all firms with $\alpha > \bar{\alpha}$ will choose the r' and all firms for which $\alpha < \bar{\alpha}$ will select into r .

Proof: Assume a firm can select to send a message such that they will be assigned either the rating r or the rating $r' > r$. A firm will be indifferent when $u(x^* | \alpha, r) = u(x^* | \alpha, r')$.

From the previous proof, $\frac{\partial x^*}{\partial \alpha} > 0$. Similarly, $\frac{\partial x^*}{\partial r} > 0$.

Then the utility functions satisfy the single crossing property, and therefore there exists some $\alpha : u(x^* | \alpha, r) = u(x^* | \alpha, r')$. Therefore, firms with $\alpha > \bar{\alpha}$ will choose the r' and all firms for which $\alpha < \bar{\alpha}$ will select into r . A visual example is shown below:



Return to (Crawford and Sobel (1982) proof:

The previous section has assumed or established the following:

- 1) $\alpha \in (\underline{\alpha}, \bar{\alpha})$
- 2) $x^*(\alpha, r, \beta) > 0$ and $x^*_{\alpha\alpha}(\alpha, r, \beta) < 0$. Therefore $x(\alpha)$ is a monotonic transformation.

Define $\bar{m} \equiv \max \{x^*|r\}$, $\underline{m} \equiv \min\{x^*|r\}$, and define the function $m \equiv \frac{(\operatorname{argmax}_x u^F|r) - \bar{m}}{\bar{m} - \underline{m}}$.

Now, for a given set of ratings, m fits the required firm definition to fit into the Crawford and Sobel (1982) model of cheap talk.

Theorem 1:

Crawford and Sobel (1982) version: The number of partitions possible in equilibrium is increasing in the misalignment between the sender and receiver.

My model: The number of partitions possible is increasing, then decreasing in β . Specifically, $\frac{\partial^2 N(\beta)}{\partial \beta^2} < 0$.

Proof: The bias term is decreasing, then increasing in β . At the point $\beta = \beta^*$ where $0 < \beta^* < \infty$, the bias term equals zero. Therefore, the number of partitions possible is increasing, then decreasing in β .

Conclusions

The point x^* is unique for a given (α, r) , and further, x^* is strictly increasing in α and r (Lemmas 1, 2 and 3). The structure of the firm's utility is such that there is a single value of x which maximizes its utility if it takes its type (α) and ratings (r) as a given. This is due to the convexity of costs to actual performance (x) , concave benefits to reducing the reported/actual gap $(r - x)$, and linear benefits to performance (x) . Partial derivatives of the firm's utility function reveal that x^* is increasing in both α and r .

The main conclusion from the two-sided model relates to ratings accuracy. If reported performance gives firms zero benefit, but there are penalties increasing in $(r - x)$, firms will be biased to report the lowest possible value they can in order to minimize the penalty $f(r - x)$. As β approaches ∞ , the firm will garner such benefits from reported performance that every firm, regardless of their ability or performance, will prefer to garner the highest rating, thus eliminating the ability for the audience to distinguish between high and low type firms. In between the extremes of $\beta = 0$ and $\beta = \infty$, there is a range of β over which the firm's preferred rating coincides sufficiently closely with their actual performance. In these cases, firms prefer to be rated accurately, and signal as such. Resultingly, Theorem 1 states that the number of partitions possible is increasing, then decreasing in β .

This model makes the following additions to the model in Crawford and Sobel (1982). In this paper's model, firms select their level of performance, whereas in Crawford and Sobel (1982) the firm (which they call the sender) has an assigned value which they cannot modify. Further, the bias term may be different signed for different firms in the distribution of firms being rated, unlike in Crawford and Sobel (1982). Finally, these two changes are made explicit by the costs and benefits to firm performance and the firm's rating. The result of all these changes is the novel Theorem 1, as well as secondary propositions, all centered upon endogenously chosen performance levels.

Appendix C: Additional Tables from Chapter 1

Examples of ESG Related Articles from TruValue

See example articles from TruValue in Table 14, referenced in Chapter 2.

Sample of ESG Related Quotes from Earnings Calls

Some companies explicitly discuss their stakeholder orientation or strategy

BP Q3 2020	“Now, from the start of this COVID-19 crisis, we’ve been guided by a clear set of priorities that are very much in line with our multi-stakeholder business model to protect our people, safeguard supply, respond to new patterns of demand, obviously support our communities and preserve cash and balance sheet strength.” -BP, 2020 Q3
Unilever Q3 2020	“It’s never been more relevant for brands to demonstrate their positive contribution to society and address the issues that our consumers care about in an authentic way. So we’re investing more of our marketing spend on communication, which is explicitly purposeful,” Unilever, 2020 Q3
Unilever Q3 2020	“In addition, while the world continues to grapple with the devastating effects of COVID- 19 and the issues of escalating inequality, it’s more important than ever that we don’t lose sight of the climate crisis and the very real and serious threats that it creates for all of us. Climate change, nature degradation, biodiversity decline, water scarcity, they’re deeply interconnected problems and we are committed to find ways to tackle them simultaneously. So in June, we announced a wide-ranging set of commitments . . .” Unilever, 2020 Q3

Discussion of ESG only in the Q&A section may suggest ESG issues are a low priority

Exxon Mobile Q3 2020	[Question by Analyst] “I guess the first question, there’s been a lot certainly that you guys have talked about and that’s been written about, carbon, and a lot of your peers who have come out with explicit carbon targets, I just wanted to know where Exxon is in that journey in terms of coming out with carbon targets and how are you thinking about setting them to the extent that’s the path you choose to go down?” [Answer by Jack Williams, Senior VP, Exxon Mobile] “Yes. Thanks, Niel. Let me just take a step back a little bit on how we see things and what we're doing. Basically, we see a world that's going to need more energy going forward; through population, through GDP growth, and a lot of that goes through the non-OECD countries. Energy consumption is tied to the population and GDP. And we don't think the current solutions set is really complete." Exxon Mobile, 2020 Q3
----------------------------	--

Summary Statistics, Normalized Variables

Statistic	ESG Policies	ESG Outcomes	Negative ESG Outcomes	Earnings Call ESG (SASB)	Earnings Call ESG (Word count)	Investor ESG Orientation	Long Term Shareholders	Employees, log	Revenue, log	Assets, log
N	4,982	4,982	4,914	4,983	4,983	4,983	4,056	4,986	4,990	4,990
Mean	0.0	0.0	0.04	-0.04	0.00	0.05	-0.08	2.05	7.61	8.28
Std. Dev.	1.0	1.0	1.12	0.94	0.99	1.00	0.95	1.46	2.02	1.80
Median	-0.3	0.1	-0.16	-0.26	-0.21	-0.11	-0.09	1.84	7.66	8.16
Minimum	-2.1	-4.4	-0.33	-1.38	-2.01	-1.43	-3.97	0.00	0.00	3.86
Maximum	3.8	2.7	14.40	8.32	4.90	15.50	7.20	7.70	13.20	14.80

Figure 13: Earnings Call ESG Discussions by Sector

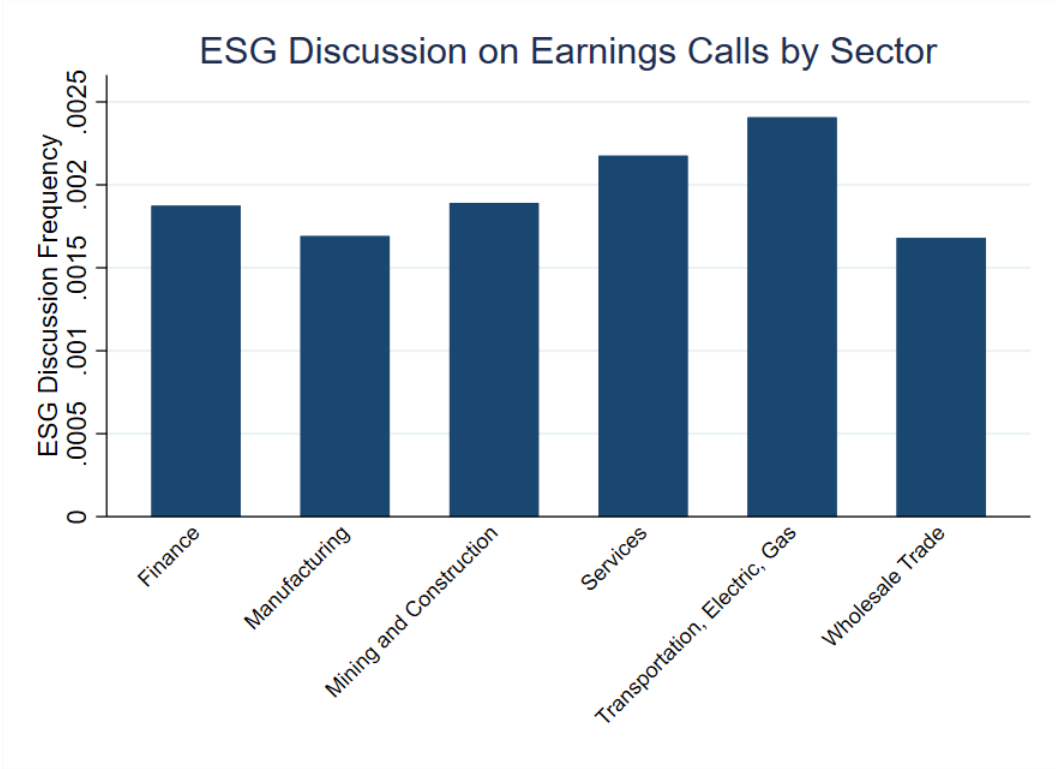


Figure 14: Earnings Call Environmental Discussions by Sector


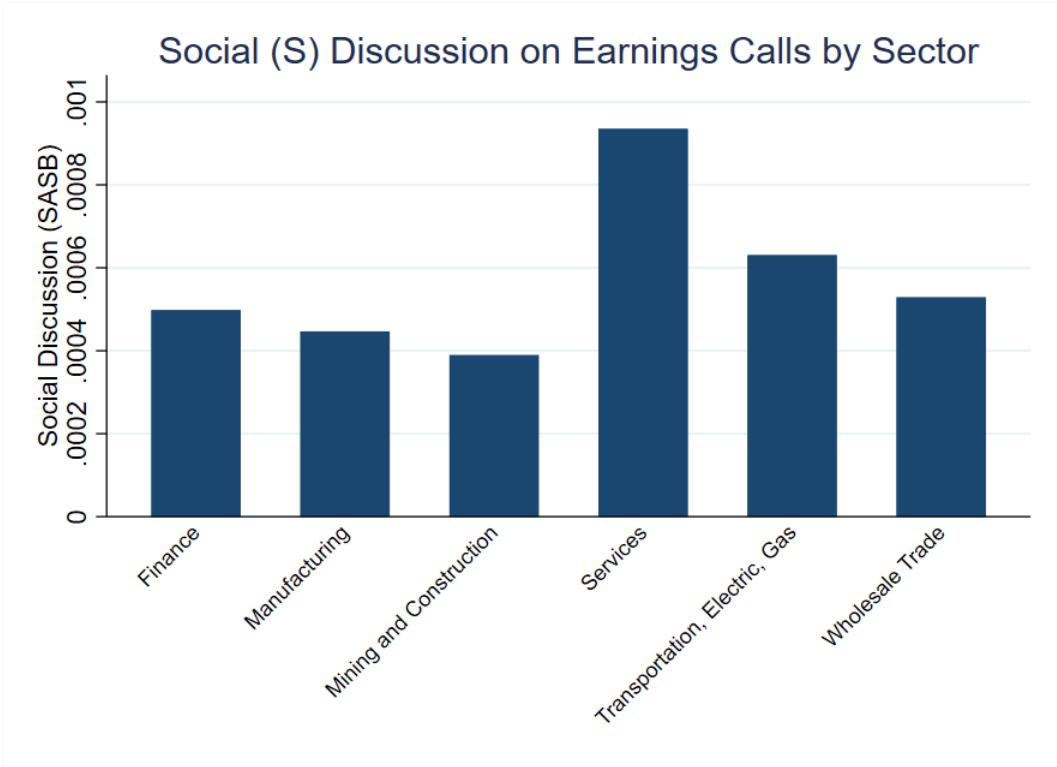


Figure 15: Earnings Call Social Discussions by Sector



Appendix E: Simulation Model of the Wage-Attrition Relationship

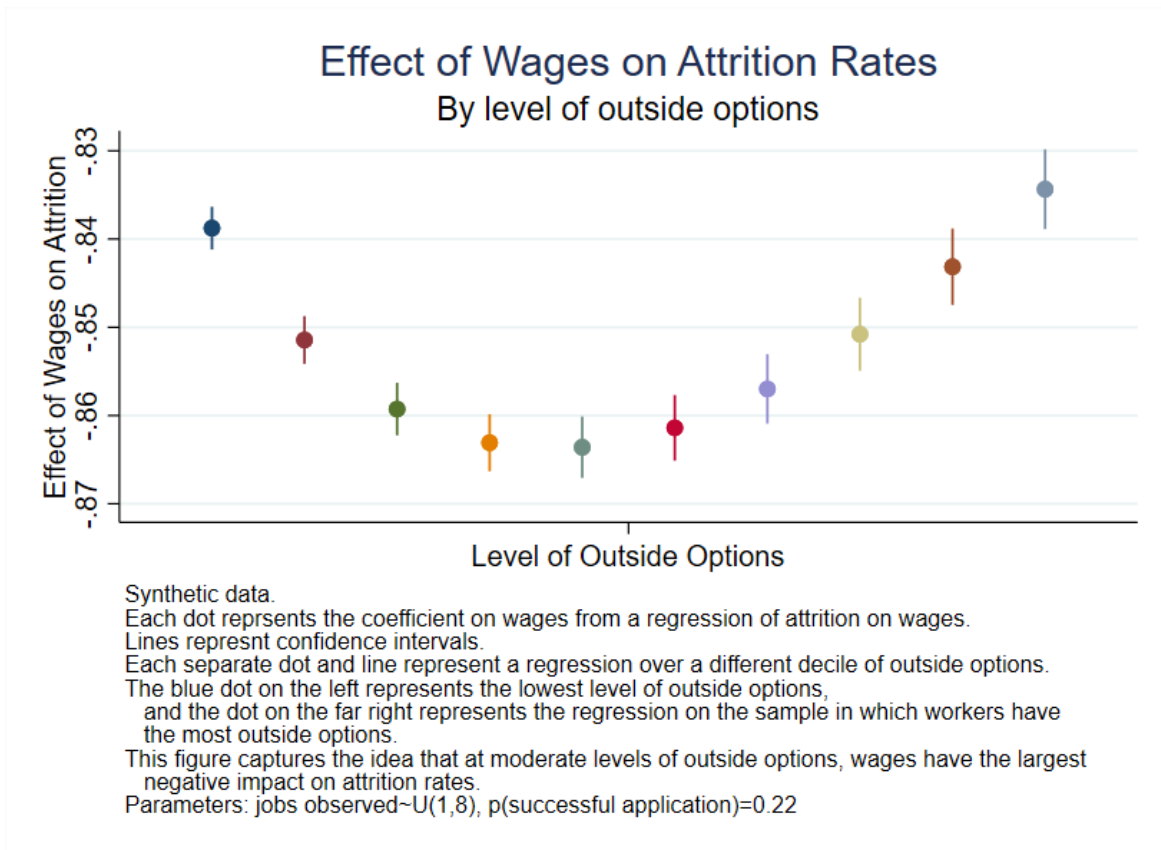
The following simple setup and simulated data add structure to the verbal arguments made in the conditionality section of the paper.

Say firms may pay workers $x_i \in (0,1)$. The level of outside options an employee has for employment is captured by the parameter t , which can be thought of as the number of jobs the employee can apply to each month. For each job the employee can apply to, they have a rate of success of receiving the job of q , which can be thought of the product of the probability of applying to a job times the likelihood of receiving a job offer conditional on applying. Assume workers only apply to jobs they observe which pay higher than their current job. In this setup, the probability that a worker paid x_i who sees a random job, which pays $x_j \sim U(0,1)$, both applies for and receives a job offer is $(1 - x_i)q$. The probability that a worker separates from their firm in a focal month, which we will call A for attrition rate, is as follows: $A = 1 - (1 - (1 - x_i)q)^t = 1 - (1 - q - qx_i)^t$

Under this simple model setup, we can generate predictions for how the relationship between wages and attrition vary depending on the level of outside options available to the employee. The relationship between wage and attrition is always negative because workers with higher wages are less likely to observe jobs which pay better than their current position. We then turn to the question of how the level of outside options influences the wage-attrition relationship. For a visual example, we fix some of our parameters and plot the wage-attrition relationships by market thickness in the

Figure 16. For the figure, we assume that workers are assigned a wage between 0 and 1, that they observe between 1 and 8 jobs per time period, and they have a 22% chance of receiving a job conditional on applying. While these parameters were chosen for their nice visual, the general patterns shown below holds for a wide range of parameters. The relationship between wages and attrition follows a U-shape when plotted against the level of employees' outside options. The relationship is most negative at moderate levels of outside options, while wages are less effective at reducing attrition when employees have very few or many outside options. The relationship is graphed visually in the figure below, which fixes some parameters to show the relationship visually.

Figure 16: Synthetic Data Example of Simulation Model



References

- Allen K. 2018, December 5. Lies, damned lies and ESG rating methodologies. *Financial Times*. Available at: <http://ftalphaville.ft.com/2018/12/06/1544076001000/Lies--damned-lies-and-ESG-rating-methodologies/>.
- Altria. 2019. *2018 Corporate Responsibility Progress Report*: 96.
- Amel-Zadeh A, Serafeim G. 2018. Why and How Investors Use ESG Information: Evidence from a Global Survey. *Financial Analysts Journal*. Routledge **74**(3): 87–103.
- Avetisyan E, Hockerts K. 2017. The Consolidation of the ESG Rating Industry as an Enactment of Institutional Retrogression: Consolidation of the ESG Rating Industry. *Business Strategy and the Environment* **26**(3): 316–330.
- Azar J, Berry ST, Marinescu IE. 2019. Estimating Labor Market Power. *SSRN Electronic Journal*. Available at: <https://www.ssrn.com/abstract=3456277>.
- Babcock AF *et al.* 2022. *Walking the Talk: Valuing a Multi-Stakeholder Strategy*. SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. Available at: <https://papers.ssrn.com/abstract=4023510>.
- Baron DP, Diermeier D. 2007. Strategic Activism and Nonmarket Strategy. *Journal of Economics & Management Strategy* **16**(3): 599–634.
- Bassier I, Dube A, Naidu S. 2021. Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies. *Journal of Human Resources*. Available at: <http://jhr.uwpress.org/lookup/doi/10.3368/jhr.monopsony.0319-10111R1>.
- Batt R, Colvin AJS, Keefe J. 2002. Employee Voice, Human Resource Practices, and Quit Rates: Evidence from the Telecommunications Industry. *Industrial and Labor Relations Review*. Sage Publications, Inc. **55**(4): 573–594.
- Berg F, Kolbel J, Rigobon R. 2020. Aggregate Confusion: The Divergence of ESG Ratings. *SSRN Electronic Journal*. Available at: <https://www.ssrn.com/abstract=3438533>.
- Berg F, Kolbel J, Pavlova A, Rigobon R. 2021. ESG Confusion and Stock Returns: Tackling the Problem of Noise. *SSRN Electronic Journal*. Available at: <https://www.ssrn.com/abstract=3941514>.
- Bowen F. 2014. *After Greenwashing: Symbolic Corporate Environmentalism and Society*. Cambridge University Press: Cambridge. Available at: <http://ebooks.cambridge.org/ref/id/CBO9781139541213>.
- Brochet F, Loumioti M, Serafeim G. 2015. Speaking of the short-term: disclosure horizon and managerial myopia. *Review of Accounting Studies* **20**(3): 1122–1163.
- Bryson A, Barth E, Dale-Olsen H. 2012. Do higher wages come at a price? *Journal of Economic Psychology* **33**(1): 251–263.
- Bundy J, Shropshire C, Buchholtz AK. 2013. Strategic Cognition and Issue Salience: Toward an Explanation of Firm Responsiveness to Stakeholder Concerns. *Academy of Management Review* **38**(3): 352–376.
- Burbano VC. 2016. Social Responsibility Messages and Worker Wage Requirements: Field Experimental Evidence from Online Labor Marketplaces. *Organization Science*. INFORMS **27**(4): 1010–1028.
- Bushee BJ. 2001. Do Institutional Investors Prefer Near-Term Earnings over Long-Run Value? *Contemporary Accounting Research* **18**(2): 207–246.
- Business Roundtable. 2019, August. Statement on the Purpose of a Corporation. Available at: <https://opportunity.businessroundtable.org/wp-content/uploads/2019/08/BRT-Statement-on-the-Purpose-of-a-Corporation-with-Signatures.pdf>.

- Carlos WC, Lewis BW. 2018. Strategic Silence: Withholding Certification Status as a Hypocrisy Avoidance Tactic. *Administrative Science Quarterly*. SAGE Publications Inc **63**(1): 130–169.
- Carlson D. 2021, February 2. BlackRock’s push on ESG and climate goals is coming at ‘a business-friendly pace’. *MarketWatch*. Available at: <https://www.marketwatch.com/story/blackrocks-push-on-esg-and-climate-goals-is-coming-at-a-business-friendly-pace-11612226740> [7 June 2021].
- Chatterji A, Toffel M. 2018. The New CEO Activists. *Harvard Business Review*.
- Chatterji AK, Durand R, Levine DI, Touboul S. 2016. Do ratings of firms converge? Implications for managers, investors and strategy researchers. *Strategic Management Journal* **37**(8): 1597–1614.
- Chava S, Du W, Malakar B. 2021, July 1. Do Managers Walk the Talk on Environmental and Social Issues? SSRN Scholarly Paper, SSRN. Available at: <https://papers.ssrn.com/abstract=3900814>.
- Chen T, Dong H, Lin C. 2020. Institutional shareholders and corporate social responsibility. *Journal of Financial Economics* **135**(2): 483–504.
- Cheng I-H, Hong H, Shue K. 2013. *Do Managers Do Good with Other People’s Money?* National Bureau of Economic Research, Cambridge, MA: w19432. Available at: <http://www.nber.org/papers/w19432.pdf>.
- Cheng I-H, Hong H, Shue K. 2014. Do Managers Do Good With Other People’s Money? Available at: <https://www-nber-org.proxy-um.researchport.umd.edu/papers/w19432.pdf>.
- Chin MK, Hambrick DC, Treviño LK. 2013. Political Ideologies of CEOs: The Influence of Executives’ Values on Corporate Social Responsibility. *Administrative Science Quarterly* **58**(2): 197–232.
- Crawford VP, Sobel J. 1982. Strategic Information Transmission. *Econometrica* **50**(6): 1431–1451.
- Dai R, Liang H, Ng L. 2020. Socially responsible corporate customers. *Journal of Financial Economics* : S0304405X20300325.
- Dale-Olsen H. 2006. Wages, fringe benefits and worker turnover. *Labour Economics* **13**(1): 87–105.
- Delmas MA, Burbano VC. 2011. The Drivers of Greenwashing. *California Management Review* **54**(1): 64–87.
- Delmas MA, Montes-Sancho MJ. 2010. Voluntary agreements to improve environmental quality: symbolic and substantive cooperation. *Strategic Management Journal* **31**(6): 575–601.
- DesJardine MR, Durand R. 2020. Disentangling the effects of hedge fund activism on firm financial and social performance. *Strategic Management Journal* **41**(6): 1054–1082.
- DesJardine MR, Marti E, Durand R. 2020. Why Activist Hedge Funds Target Socially Responsible Firms: The Reaction Costs of Signaling Corporate Social Responsibility. *Academy of Management Journal*. Academy of Management. Available at: <https://journals.aom.org/doi/10.5465/amj.2019.0238>.
- Dess GG, Shaw JD. 2001. Voluntary Turnover, Social Capital, and Organizational Performance. *The Academy of Management Review*. Academy of Management **26**(3): 446–456.
- DiMaggio PJ, Powell WW. 1983. The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields. *American Sociological Review* **48**(2): 147–160.

- Drempetic S, Klein C, Zwergel B. 2020. The Influence of Firm Size on the ESG Score: Corporate Sustainability Ratings Under Review. *Journal of Business Ethics* **167**(2): 333–360.
- Duhautois R, Gilles F, Petit H. 2016. Decomposing the relationship between wage and churning. *International Journal of Manpower*. Emerald Group Publishing Limited **37**(4): 660–683.
- Durand R, Hawn O, Ioannou I. 2019. Willing and Able: A General Model of Organizational Responses to Normative Pressures. *Academy of Management Review* **44**(2): 299–320.
- Dyck A, Lins KV, Roth L, Wagner HF. 2019. Do institutional investors drive corporate social responsibility? International evidence. *Journal of Financial Economics* **131**(3): 693–714.
- Fancy T. 2021, March 16. Financial world greenwashing the public with deadly distraction in sustainable investing practices. *USA TODAY*. Available at: <https://www.usatoday.com/story/opinion/2021/03/16/wall-street-esg-sustainable-investing-greenwashing-column/6948923002/>.
- Fiss PC, Zajac EJ. 2004. The Diffusion of Ideas over Contested Terrain: The (Non)adoption of a Shareholder Value Orientation among German Firms. *Administrative Science Quarterly*. [Sage Publications, Inc., Johnson Graduate School of Management, Cornell University] **49**(4): 501–534.
- Flammer C. 2015. Does Corporate Social Responsibility Lead to Superior Financial Performance? A Regression Discontinuity Approach. *Management Science* **61**(11): 2549–2568.
- Flammer C, Hong B, Minor D. 2019. Corporate governance and the rise of integrating corporate social responsibility criteria in executive compensation: Effectiveness and implications for firm outcomes. *Strategic Management Journal* **40**(7): 1097–1122.
- Gentry RJ, Harrison JS, Quigley TJ, Boivie S. 2021. A database of CEO turnover and dismissal in S&P 1500 firms, 2000–2018. *Strategic Management Journal* **42**(5): 968–991.
- Glebbeck AC, Bax EH. 2004. Is High Employee Turnover Really Harmful? An Empirical Test Using Company Records. *The Academy of Management Journal*. Academy of Management **47**(2): 277–286.
- Gond J-P, Piani V. 2013. Enabling Institutional Investors' Collective Action: The Role of the Principles for Responsible Investment Initiative. *Business & Society* **52**(1): 64–104.
- Gray A, Badkar M. 2018, December 20. Marlboro maker raises eyebrows with Juul price tag. *Financial Times*. Available at: <https://www.ft.com/content/f83022dc-0452-11e9-9d01-cd4d49afb3>.
- Greenwood R, Díaz AM, Li SX, Lorente JC. 2010. The Multiplicity of Institutional Logics and the Heterogeneity of Organizational Responses. *Organization Science* **21**(2): 521–539.
- Hart O, Zingales L. 2017. Companies Should Maximize Shareholder Welfare Not Market Value. *Journal of Law, Finance, and Accounting* **2**(2): 247–275.
- Hartzmark SM, Sussman AB. 2019. Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows. *The Journal of Finance* **74**(6): 2789–2837.
- Hausknecht JP, Trevor CO, Howard MJ. 2009. Unit-level voluntary turnover rates and customer service quality: Implications of group cohesiveness, newcomer concentration, and size. *Journal of Applied Psychology* **94**(4): 1068–1075.
- Hawn O, Ioannou I. 2016. Mind the gap: The interplay between external and internal actions in the case of corporate social responsibility. *Strategic Management Journal* **37**(13): 2569–2588.

- Heavey AL, Holwerda JA, Hausknecht JP. 2013. Causes and consequences of collective turnover: A meta-analytic review. *Journal of Applied Psychology* **98**(3): 412–453.
- Henisz WJ, McGlinch J. 2019. ESG, Material Credit Events, and Credit Risk. *Journal of Applied Corporate Finance* **31**(2): 105–117.
- hiQ Labs. 2021. *Are Millennials Less Loyal To Their Employers?* Available at: <https://static1.squarespace.com/static/5803b57737c581885cbd0667/t/5a638a5e71c10be0a05eb0ea/1516473156462/AMLL.hiQLabs.pdf>.
- Hölmstrom B. 1979. Moral Hazard and Observability. *The Bell Journal of Economics*. [RAND Corporation, Wiley] **10**(1): 74–91.
- Hong H, Kacperczyk M. 2009. The price of sin The effects of social norms on markets. *Journal of Financial Economics* **93**: 15–36.
- Ilmakunnas P, Maliranta M. 2005. Worker inflow, outflow, and churning. *Applied Economics*. Routledge **37**(10): 1115–1133.
- Jebb AT, Tay L, Diener E, Oishi S. 2018. Happiness, income satiation and turning points around the world. *Nature Human Behaviour*. Nature Publishing Group **2**(1): 33–38.
- Jouvenot V, Krueger P. 2020. Mandatory Corporate Carbon Disclosure: Evidence from a Natural Experiment. *SSRN Electronic Journal*.
- Kacmar KM, Andrews MC, Van Rooy DL, Steilberg RC, Cerrone S. 2006. Sure Everyone Can Be Replaced... but at What Cost? Turnover as a Predictor of Unit-Level Performance. *The Academy of Management Journal*. Academy of Management **49**(1): 133–144.
- Kahneman D, Deaton A. 2010. High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*. Proceedings of the National Academy of Sciences **107**(38): 16489–16493.
- Kaplan S. 2019, January 4. F.D.A. Accuses Juul and Altria of Backing Off Plan to Stop Youth Vaping. *New York Times*. Available at: <https://www.nytimes.com/2019/01/04/health/fda-juul-altria-youth-vaping.html>.
- Killingsworth MA. 2021. Experienced well-being rises with income, even above \$75,000 per year. *Proceedings of the National Academy of Sciences*. Proceedings of the National Academy of Sciences **118**(4): e2016976118.
- Kim BK, Jensen M. 2011. How Product Order Affects Market Identity: Repertoire Ordering in the U.S. Opera Market. *Administrative Science Quarterly*. [Sage Publications, Inc., Johnson Graduate School of Management, Cornell University] **56**(2): 238–256.
- Kim E-H, Lyon TP. 2011. Strategic environmental disclosure: Evidence from the DOE's voluntary greenhouse gas registry. *Journal of Environmental Economics and Management* **61**(3): 311–326.
- Kim E-H, Lyon TP. 2015. Greenwash vs. Brownwash: Exaggeration and Undue Modesty in Corporate Sustainability Disclosure. *Organization Science* **26**(3): 705–723.
- King A, Pucker K. 2021. The Dangerous Allure of Win-Win Strategies. *Stanford Social Innovation Review* **Winter 2021**.
- King AA, Lenox MJ. 2000. Industry Self-Regulation without Sanctions: The Chemical Industry's Responsible Care Program. *The Academy of Management Journal*. Academy of Management **43**(4): 698–716.
- King B, McDonnell M-H. 2012. Good Firms, Good Targets: The Relationship between Corporate Social Responsibility, Reputation, and Activist Targeting. *SSRN Electronic Journal*. Available at: <http://www.ssrn.com/abstract=2079227>.

- Loughran T, McDonald B. 2010. *When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks*. SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. Available at: <https://papers.ssrn.com/abstract=1331573>.
- Lyon TP, Maxwell JW. 2011. Greenwash: Corporate Environmental Disclosure under Threat of Audit. *Journal of Economics & Management Strategy* **20**(1): 3–41.
- Lyon TP, Montgomery AW. 2015. The Means and End of Greenwash. *Organization & Environment*. Sage Publications, Inc. **28**(2): 223–249.
- Maas K. 2018. Do Corporate Social Performance Targets in Executive Compensation Contribute to Corporate Social Performance? *Journal of Business Ethics* **148**(3): 573–585.
- Maas K, Rosendaal S. 2016. Sustainability Targets in Executive Remuneration: Targets, Time Frame, Country and Sector Specification. *Business Strategy and the Environment* **25**(6): 390–401.
- Mackintosh J. 2018, September 17. Is Tesla or Exxon More Sustainable? It Depends Whom You Ask. *The Wall St Journal*. Available at: <https://www.wsj.com/articles/is-tesla-or-exxon-more-sustainable-it-depends-whom-you-ask-1537199931>.
- Marquis C, Qian C. 2014. Corporate Social Responsibility Reporting in China: Symbol or Substance? *Organization Science* **25**(1): 127–148.
- Marquis C, Toffel MW, Zhou Y. 2016. Scrutiny, Norms, and Selective Disclosure: A Global Study of Greenwashing. *Organization Science* **27**(2): 483–504.
- McDonnell M-H, King B. 2013. Keeping up Appearances: Reputational Threat and Impression Management after Social Movement Boycotts. *Administrative Science Quarterly*. SAGE Publications Inc **58**(3): 387–419.
- McDonnell M-H, King BG, Soule SA. 2015. A Dynamic Process Model of Private Politics: Activist Targeting and Corporate Receptivity to Social Challenges. *American Sociological Review*. SAGE Publications Inc **80**(3): 654–678.
- Meyer JW, Rowan B. 1977. Institutionalized Organizations: Formal Structure as Myth and Ceremony. *American Journal of Sociology* **83**(2): 340–363.
- Mueller HM, Ouimet PP, Simintzi E. 2017. Within-Firm Pay Inequality. *The Review of Financial Studies* **30**(10): 3605–3635.
- Muller A, Kolk A. 2010. Extrinsic and Intrinsic Drivers of Corporate Social Performance: Evidence from Foreign and Domestic Firms in Mexico. *Journal of Management Studies* **47**(1): 1–26.
- Nason D. 2020, December 21. ‘Sustainable investing’ is surging, accounting for 33% of total U.S. assets under management. *CNBC*. Available at: <https://www.cnbc.com/2020/12/21/sustainable-investing-accounts-for-33percent-of-total-us-assets-under-management.html> [1 February 2021].
- Oikonomou I, Yin C, Zhao L. 2020. Investment horizon and corporate social performance: the virtuous circle of long-term institutional ownership and responsible firm conduct. *The European Journal of Finance*. Routledge **26**(1): 14–40.
- Peirce H. 2019, June 18. Scarlet Letters: Remarks before the American Enterprise Institute. Washington D.C. Available at: <https://www.sec.gov/news/speech/speech-peirce-061819>.
- Petrenko O, Aime F, Ridge J, Hill A. 2016. Corporate Social Responsibility or CEO Narcissism? CSR motivations and organizational performance. *Strategic Management Journal* **37**: 262–279.
- Principles for Responsible Investment. 2019. Taking Stock: Sustainable Finance Policy Engagement and Policy Influence. Available at:

- https://d8g8t13e9vf2o.cloudfront.net/Uploads/c/j/u/pripolicywhitepapertakingstockfinal_335442.pdf.
- Ramanna K. 2021, January 17. ESG accounting needs to cut through the greenwash. *Financial Times*. Available at: <https://www.ft.com/content/99b7a241-dfd4-4770-82b6-13de5a2804bf>.
- Ramus CA, Montiel I. 2005. When Are Corporate Environmental Policies a Form of Greenwashing? *Business & Society* **44**(4): 377–414.
- Sandvik J, Saouma R, Seegert N, Stanton C. 2021. Employee Responses to Compensation Changes: Evidence from a Sales Firm. *Management Science*. INFORMS **67**(12): 7687–7707.
- Semega J, Kollar M, Shrider EA, Creamer JF. 2021. *Income and Poverty in the United States: 2019*. Current Population Reports. Census Bureau: 88.
- Serafeim G. 2015. Integrated Reporting and Investor Clientele. *Journal of Applied Corporate Finance* **27**(2): 34–51.
- Shapiro C, Stiglitz JE. 1984. Equilibrium Unemployment as a Worker Discipline Device. *The American Economic Review*. American Economic Association **74**(3): 433–444.
- Shaw JD. 2014. Pay Dispersion. *Annual Review of Organizational Psychology and Organizational Behavior* **1**(1): 521–544.
- Shin S, Lee J, Bansal P. 2021. Moving Towards a Stakeholder Orientation: Evidence from the Analyses of Chief Executive Dismissal. *Academy of Management Proceedings* **2021**(1): 12225.
- Siebert WS, Zubanov N. 2009. Searching for the Optimal Level of Employee Turnover: A Study of a Large U.K. Retail Organization. *The Academy of Management Journal*. Academy of Management **52**(2): 294–313.
- Smeuninx N, De Clerck B, Aerts W. 2020. Measuring the Readability of Sustainability Reports: A Corpus-Based Analysis Through Standard Formulae and NLP. *International Journal of Business Communication*. SAGE Publications Inc **57**(1): 52–85.
- Sockin J. 2021. *Show Me the Amenity: Are Higher-Paying Firms Better All Around?* SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. Available at: <https://papers.ssrn.com/abstract=3957002>.
- Song K, Tan X, Qin T, Lu J, Liu T-Y. 2020. MPNet: Masked and Permuted Pre-training for Language Understanding. In *NeurIPS 2020*. Vancouver, Canada. Available at: <http://arxiv.org/abs/2004.09297>.
- Starks LT, Venkat P, Zhu Q. 2017. Corporate ESG Profiles and Investor Horizons. *SSRN Electronic Journal*. Available at: <https://www.ssrn.com/abstract=3049943>.
- Stevens JM, Kevin Steensma H, Harrison DA, Cochran PL. 2005. Symbolic or substantive document? The influence of ethics codes on financial executives' decisions. *Strategic Management Journal* **26**(2): 181–195.
- Sull D, Sull C, Zweig B. 2022. Toxic Culture Is Driving the Great Resignation. *MIT Sloan Management Review*. Available at: <https://sloanreview.mit.edu/article/toxic-culture-is-driving-the-great-resignation/>.
- Tang Y, Mack DZ, Chen G. 2018. The differential effects of CEO narcissism and hubris on corporate social responsibility. *Strategic Management Journal* **39**(5): 1370–1387.
- Terrachoice. 2010. *The Sins of Greenwashing Home and Family Edition*. Available at: <http://sinsofgreenwashing.com/index35c6.pdf>.

- Tyko K. 2021, February 25. Costco to increase its minimum wage to \$16 per hour, but CEO says the average pay for hourly workers is \$24. *USA TODAY*. Available at: <https://www.usatoday.com/story/money/shopping/2021/02/25/costco-minimum-wage-hourly-pay-increase-2021/6816327002/> [20 April 2022].
- Vracheva V, Judge WQ, Madden T. 2016. Enterprise strategy concept, measurement, and validation: Integrating stakeholder engagement into the firm's strategic architecture. *European Management Journal* **34**(4): 374–385.
- Werbel JD, Carter SM. 2002. The CEO's Influence on Corporate Foundation Giving. *Journal of Business Ethics*. Springer **40**(1): 47–60.
- Xia C, Meyer-Doyle P. 2021. Who Bleeds When Activists Attack? Activist Institutional Ownership And Employee Safety. *Academy of Management Proceedings* **2021**(1): 10538.
- Yang Y, UY MCS, Huang A. 2020. FinBERT: A Pretrained Language Model for Financial Communications. *arXiv:2006.08097 [cs]*. Available at: <http://arxiv.org/abs/2006.08097>.