Essays in Environmental Finance

Mihir Tirodkar

A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy in Finance

The University of Auckland, 2020

Abstract

This thesis comprises of three empirical studies that examine separate interactions between environmental variables and financial markets. Set within a U.S. backdrop, the studies integrate and analyse relationships between asset pricing, corporate finance, behavioural finance, and environmental economics. The first chapter examines whether low frequency temperature risk, a component of climate risk, is a priced risk factor in equity markets. Rising temperatures are associated with states of poor consumption and potential disasters; consumption-based asset pricing theories suggest investors prefer investments which pay off in these poor states. I estimate low frequency temperature shocks and employ them in asset pricing tests. Results provide no evidence of a low frequency temperature risk premium in U.S. equity markets. I discuss possible reasons as to why results may diverge from the asset pricing theory. The second chapter tests whether institutional investors are reluctant to own polluter 'sin' stocks. I hypothesise that sensitivity to social norms restricts institutional ownership of polluters. Using toxic emissions data from the Toxic Release Inventory, I find results that are consistent with the hypothesis. Furthermore, I find that institutions with long-term investment horizons and exposure to public scrutiny display a stronger reluctance to own polluters. I find no evidence of positive abnormal performance of polluter stocks, as hypothesised by the 'shunned-stock' theory. The final chapter examines security analyst earnings forecasts for polluter firms. Polluters are negatively exposed to increased regulations and consumer backlash; however, security analysts may misestimate associated costs. Tests show that analysts generate systematically pessimistic forecasts for polluter firm earnings on average; behavioural theories suggest that this pessimism is due to cognitive constraints. I also find evidence of persistent analyst bias for polluters, consistent with the conservatism bias theory. Results provide no evidence of polluter abnormal returns resulting from positive earnings surprises around earnings announcements.

Acknowledgements

A Doctoral thesis is rarely the work of a sole individual, and often reflects the advice, opinions, and support of others. This research has greatly benefited from the help of several people; it is very difficult to convey their full contributions within this limited space.

First and foremost, I would like to thank both of my supervisors, Henk Berkman and Paul Geertsema. I am very lucky to have been able to work with both Henk and Paul, and I consider them to be good friends.

I could not have written this thesis without the insights of Henk, which included research advice, constructive criticisms, and suggestions drawn from a depth of research experience; this thesis is all the better for it. Henk has been an outstanding teacher and I have enjoyed our discussions on a wide variety of topics over the past few years. I would like to thank Henk for guiding the beginning of my journey as a researcher and encouraging me to explore difficult questions.

Paul has been an invaluable source of expertise in all areas of asset pricing and modelling. I truly appreciate the many technical discussions we have had over the past years. Paul's teachings to me began prior to my time in graduate school and have been of great benefit in the development of my economic, logical, and quantitative skillsets.

The University of Auckland has been my second home for some time now and is an institution to which I attribute much of my personal growth. I would like to thank the University for providing me with a Doctoral Scholarship and generous funding for conference travel. I thank staff members from the department of Accounting and Finance for their helpful comments and general support in writing this thesis. I would also like to thank all the friends I made whilst completing my research, I am grateful for your company and support.

I thank all seminar participants at the 2017 ACFR doctoral symposium, 2019 and 2020 NZFC doctoral symposiums, and at the University of Auckland for their helpful comments and feedback.

I would like to express my gratitude to my family, to whom I am beyond grateful. To my parents for supporting me in my education since day one, providing me with countless opportunities, and for encouraging me to pursue my interests. To my grandfather, for all the proof-reads and showing genuine interest in my research. To Hobbes, for your positive companionship. Lastly, to Prajakta; you have been an unwavering source of support and motivation during the good times as well as the hard. This journey would not have been possible without your help and encouragement.

Mihir Tirodkar

April 2020

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Introduction

Central to environmental economics are the concepts of value and market failure. The natural environment is a source of value for any economy; however, unsustainable practices and misuse of natural resources can create economic risk and welfare loss. As a subset of environmental economics, studies in environmental finance examine interactions between the natural environment, financial securities, investors, and institutions that operate in financial markets. Finance is concerned with the optimisation of some economic objective function (Miller, 1999); these objectives often include firm value or personal wealth. Environmental variables have a clear interaction with such concepts of economic optimisation.

This thesis empirically examines three separate environmental topics within U.S. financial markets, where each is set out as an independent chapter. The first chapter studies the role of temperature, a subset of aggregate climate change, as a source of macroeconomic risk in equity markets, and empirically estimates the price of a temperature risk factor. The second chapter investigates institutional investors and their equity ownership of firms that pollute toxic emissions. The final chapter examines the relationship between security analysts, behavioural biases, and forecasted earnings of polluter firms.

Chapter 1 is the first study in this thesis on the U.S. equity market and low frequency (i.e. long-run) temperature risk. Climate change is driven by increasing average temperatures, which are expected to bring negative long-term economic consequences as certain temperature thresholds are crossed (Nordhaus, 2013; IPCC, 2014). Economic costs arise due to a declining natural environment as well as an increasing probability of natural disasters. The consumption-based and rare disaster asset pricing literature asserts that investors with consumption smoothing preferences should prefer investments that pay off in such disaster states. In contrast, securities that are expected to perform poorly in the worst climate scenarios require a risk premium for investment.

The primary contribution of this research is testing for a low frequency temperature risk premium, as hypothesised under consumption-based asset pricing, using empirical testing techniques consistent with the literature. The study highlights the economic rationale as to why temperature may be a priced risk factor, and empirically estimates a temperature risk premium. In recent studies, Bansal, Kiku, & Ochoa (2016) and Balvers, Du, & Zhao (2017) find evidence of a temperature risk premium in U.S. equity markets. Building on their work, I use a range of alternative empirical techniques to test for a similarly defined risk factor.

Using temperature data for the contiguous U.S. states obtained from the National Oceanic and Atmospheric Administration, I estimate low frequency temperature shocks through time series analysis. Within the data sample of 1988 to 2016, I test for a temperature risk premium using a range of benchmark asset pricing models and empirical techniques, which include pooled panel regressions, Fama-MacBeth regressions, and long-short portfolio tests. All techniques fail to provide any evidence of a low frequency temperature risk premium in equity markets. Furthermore, additional analysis on temperature exposures using an event study on the Paris Agreement and an alternative climate risk metric suggest that estimated temperature betas are poor measures of climate change risk. I provide some qualitative reasoning as to why estimated temperature risk premiums may be insignificant in tests.

Chapter 2 moves away from climate risk and asset pricing, and towards environmental pollution and corporate finance. Hong & Kacperczyk (2009) reveal that as a result of societal discrimination, institutional investors avoid investments in the traditional ‘sin’ stocks of alcohol, gambling and tobacco firms. This effect is less pronounced for institutions with aggressive investment strategies that are more likely play the role of market arbitrageur, and leads to reduced analyst coverage along with positive abnormal returns in accordance with the shunned-stock hypothesis (Angel & Rivoli, 1997). As with traditional sin stocks, polluting companies may be discriminated by society due to being perceived as irresponsible and costly to societal welfare; I test whether this is the case.

The contribution of this study is to assess whether institutional investors hold relatively fewer shares of polluter firms due to societal pressures. While anecdotal evidence may suggest there is an increasing environmental awareness amongst society and institutions, I empirically assess whether this is indeed the case. Furthermore, the study examines whether polluting firms exhibit other characteristics that are found for traditional sinners, such as reduced analyst coverage and abnormal stock returns.

Using data from the Toxic Release Inventory (TRI), I identify the greatest polluting firms in the sample of 1987 to 2014. Tests provide evidence of reduced aggregate institutional ownership of the top polluters relative to their other investments; I also find evidence of an increasing institutional reluctance in polluter equity investment over time. Results indicate that while publicly scrutinised institutions have the greatest relative reduction in polluter stock ownership, aggressive arbitrageur institutions are not as constrained. Institutions with shorter investment horizons disproportionately own more polluter equity. Furthermore, tests reveal that polluters receive reduced analyst coverage. In contrast to the shunned-stock hypothesis, I find no evidence of any abnormal returns generated from a long-short polluter trading strategy.

The third and final chapter of this thesis shifts the focus to security analysts, and the biases present in their forecasts of polluter firm earnings. Building on prior behavioural theories and empirical findings, such as Tversky & Kahneman (1974, 1979, 1981) and De Bondt & Thaler (1985, 1987, 1990), I hypothesise that analyst forecasted earnings for polluting firms are systematically pessimistic. Pessimistic biases may arise from analysts overweighting low-probability regulatory or demand-side costs, being overly influenced by memorable past events, and generating inaccurate conditional expectations of performance.

This research examines the effects of pollution on firm value from the perspective of security analysts. This study contributes to the literature by examining whether security analysts, who are professional estimators of value, exhibit biases when forecasting for polluting firms. Under the framework of rational expectations, there should be no systematic bias present in analyst estimates, however behavioural theories and empirical evidence indicate that this may not be the case. This research also examines whether individual analysts display persistent biases towards polluters, and whether aggregate analyst bias results in abnormal returns around earnings announcements.

Combining data on analyst forecasts, firm fundamentals, and the TRI, I find evidence of a systematic pessimistic bias within analyst forecasts of next period earnings. I find that this pessimism disappears when firm pollution is scaled by firm sales. When disaggregated by the types of toxic chemicals in the dataset, tests indicate that compared to persistent bio-accumulative chemicals and dioxins, standard TRI chemical releases are most associated with significant forecast pessimism despite being relatively less toxic. At the individual analyst level, I find evidence of persistence in analyst bias. Analysts that are relatively pessimistic (optimistic) towards polluting firms again exhibit pessimism (optimism) for subsequent forecasts. Despite evidence of an aggregate analyst pessimism in polluter earnings forecasts, I find no evidence of positive earnings surprises and abnormal returns around earnings announcements.

Chapter 1

Global Warming Asset Pricing: Estimating a Temperature Risk Premium

1.1. Introduction

Climate change is a significant economic concern due to the expected consequences of a deteriorating environment on variables such as consumption, output, employment and productivity. This research includes financial market considerations into the climate change discussion by examining the explanatory power of temperature as a risk factor in the cross-section of U.S. equity returns. Temperature change is a core driver of other climate change phenomena and is the main climate variable examined in this study. I focus on the impact of low frequency temperature shocks on equities.[1] The main contribution of this study is testing the hypothesis of a priced temperature risk factor in financial markets with an approach that is consistent with the asset pricing literature. I take a consumption-based asset pricing approach motivated by rare disaster models, and test for the existence of a risk premium for temperature sensitive stocks. Temperature rise is strongly linked to the physical risk of climate disaster events that ultimately reduce consumption (IPCC, 2014). Assets that perform poorly in states of unexpectedly high temperatures and low consumption are less attractive; I therefore test whether investors require a higher premium for stocks with negative temperature loadings.

I create a low frequency temperature shock variable by transforming raw U.S. temperature data and use it to proxy for shocks to expectations of long-term temperature trends. I estimate the sensitivities of U.S. industry equity returns to low frequency temperature shocks and find no evidence to suggest that average industry temperature betas have been decreasing with time over the sample. I also find no evidence of any interaction between temperature betas and the underlying low frequency temperatures themselves.

The main empirical tests estimate a cross-sectional temperature risk premium that is required by investors as compensation for exposure to low frequency temperature risk. I focus on the period post

[1] Low frequency temperatures are defined as long-run temperature averages. Low frequency temperature shocks are the differences between expected and realised low frequency temperatures. In contrast, high frequency temperatures are short-run temperatures which fluctuate around the long-run average.

1988; this date is chosen to coincide with the establishment of the Intergovernmental Panel on Climate Change (IPCC)[2] and is set as the cut-off date from which climate change awareness rapidly increased.[3] I find no evidence of a low frequency temperature risk priced into U.S. stock returns. Traditional portfolio tests are also used to examine the relationship between returns and temperature risk. I create portfolios based on temperature sensitivity to test the returns of a temperature sensitive long-short position, and to serve as a robustness check for main results. Neither equal nor value-weighted portfolios are found to generate significant average returns in accordance with the hypothesis, nor are portfolio alphas significant after benchmark risk factors are controlled for. Results provide no evidence of a low frequency temperature risk priced into U.S. stock returns.

Additional tests conducted on temperature betas include an event study on the 2015 Paris Agreement and temperature beta correlations with firm-specific climate disclosures. Industry temperature betas are not correlated with the outcomes of the Paris Agreement, nor are firm exposures to temperature shocks correlated with firm disclosure to total climate risk. Results suggest that historically estimated temperature betas are weak measures of return sensitivity to climate change.

I reconcile the lack of evidence of a priced temperature risk factor with three possible reasons. The first is the very long time horizons of climate-related disasters. If extreme climate disasters are expected to occur in the distant future, today's investors may not price this risk. Secondly, if temperature exposure is diversifiable then temperature risk is not systematic, and exposure will not generate a premium. Finally, long-run climate exposures may be reduced through firm and industry adaptabilities.

[2] The IPCC was established in 1988 by the World Meteorological Organization. The IPCC has issued a series of Assessment Reports that highlight the causes and consequences of climate change, and have increased climate awareness. The First Assessment Report was produced in 1990.

[3] James E. Hansen of NASA also gave testimony to Congress in 1988 that largely raised climate change awareness.

1.2. Literature review

The primary aim of this research is to estimate the market price of risk associated with exposure to temperature change factors. Specifically, the alternative hypothesis is the existence of a negative temperature risk premium that compensates investors for exposure to low frequency temperature risk.[4] Low frequency temperature shocks are assumed to affect realised returns by shocking both contemporaneous and future cash flows. I test whether temperature exposure helps explain cross-sectional variation in expected returns.

In this literature review, I examine the price of risk associated with temperature change through a consumption-based approach. I initially highlight some recent studies on climate risk in financial markets, and then consider the relationship between temperature, consumption and disasters. Thereafter, I review the literature for examples of interactions between temperatures and industry performance.

Bansal, Kiku, & Ochoa (2016) and Balvers, Du, & Zhao (2017) find evidence of equity market sensitivity to changes in average temperatures and a negative price of temperature risk. Using consumption-based pricing theory, I also hypothesise the existence of a negative temperature risk premium; however, I improve upon the limitations of the prior studies using an augmented modelling approach along with differences in assumptions and data.

Bansal et al. (2016) empirically test whether exposure to first order differences in low frequency temperatures is a source of risk, however, I argue that the first order differences may be predictable. Predictable changes in temperatures are unlikely to create systematic risk; expected changes should already be priced according to theories of efficient markets and rational expectations. Instead, return sensitivity to temperature *shocks* generates risk. Bansal et al. (2016) do not control for popular risk factors, use yearly frequency data, and estimate temperature betas without allowing for time-variation

[4] In the Arbitrage Pricing Theory example $E[R] = \beta * \gamma$ where expected excess returns R are a function of temperature sensitivity β, the price of temperature risk is γ. The temperature factor risk premium is equivalent to the price of temperature risk and the two terms are used interchangeably in this study.

in temperature sensitivities. Their regressions may therefore suffer from omitted variable bias from missing benchmark factors and measurement bias from the low frequency of returns and time-invariant estimates of temperature risk. Furthermore, their empirical models use data dating from 1934, which may generate incorrect temperature risk premium estimates given that expectations around climate change are a much more recent phenomenon.

Balvers et al. (2017) estimate whether temperature is a priced risk factor, primarily using a mimicking portfolio approach. Balvers et al. (2017) focus on average annual temperatures, which may be too high a frequency to proxy for climatic shifts. Their study controls for the Fama-French 3-factor model, however, the literature has developed stronger benchmark factor models. Factor exposures are not generated with rolling windows, leading to static estimates of sensitivity to temperature shocks. Balvers et al. (2017) also use data dating from 1953 when climate change was not a known concern and therefore unlikely to have been priced.

I develop on Bansal et al. (2016) and Balvers et al. (2017) with stronger testing methodologies, which include two-way clustered panel regressions and long-short portfolio alpha tests. I also use an alternative climate data set and an event study to cross-check estimated temperature betas.

Other notable work in this literature includes Donadelli, Jüppner, Paradiso, & Schlag (2019), who argue that inter-annual temperature *volatility* affects macroeconomic variables and is a priced risk factor. Griffin, Lont, & Lubberink (2019) find that extreme temperatures in the U.S. generate negative equity returns for firms with local operations. Engle, Giglio, Lee, Kelly, & Stroebel (2020) develop a strategy of hedging the impacts of short-term climate change news on realised returns. Choi, Gao, & Jiang (2020) take a behavioural approach and argue that investors adjust their beliefs about climate change following high frequency temperature shocks. While relevant, these studies do not specifically focus on estimating the risk premium of exposure to low frequency temperature shocks as based on classical economic and asset-pricing conventions, which is the primary contribution of this research.

1.2.1. Temperature and consumption

Campbell (2003) states that assets that are expected to perform relatively better in states of poor consumption should be in greater demand, with investors willing to pay higher prices or equivalently require lower average returns as compensation.[5] Campbell (2003) highlights that because investors attempt to smooth consumption through time, the equity premium is the covariance between stock excess returns and consumption growth multiplied by investor risk aversion. Empirically, models that estimate the equity premium only reconcile with observed premiums if unreasonable risk aversion parameters are introduced (Mehra & Prescott, 1985). Temperature effects have been found to shock consumption growth rates (Bansal & Ochoa, 2011); this negative relationship combined with rare disaster theory may help explain the equity premium puzzle.

Bansal & Ochoa (2011) find that rising temperatures impact world GDP growth negatively. Rising global temperature trends deteriorate aggregate growth and may lead to states of reduced consumption. Bansal & Ochoa (2011) also show that temperature betas contain information about differences in cross-country risk premiums. Global temperature rise has stronger negative impacts on economic growth for countries that are closer to the equator; market correlations with temperature shocks are found to vary between countries based on geography. This supports evidence presented by Dell, Jones, & Olken (2009) who find that temperature has a negative relationship with cross-sectional income at both the international and domestic levels. Additionally, Dell, Jones, & Olken (2012) reveal that growth is negatively affected when poorer countries have unusually hotter years, and is correlated with decreased investments and increased political instability.[6] Donadelli, Jüppner, Riedel, & Schlag (2017) find that temperature shocks in the U.S. lead to drops in consumption, output, investment and labour productivity growths. Consumption-smoothing investors value returns more

[5] More specifically, stochastic discount factor pricing illustrates that financial assets with low covariances with the future marginal utility of consumption must have high expected returns, as they tend to have poor pay-offs when investors need returns the most; when future marginal utility is high and future consumption is low.

[6] Using data on the Philippines, Crost, Duquennois, Felter, & Rees (2018) find that wetter wet seasons and drier dry seasons resulting from temperature rise may lead to increased civil conflict. Similarly, Ranson (2014) finds that under the IPCC's A1B climate scenario the U.S. will experience a large increase in crime. Climate change is therefore expected to impact socio-economic variables that are linked to consumption.

in poorer states and therefore should place a negative price on the covariance of future equity returns and unexpected temperature changes.

1.2.2. Rare disaster risk

Rare disasters are defined as improbable states in which consumption and output fall sharply (Barro & Jin, 2011). The disaster asset pricing literature primarily focuses on economic and wartime disasters such as the Great Depression, World Wars, and disease epidemics. I consider the potential consequences of climate change that will ultimately shock consumption, which include sea level rise, ocean acidification, permafrost thawing, and an increased intensity and prevalence of hurricanes, storm surges, wildfires, droughts and coastal flooding (IPCC, 2014; Jaffe & Kerr, 2015). Pricing of both actual and potential rare disaster events are of importance, as asset prices are set ex-ante on forward-looking expectations of future states (Berkman, Jacobsen, & Lee, 2011). The unmanageable events of Nordhaus (2013) are examples of disastrous climate change events that will negatively shock consumption and production in the long and very long-run. The IPCC (2014) states that higher temperatures are correlated with more frequent and intense natural disasters. The true disaster states resulting from climate change are expected to occur in the distant future as temperatures exceed safe bounds. Tail event disaster states shock consumption and should be priced into equity returns (Rietz, 1988); equities that are sensitive to short-run and long-run disaster states resulting from temperature shocks should require greater returns for exposure to climate disaster risk.

Disaster asset pricing provides a potential solution to the equity premium puzzle which is also grounded in consumption-based asset pricing theory. Rare disaster models incorporate the demands of risk-averse investors who are averse to extreme losses that may be incurred during disastrous events. Even if next period disasters do not actually occur ex-post, equity owners must be compensated with a premium for ex-ante exposure. Rietz (1988) finds that given reasonable estimates of risk aversion and investor impatience, an Arrow-Debreu approach that accounts for probabilities of market crashes can explain high equity premiums. Similarly, Barro (2006) calibrates a model to estimate an average probability of disasters that reduce GDP per capita, providing an explanation for high equity premiums and low interest rates in the U.S. during major wars. Copeland & Zhu (2007), however, argue that rare disaster risk is diversifiable to the extent that correlations between international disasters are less than perfect. Extending their argument to climate change, if climate disaster exposures are diversifiable between countries or industries then the effect on required premiums will be constrained.

The Rietz-Barro hypothesis is limited by the assumption of constant probability of disaster. Disaster probabilities can instead be modelled as a dynamic variable that adjusts based on investor expectations of future states (Wachter, 2013), while the corresponding equity exposures may vary in both the cross-section and time series (Gabaix, 2012). Variation in the cross-section and time series of climate disaster exposure is realistic due to varying correlations between industry performance and temperature. The probability of extreme climate disasters is likely conditional on ex-ante temperatures (IPCC, 2014), while the cash flow impact of climate shocks is unequally distributed among industries (Hitz & Smith, 2004; Schaeffer et al., 2012).

Models that allow for variable rare disaster risk also allow for volatile asset prices and can incorporate time-varying risk premiums. Berkman et al. (2011) empirically test this approach by using a time-varying index on perceptions of political disaster probabilities and find evidence in the cross-section for priced crisis risk; industries that are more exposed to rare crises are found to yield higher returns on average. My approach estimates sensitivities to long-run climate disasters based on historical correlations between equity returns and temperature shocks. Instead of explicitly estimating the probabilities and consumption costs of long-run climate disasters, I estimate cross-sectional exposures of diversified portfolios to low frequency temperature risk and the associated price of this risk. Implicit in my model is the central assumption that historical correlations between realised returns and temperature shocks are realistic estimates of forward-looking correlations. Bansal et al. (2016) find that the price of temperature risk has both a constant component and a time-varying component. Their results indicate that the temperature risk premium is dependent on temperatures in the current state, however, this premium is estimated using constant temperature betas. Developing on their findings, I additionally allow for time-varying temperature betas and include tests for a time-varying temperature risk premium. Incorporating dynamic estimates of temperature betas and risk premiums in models accounts for industry adaptation effects and time-varying prices of future climate disaster risk.

In summary, assets which are expected to perform well during states in which consumption is reduced through temperature-related rare disasters should require lower returns in the cross-section, ceteris paribus. This leads to a hypothesised negative temperature risk premium. If higher temperatures correlate with states of lower consumption, assets that have positive sensitivities to temperature shocks should generate less returns in equilibrium. Alternatively, investments that perform poorly in states when temperatures are high and consumption is low introduce greater volatility in future consumption, are unattractive to investors, and require a premium in returns.

1.2.3. Industry exposure to temperature

Temperature shocks are mostly related to physical risks which impact asset values by introducing exposure to climate factors that damage tangible assets, disrupt supply chains, and shift consumer preferences.[7] Not all firms are equally exposed to these channels; variation in temperature sensitivity among industries allows better estimation of a temperature risk premium. Through these channels, climate change winners and losers emerge within the economy. The immediate consequences of climate change are not always negative, and in the short-run certain industries and firms may be positively exposed to low frequency temperature shocks (Hitz & Smith, 2004). I highlight this possibility with the following industry examples.

[7] In the framework of Carney (2015), aggregate climate risk can be disaggregated into physical, transition and liability risks. Transition risk refers to the potential shift to a low-carbon economy and reflects the valuation impacts of changes in policy, technology, and firm reputation. Liability risk refers to the potential for firms to be held legally responsible and owe compensation to parties that suffer losses through climate change.

Schaeffer et al. (2012) summarise the consequences of climate change on various industries. They argue that agricultural considerations of long-term temperature rise include effects on precipitation, evapotranspiration and the reproduction rates of pests, all of which are expected to worsen cash flows to the industry. The increasing probability of tail events that include droughts, frosts and floods are also material considerations. However, higher CO_2 levels can positively improve the photosynthesis of crops (Schaeffer et al., 2012). As each crop category has an ideal temperature range in which productivity is maximised, gradual increases in temperature are likely to have parabolic relationships with crop output. Mendelsohn, Nordhaus, & Shaw (1994) argue that although the major grain groups are negatively exposed to temperature rise, they represent a small proportion of the American farm market. Warmer temperatures may improve alternative agricultural produce; the total impact of temperature on agriculture is dependent on the composition of the industry.

Climate change impacts on the energy and utility sector are also nonconstant. Within the energy and utilities sector, Schaeffer et al. (2012) further break down the sub-industries of hydropower, wind power, biofuels, solar energy, marine energy, oil, gas and coal into their resource endowments, energy supply and energy distributions supply chains. Even at this relatively broad level, the number of moving variables is large, and it is immediately obvious that climate factors have varying effects on operations. For example, solar energy generation is dependent on atmospheric water vapour content, cloud characteristics and atmospheric transmissivity. Climate change therefore has differing implications on solar energy generation at the country level; while positive impacts of temperature rise are reported in south-eastern Europe, negative impacts are noted in Canada as a result of decreasing solar radiation (Schaeffer et al., 2012). With rising temperatures, energy demand is also found to increase for cooling and decrease for heating. Total energy demand for temperature control is a parabolic function of average temperatures. Similar non-linear structures are also found in the demand for motors, engines and water. Similarly, Schaeffer et al. (2012) suggest that rising temperatures will increase the demand for vehicular air conditioning, leading to an increase in demand for fuel and efficient vehicles. Demand-side implications for energy are thus influenced by regional effects; rising temperatures in tropical climates would likely increase cooling energy demand while colder regions would see reductions in heating energy demand.

1.3. Data

1.3.1. Low frequency temperature shock

I create a proxy for low frequency temperature shocks that is used in the asset pricing models that follow. I focus on shocks to long-run temperatures, as high frequency temperature surprises are unlikely to reflect average temperatures shifts or climate change trends (IPCC, 2014). Raw temperature data consists of U.S. temperature observations in degrees Fahrenheit, obtained from the National Oceanic and Atmospheric Administration.[8] The data is of average monthly temperature observations for the contiguous 48 states. Using relatively less granular temperatures is justified through the demand-side motivation for the temperature risk premium itself; consumption-based asset pricing models assume that financial security prices are set based on macroeconomic consumption, which is in turn more likely to be influenced by average temperatures.[9] While temperature warming effects may vary across specific regions, climate change is itself representative of shifts in long-run temperature *averages*.

Low frequency temperature shock data requires a transformation of raw temperature data into a new variable that estimates the differences between observations and investor expectations of average temperatures. I create a low frequency temperature shock variable, *Temp*, from the residuals of an autoregressive temperature model.

Following Bansal et al. (2016), first a simple 5-year moving average is calculated for U.S. temperatures, representing low frequency temperatures. I also choose 5 years to maintain consistency

[8] Temperature data is sourced from URL https://www.ncdc.noaa.gov/cag/time-series/us. Monthly data for the observed average temperature in degrees Fahrenheit is retrieved for the period January 1895 to April 2017. No base period is subtracted from raw data and thus temperature data is not a meteorological temperature anomaly. Data has, however, been adjusted to remove artificial effects created by instrument changes, station relocation, urbanisation and observer practice changes, and may thus differ from official observations located elsewhere.

[9] Region-specific temperatures may influence regional firm operations, but have no impact on the consumption of investors, many of whom live elsewhere, and does not reconcile with the need for a risk premium.

with the sample size of rolling average regressions in later empirical tests. The moving average is not biased by monthly or quarterly seasonality and has a smoother trend than raw high frequency temperatures. I make the modelling assumption that low frequency temperatures are generated through the following autoregressive process.

$$MA_t = \alpha + MA_{t-1} + \beta^{\Delta MAt-1} * \Delta MA_{t-1} + \varepsilon_t \qquad (1.1)$$

This process assumes that moving average temperatures follow an augmented random walk, in which contemporaneous low frequency temperatures are equal to the sum of a deterministic drift, previous low frequency temperatures, feedback from prior changes, and an unexpected error term. I estimate and validate the parameters of model *(1.1)* as follows.

Using data from 1980 to 2016,[10] I take first order differences of the moving average to estimate observed changes in low frequency temperatures.

$$\Delta MA_t = MA_t - MA_{t-1} \qquad (1.2)$$

ΔMA is not a shock if it is not entirely random, as highlighted in the full temperature model. This is a realistic assumption, given that the very nature of temperature rise implies an expectation of positive shifts in low frequency temperatures. I use an autoregressive model to decompose *ΔMA* into expected and unexpected components. *ΔMA* autocorrelations with up to 12-month lagged *ΔMA* values are plotted using a partial autocorrelation function (PACF) in Figure 1.1. The estimated partial correlations between *ΔMA* and lagged *ΔMA* values are plotted with a 95% confidence band. *ΔMA* and 1-month lagged *ΔMA* show a significantly positive correlation with each other. The correlation between *ΔMA* and 2-month lagged *ΔMA* values is also positive and borderline significant at the 5% level, however for simplicity I ignore this marginally significant correlation and account for only a 1-month autoregressive term in my model.

[10] More specifically, I create the temperature anomaly variable from January 1980 to April 2017. Temperature betas are later calculated in a 60-month rolling window, and cross-sectional regressions run from 1988 to the end of 2016. This provides enough data points for risk premia to be estimated without losing any observations post 1988 in later regressions.

Figure 1.1: Partial autocorrelation function plot of ΔMA_t with 12 lags.

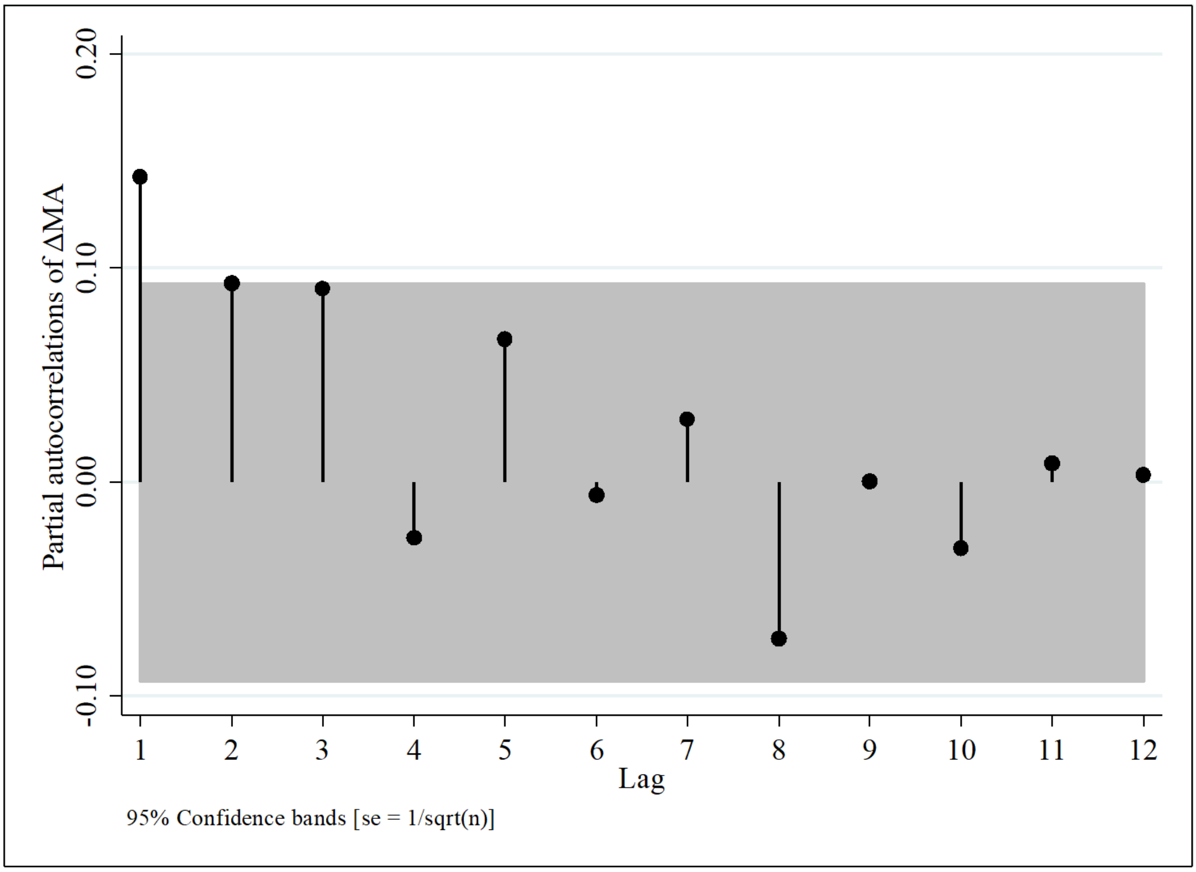

Following the output of the PACF function, I use an AR(1) process to model the time series of low frequency temperatures. I regress the changes in the moving average temperature against 1-month lagged values in the following autoregressive model.[11]

$$\Delta MA_t = \alpha + \beta^{\Delta MAt-1} * \Delta MA_{t-1} + \varepsilon_t \qquad (1.3)$$

This model accounts for an average temperature change with intercept α, and for feedback effects from prior temperature changes with ΔMA_{t-1}. The coefficient $\beta^{\Delta MAt-1}$ is estimated as 0.143 and is significant at the 1% level, while the constant is estimated as 0.005 and is significant at the 5% level; the regression R^2 is 0.02. The residual term reflects unexpected shifts in low frequency temperatures after accounting for the expected constant and feedback effects. I store and label the residuals ε_t as *Temp*, the low frequency temperature shock variable.

[11] Model *(1.3)* is equivalent to model *(1.1)* after subtracting MA_{t-1} from both sides of the equation.

$$Temp_t = \varepsilon_t \tag{1.4}$$

Conceptually, *Temp* is a proxy for shocks to 5-year temperature averages after adjusting for drift and feedback effects in temperatures through an autoregressive structure in temperature changes. I assume investor expectations are removed through the data transformation; *Temp* is thus assumed to reflect both unanticipated and exogenous low frequency temperature shocks.

Bansal et al. (2016) use first order differences in the moving average of temperature to proxy for low frequency temperature shocks; however, first order differences are not representative of shocks to investor expectations. First order differences may be predictable if the temperature trend can be estimated and therefore are not representative of temperature *risk*. I use a different approach by proxying investor expectations with my AR(1) model. *Temp* has a smaller average magnitude than the first order differences of Bansal et al. (2016), as expected components are subtracted from observed changes. I assume that this approach creates a more appropriate explanatory risk variable than first order differences in moving averages.

Table 1.1: Summary statistics for *Temp*, the variable used to proxy for low frequency temperature shocks.

		***Temp* summary statistics**					
Date range	N	Mean	Median	Min	Max	Std Dev	Skewness
March 1980 - April 2017	446	0.000	0.002	-0.158	0.181	0.048	0.231

Due to the moving average transformation, the first two months of data are lost in creating *Temp*. On average, the unexpected monthly temperature innovation is 0 degrees Fahrenheit as seen in Table 1.1; this is facilitated by the construction of *Temp* as regression residuals which is consistent with the assumption of rational expectations of temperatures. There is only weak positive skew in the *Temp* data, and the mean and median are similar. Both Pearson and Spearman correlations between *Temp* and shifts in the 5-year temperature MA used in Bansal et al. (2016) are 0.99 and are significant at the 1% level. The high correlation reveals that movements in *Temp* closely mirror the movements in the temperature variable used in Bansal et al. (2016), and that a high proportion of changes in temperature averages are contributed by the regression residuals.

I illustrate the *Temp* time series in Figure 1.2. First order differences in the temperature moving average are of low magnitudes, and thus the shocks to expectations of temperature change are also of low magnitudes. As drift and feedback effects have been removed from *Temp*, there is no clear deterministic tendency in the data.[12]

For robustness, I recreate *Temp* using alternative time series models and rerun main results. I find that the alternative measures of low frequency temperature shocks produce estimates of temperature risk premiums which are qualitatively the same in statistical and economic significance. Alternative measures of *Temp* and subsequent premium estimates using the Carhart 4-factor model are provided in the appendix.

[12] Using an augmented Dickey-Fuller test I reject the null hypothesis that *Temp* has a unit root at the 1% level. *Temp* is therefore accepted as a stationary time series. *Temp* does not have a significant coefficient when regressed on lagged values of itself either, providing no evidence of additional layers of autocorrelation.

Figure 1.2: The time series of *Temp*, the shocks to low frequency temperatures. The 5-year moving average of *Temp* is plotted in bold.

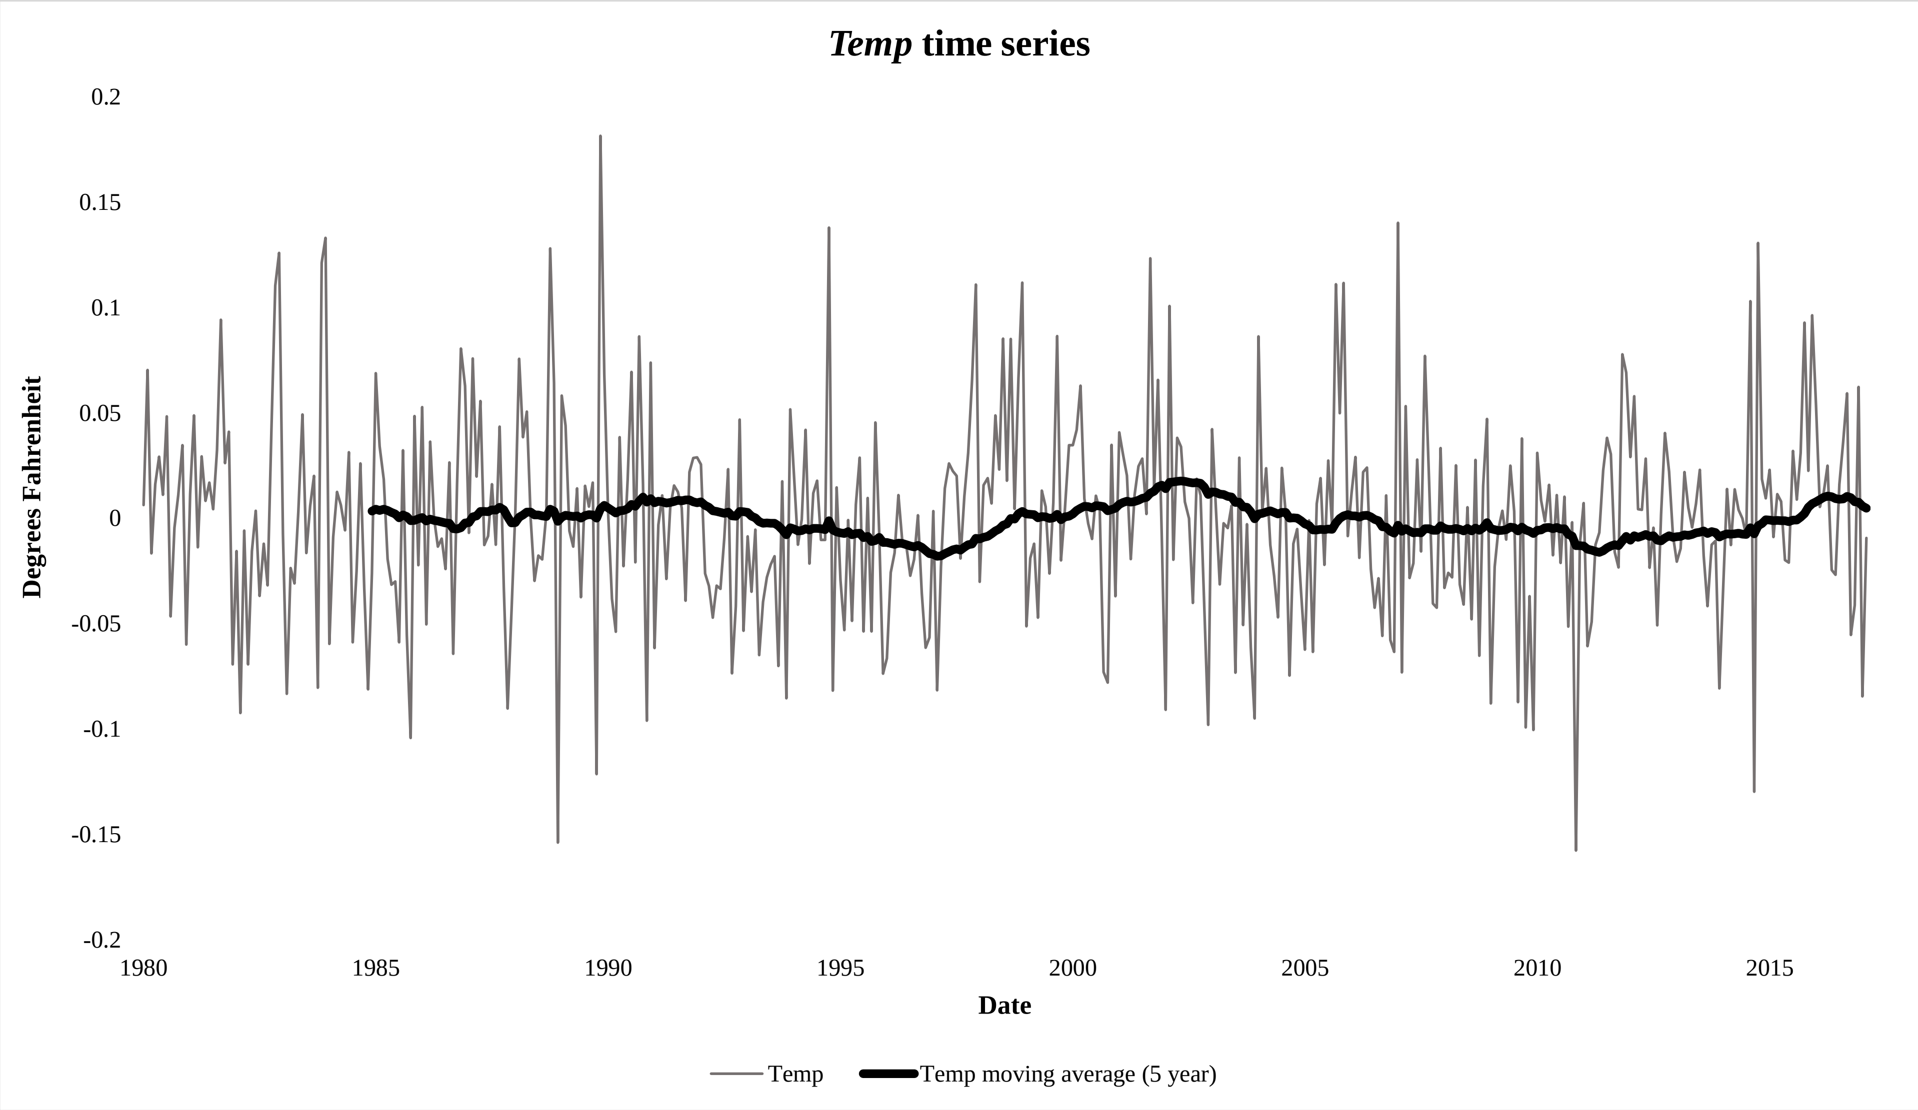

1.3.2. Returns and control risk factors

I obtain monthly and daily returns for the Fama-French 49 industry value-weighted portfolios (Fama & French, 1997); portfolios are used as test assets to reduce the effects of idiosyncratic risk on returns.[13] I choose industry portfolios as industries are intuitively expected to have natural variation in their sensitivity to temperatures. Furthermore, popular benchmark models such as the Fama-French factors do a poor job of explaining cross-sectional variation in industry portfolio returns (Berkman et al., 2011). For long-short portfolio tests, I use monthly realised returns on domestic U.S. equities.[14]

I obtain monthly data for control risk factor portfolios.[15] The five control factor models used in this study include the Capital Asset Pricing Model, the Fama-French 3-factor and 5-factor models (Fama & French, 1993, 2015), the Carhart 4-factor model (Carhart, 1997), and the Hou-Xue-Zhang q-factor model (Hou, Xue, & Zhang, 2014). All dependent and independent variable returns data are stored in percentage format.

In Table 1.2, I present estimated correlations between control risk factors and estimated *Temp* shocks from 1980 to 2016. Correlations are weak and indicate that the *Temp* variable is orthogonal to the risk factors.

---

[13] Industry portfolio returns were sourced from Kenneth R. French's data library, URL: *http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html*. Daily returns of industry portfolios are used in the event study.

[14] Monthly realised return data for domestic common equities are obtained from the CRSP database, with CRSP share codes of 10 or 11. I follow Shumway (1997) in correcting for delisting biases. If delisting returns in the panel data have a delisting stock code of 500, 520, between 551 and 573 inclusive, 574, 580 or 584, returns are set to -30%; while a missing delisting return with an available delisting code has returns set to -100%. Micro-caps are excluded from the equal-weighted portfolio. The micro-cap exclusion process involves dropping stocks with market capitalisations in the lowest decile in each month. Micro-caps are given immaterial weights in value-weighted portfolios and therefore are not excluded.

[15] *MKT*, *SMB*, *HML*, *RMW*, *CMA*, *MOM* and the risk-free rate data were sourced from Kenneth R. French’s data library. I obtain daily frequency data for the Carhart 4-factors to use in the event study. I thank Lu Zhang for data on the q-factor portfolios. Minor deviations exist between the q-factor and Fama-French market and size premiums due to slight differences in portfolio construction methods.

Table 1.2: Correlation matrix of explanatory risk factors. Correlations between temperature anomalies and control risk factors are weak, suggesting orthogonality and a low chance of collinearity issues in tests. Correlation coefficients significant at the 5% level are in bold.

	Temp	*MKT (FF)*	*SMB*	*HML*	*MOM*	*RMW*	*CMA*	*MKT (HXZ)*	*ME*	*I/A*	*ROE*
Temp	1										
MKT (FF)	-0.035	1									
SMB	0.013	**0.210**	1								
HML	-0.017	**-0.277**	**-0.154**	1							
MOM	0.007	**-0.128**	0.027	**-0.213**	1						
RMW	-0.052	**-0.316**	**-0.426**	**0.246**	**0.107**	1					
CMA	-0.014	**-0.391**	-0.061	**0.684**	0.007	**0.126**	1				
MKT (HXZ)	-0.033	**0.998**	**0.217**	**-0.276**	**-0.134**	**-0.323**	**-0.390**	1			
ME	0.013	**0.216**	**0.972**	**-0.105**	0.060	**-0.419**	-0.034	**0.222**	1		
I/A	-0.030	**-0.374**	**-0.169**	**0.684**	0.006	**0.229**	**0.908**	**-0.376**	**-0.138**	1	
ROE	-0.005	**-0.252**	**-0.366**	0.000	**0.510**	**0.700**	-0.008	**-0.265**	**-0.292**	**0.101**	1

1.3.3. Firm-specific climate disclosure

For secondary tests on climate sensitivity, I use a measure of firm-specific exposure to aggregate climate risk. Generated from sustainability disclosures on Ceres,[16] climate risk is measured from 10-K filings following the 2010 SEC ruling stating that material information on climate risk should be included in reports.[17] Prior to 2010, less than 24% of companies include any discussion of climate risk in their 10-K's, whereas in 2011 to 2015 56% of companies disclose some type of material climate risk (Berkman, Jona, & Soderstrom, 2019); hence the sample size is limited to this range.

ClimateScore is a proxy for firm-specific climate exposure from textual analysis of climate risk disclosures for Russell 3000 firms, and is based on the language and length of climate disclosure in firm's 10-K reports. Higher values of *ClimateScore* represent greater firm-specific climate risk. *ClimateScore* is an aggregate climate risk measure, and comprises of consolidated disclosures of firm-specific physical, transition and liability climate risks.

Summary statistics for *ClimateScore* are presented in Table 1.3. The positive skew in the data is largely due to variation between industry averages. In tests involving *ClimateScore*, I control for firm size and book-to-market ratios as at fiscal year-end sourced from CRSP and Compustat.

Table 1.3: Summary statistics for *ClimateScore*, values of self-disclosed climate exposure extracted using textual analysis of individual firm 10-K reports. Data is of yearly frequency.

			***ClimateScore* summary statistics**				
Date range	N	Mean	Median	Min	Max	Std Dev	Skewness
2011-2014	5,561	20.402	2.000	0.000	961	50.900	6.145

[16] I thank Henk Berkman for providing this climate disclosure data, which is also publicly available at https://www.ceres.org/issues/resources/tools/sec-sustainability-disclosure. See Berkman et al. (2019) for an in-depth discussion of *ClimateScore*.

[17] See SEC (2010) for full guidelines on required climate disclosure.

1.4. Average industry temperature exposure over time

Using the Fama-French 49 industry value-weighted portfolios as test assets, I initially estimate time-varying temperature betas and examine their trends over time. Temperature betas represent the sensitivities of industry realised returns to any deviations from expected low frequency temperatures. I estimate industry temperature betas with the following time series regression for each industry *i*.

$$R_{i,t} = \alpha_{i,t} + \beta_{i,t}^{temp} * Temp_t + \beta_{i,t}^{mkt} * MKT_t + \beta_{i,t}^{smb} * SMB_t + \beta_{i,t}^{hml} * HML_t + \beta_{i,t}^{mom} * MOM_t + \varepsilon_{i,t} \quad (1.5)$$

Industry excess realised returns $R_{i,t}$ are regressed on $Temp_t$ while controlling for the Carhart 4-factors in a 60-month rolling window regression.[18] Rolling windows provide time-varying estimates of exposure to all risk factors employed in the regression. Betas estimated with less than 30 observations in a window are set as missing.[19] I then average the individual industry temperature betas $\beta_{i,t}^{temp}$ for each month in the cross-section from 1988 onwards. This procedure generates an equal-weighted measure representative of average industry exposure to temperature. I plot average industry temperature betas in Figure 1.3.

[18] I choose 60-month rolling windows following the methodology of other research in this area, such as Petkova & Zhang (2005), Berkman et al. (2011), and Franzoni (2002). For illustrative purposes, I present the average temperature beta estimates for each industry based on these rolling window regressions in the appendix.

[19] The panel of observations with industry and month dimensions is strongly balanced from 1988 onwards.

Figure 1.3: An illustration of estimated average industry temperature betas β^{temp} over the sample period.

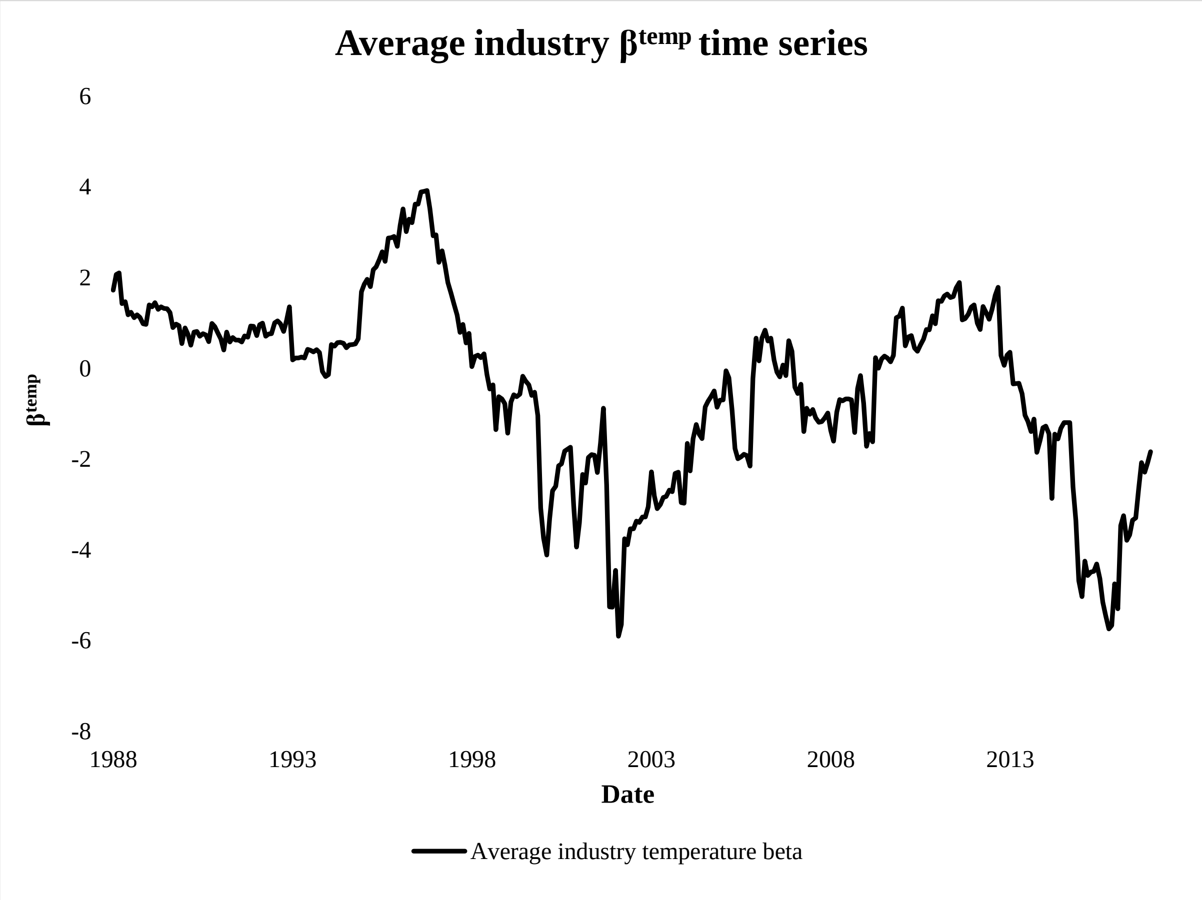

Average industry temperature betas appear to be decreasing over the sample. A negative temperature beta trend is intuitive. With rising average temperatures over time, the consequences of temperature shocks on returns may be larger; average industry returns may have a larger negative exposure to low frequency temperature shocks as average temperatures rise. I test for both trends in the average industry temperature beta and interactions with average temperatures in this section.

I first estimate trends in the time series of average industry temperature betas with the following regression.

$$\Delta\overline{\beta_t^{temp}} = \alpha + \gamma^{time} * t + \varepsilon_t \tag{1.6}$$

First order differences in the estimated average industry temperature betas $\Delta\overline{\beta_t^{temp}}$ are regressed against time t, from which a trend coefficient γ^{time}, intercept α and errors ε_t are estimated.[20] As the dependent variable $\Delta\overline{\beta_t^{temp}}$ is transformed by differencing the underlying $\overline{\beta_t^{temp}}$, the trend coefficient is interpreted as an exponential trend in the underlying average industry temperature beta, while the intercept is interpreted as a linear drift. Standard errors are adjusted for Newey-West 5-month lags.[21] I present the regression coefficients in Table 1.4.

Table 1.4: Estimated average temperature beta linear trend coefficient and constant values. P-values are Newey-West adjusted with 5-month lags. P-values in bold denote significance at the 10% level.

Average industry temperature beta trend

Variable	Coefficient	P-value
Trend	0.000	(0.882)
Intercept	-0.016	(0.671)
N	347	

I find no evidence of either an exponential or linear negative trend in the average industry temperature beta. Neither the trend nor the intercept coefficients are statistically significant, suggesting that the average industry temperature betas do not display a deterministic exponential or drift-like trend.

I also consider the correlation between average industry temperature betas and the underlying average temperature. I hypothesise a negative relationship between estimated temperature betas and average temperatures, meaning that the marginal costs of temperature shocks increase as base temperatures rise. I use the following regression to test this hypothesis.

$$\Delta\overline{\beta_t^{temp}} = \mu + \pi * \Delta MA_t + \eta_t \tag{1.7}$$

[20] The results of a PACF function on $\Delta\overline{\beta_t^{temp}}$ do not indicate the presence of first order autocorrelation; therefore I do not include an autoregressive difference term in this regression. The model effectively imposes a random walk process on $\Delta\overline{\beta_t^{temp}}$. An alternative approach is to regress $\Delta\overline{\beta_t^{temp}}$ on a trend and include $\overline{\beta_t^{temp}}$ as an explanatory variable; I find similar non-results using this specification.

[21] Following the literature I set the lag equal to $4(T/100)^a$ where T = 347 time periods and a = 4/25 using the quadratic spectral kernel. The output equals 4.88, which I round up to 5.

First order differences in average industry temperature betas $\Delta\overline{\beta_t}^{temp}$ are regressed against the first order differences in the 5-year moving average temperature, ΔMA_t.[22] The model provides estimates of the intercept μ, the slope parameter π and the error term η_t. If the hypothesised relationship between average industry temperature betas and average temperatures is true, a negative value for π is expected. I present the regression results in Table 1.5.

Table 1.5: Estimated relationship between first order differences in average industry temperature betas and a temperature 5-year moving average. P-values are Newey-West adjusted with 5-month lags. P-values in bold denote significance at the 10% level.

Average temperature beta relationship with low frequency temperatures		
Parameter	Coefficient	P-value
π	0.732	(0.363)
μ	-0.014	(0.567)
N	347	

Results provide no evidence of a negative relationship between low frequency temperatures and average industry temperature betas, revealing that the changes in estimated temperature exposures are not driven by changes in average temperatures.[23]

1.5. Main Results

In primary tests, I use estimated temperature betas to test for evidence of a priced temperature risk factor. Specifically, I examine whether stocks with greater temperature betas generate lower expected returns on average, and whether *Temp* adds explanatory power to asset pricing models. Tests include pooled panel regressions, Fama-MacBeth regressions, and a portfolio test.

1.5.1. Pooled panel regressions

I employ a pooled panel regression model with two-way clustered standard errors to estimate the temperature risk premium. The approach follows two stages. In the first stage I estimate industry

[22] Using augmented Dickey-Fuller tests on $\Delta\overline{\beta}^{temp}$ and ΔMA, I reject the null hypothesis of unit roots in the transformed variables at the 1% level. The results of the Durbin-Watson test do not indicate autocorrelation in the error terms of the specified regression.

[23] In unreported results, I find no evidence of a relationship between average industry temperature betas and temperatures even when interaction effects between ΔMA_t and MA_t are included. Similar non-results are obtained when estimating the relationship between first order differences in average temperature betas and lagged temperatures.

return sensitivities to *Temp*, and in the second stage I conduct a pooled panel regression of industry returns on estimated betas.

In the first stage, I conduct the following rolling window regression for each portfolio *i*.

$$R_{i,t} = \alpha_{i,t} + \beta_{i,t}^{temp} * Temp_t + \boldsymbol{\beta_{i,t}^{cont}} * \boldsymbol{cont_t} + \varepsilon_{i,t} \qquad (1.8)$$

Industry excess portfolio returns $R_{i,t}$ are regressed against the low frequency temperature shock $Temp_t$ and a vector of control risk factors $\boldsymbol{cont_t}$ in a 60-month rolling window time series regression. For each portfolio *i* during month *t*, $\alpha_{i,t}$ is the regression constant and $\varepsilon_{i,t}$ is the error term. $\beta_{i,t}^{temp}$ and $\boldsymbol{\beta_{i,t}^{cont}}$ are the estimated factor loadings of the temperature shock and control risk factors respectively, and are stored for the second stage. Betas estimated with less than 30 observations in a window are set as missing. Estimated sensitivities to temperature innovations vary dependent on the benchmark control risk factors used to estimate betas. The sample is constrained to 1988 - 2016; with this reduced sample I conduct the second stage pooled panel regression.[24]

$$R_{i,t} = \mu + \gamma^{temp} * \beta_{i,t-1}^{temp} + \boldsymbol{\gamma^{cont}} * \boldsymbol{\beta_{i,t-1}^{cont}} + \eta_{i,t} \qquad (1.9)$$

Industry excess returns $R_{i,t}$ are regressed against lagged beta estimates $\beta_{i,t-1}^{temp}$ and $\boldsymbol{\beta_{i,t-1}^{cont}}$ in a pooled panel regression with two-way clustered standard errors (Petersen, 2009). The model adjusts standard errors for clustering on both the industry and time dimensions. $\eta_{i,t}$ captures the error term of the pooled panel regression, while γ^{temp} is the estimated temperature risk premium and $\boldsymbol{\gamma^{cont}}$ is the estimated vector of premiums for control factor risk. I repeat the entire two-stage approach separately using each of the five control risk factor models. Table 1.6 presents the results of the pooled panel regressions.

[24] Results of the Harris-Tzavalis, Breitung, Im-Presaran-Shin and Fisher panel data unit root tests on estimates of β_t^{temp} generated with the Carhart 4-factors are significant at the 1% level; therefore the null hypothesis that temperature beta panels contain unit roots can be rejected. Similarly, the error terms of the pooled panel regression are also significant at the 1% level when used in the same panel data unit root tests, and are also be assumed to be stationary.

Table 1.6: Pooled panel regression results with two-way clustered standard errors. Industry portfolios are used as test assets. Pooled panel tests are run with CAPM, FF 3-factor, 5-factor, Carhart 4-factor and HXZ q-factor benchmark models. Monthly industry excess returns are regressed against temperature and control risk factor loadings to estimate the corresponding risk premiums. P-values are based on two-way clustered standard errors by industry and time and are shown in brackets below estimates. P-values in bold denote significance at the 10% level.

	Pooled panel regression results				
	CAPM	FF 3	Carhart	FF 5	HXZ
Constant	0.549	0.561	0.732	0.852	0.740
	(0.043)	**(0.022)**	**(0.001)**	**(0.001)**	**(0.005)**
Temp	-0.005	-0.002	0.000	0.001	0.000
	(0.317)	(0.686)	(0.954)	(0.912)	(0.952)
MKT	0.176	0.132	-0.041	-0.163	-0.038
	(0.583)	(0.680)	(0.894)	(0.623)	(0.902)
SMB		0.115	0.217	0.174	
		(0.466)	(0.149)	(0.222)	
HML		0.092	0.006	0.058	
		(0.647)	(0.975)	(0.769)	
MOM			0.106		
			(0.754)		
RMW				0.090	
				(0.563)	
CMA				-0.047	
				(0.722)	
ME					0.108
					(0.526)
I/A					0.155
					(0.265)
ROE					-0.015
					(0.927)
N	17,003	17,003	17,003	17,003	17,003

Results of the pooled panel regressions do not provide any evidence of a cross-sectional temperature risk premium. Exposure to low frequency temperature shocks are only estimated to generate a negative premium in two tests, and all estimates are insignificant at the 10% level. The regression constant captures a much larger premium than that generated from exposure to *Temp* or any of the control risk factors. This contrasts with arbitrage pricing theory, and supports prior evidence suggesting that common risk factors do not explain industry portfolio expected returns very well (Berkman et al., 2011).

I also test whether using a more recent sample generates a different estimate of the temperature risk premium. I reduce the sample to 2000 – 2016 and again conduct the pooled regressions, controlling for the Carhart 4-factors. For this reduced sample I estimate a temperature risk premium of 0.003 with an insignificant p-value of 0.363. Though climate awareness has grown in this period, estimates of the temperature risk premium are still insignificant. In another unreported robustness test, I transform temperature beta estimates into decile ranks and rescale them between 0 and 1.[25] This transformation serves to aid the interpretation of coefficient estimates and reduces second stage estimate sensitivity to measurement errors. I again do not find evidence of a temperature risk premium in this test. I also conduct the prior panel regressions on an alternate test sample of 25 portfolios that are created by independently double-sorting individual stocks on *Temp* and *SMB* betas estimated with the 4-factor model, however, estimated temperature risk premiums are still insignificant.

1.5.2. Time-varying temperature risk premium: Fama-MacBeth regressions

In this section, I use the Fama-MacBeth regression methodology to test for temperature risk premiums (Fama & MacBeth, 1973). The Fama-MacBeth approach allows for time-varying temperature risk premium estimates. I again use the Fama-French 49 industry value-weighted portfolios in a two-stage regression approach. In the first stage I calculate temperature betas for each of the *i* portfolios in separate rolling window regressions.

[25] For example, see Nagel (2005) or Berkman et al. (2011).

$$R_{i,t} = \alpha_{i,t} + \beta_{i,t}^{temp} * Temp_t + \boldsymbol{\beta_{i,t}^{cont}} * \boldsymbol{cont_t} + \varepsilon_{i,t} \quad (1.10)$$

I regress industry excess returns $R_{i,t}$ on $Temp_t$ and control risk factors $\boldsymbol{cont_t}$ in the first stage 60-month rolling window regression, in the same manner as the pooled panel regressions. Betas estimated with less than 30 observations in a window are set as missing. After the first stage procedure I reduce the sample period to observations from 1988 - 2016. The estimated betas from the first stage are stored; I then conduct the following second stage cross-sectional regressions for each month *t*.

$$R_{i,t} = \mu_t + \gamma_t^{temp} * \beta_{i,t-1}^{temp} + \boldsymbol{\gamma_t^{cont}} * \boldsymbol{\beta_{i,t-1}^{cont}} + \eta_{i,t} \quad (1.11)$$

In contrast to the single pooled panel regression conducted at the second stage for the pooled panel regression, in the Fama-MacBeth approach regressions are conducted for each time period. Excess returns are regressed against 1-month lagged β^{temp} and $\boldsymbol{\beta^{cont}}$ variables in each cross-section to obtain a monthly risk premium estimate for each risk factor, labelled as γ_t^{temp} and $\boldsymbol{\gamma_t^{cont}}$ respectively. μ_t captures the constant risk premium term in the model, while $\eta_{i,t}$ are the estimated error terms.

The resulting time series of estimated risk premium for the temperature factor and control factors, γ_t^{temp} and $\boldsymbol{\gamma_t^{cont}}$, are then averaged with Newey-West standard error corrections for 5-month lags. The entire two-step procedure is conducted separately using each of the five control risk models. Results are presented in Table 1.7.

Table 1.7: Second stage Fama-MacBeth regression results. Industry portfolios are used as test assets. Results are the average risk premiums of temperature and control risk factors over the time series. Fama-MacBeth tests are run with CAPM, FF 3-factor, 5-factor, Carhart 4-factor and HXZ q-factor benchmark models. P-values calculated from Newey-West adjusted standard errors with 5-month lags are shown in brackets below estimates. P-values in bold denote significance at the 10% level.

Second stage Fama-MacBeth regression results

	CAPM	FF 3	Carhart	FF 5	HXZ
Constant	0.361	0.346	0.333	0.422	0.335
	(0.180)	(0.132)	(0.145)	**(0.057)**	(0.165)
Temp	-0.003	0.002	0.003	0.001	0.003
	(0.572)	(0.704)	(0.545)	(0.855)	(0.535)
MKT	0.338	0.306	0.326	0.253	0.340
	(0.308)	(0.318)	(0.302)	(0.379)	(0.274)
SMB		0.039	0.089	-0.014	
		(0.788)	(0.540)	(0.926)	
HML		0.307	0.266	0.263	
		(0.101)	(0.152)	(0.145)	
MOM			0.317		
			(0.253)		
RMW				0.096	
				(0.447)	
CMA				0.042	
				(0.782)	
ME					0.097
					(0.521)
I/A					0.103
					(0.494)
ROE					0.125
					(0.437)
N	347	347	347	347	347

Fama-MacBeth results do not show a significant average temperature risk premium for exposure to low frequency temperature shocks. The control risk factor coefficients do not provide significant estimates of other risk premiums either, again illustrating the difficulties in explaining the expected return characteristics of industry portfolios (Berkman et al., 2011).

With increasing climate change awareness, it is possible that the temperature risk premium has become more negative over time. In Figure 1.4 I illustrate the temperature risk premium time series by plotting cross-sectional Fama-MacBeth estimates of γ_t^{temp}, after controlling for the Carhart 4-factors, from 1988 to 2016.

Figure 1.4: Monthly cross-sectional estimates of the temperature risk premium γ_t^{temp}, estimated after controlling for the Carhart 4-factors. A 5-year moving average of γ_t^{temp} is plotted with the darker thicker line.

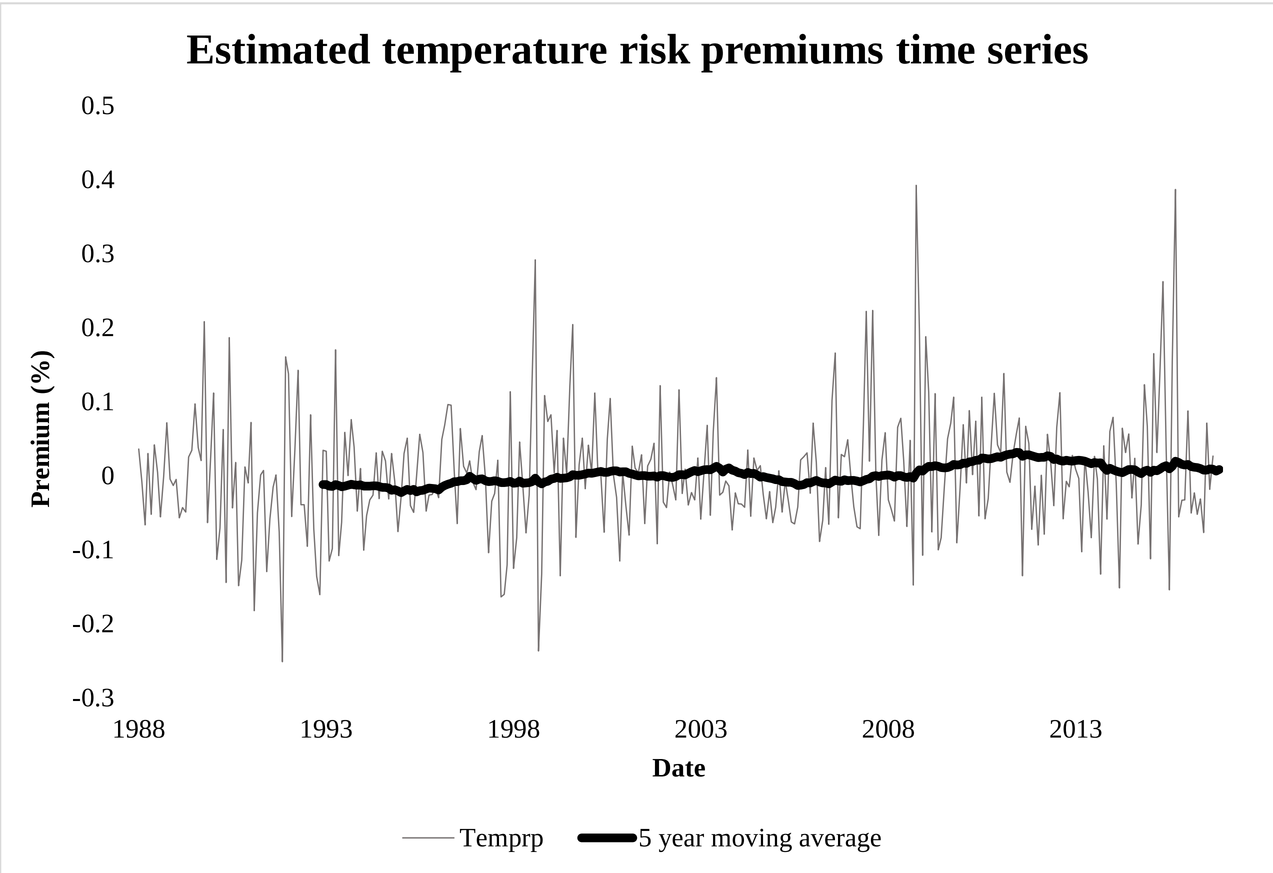

Contrary to a-priori expectations of increasing climate awareness, temperature risk premium estimates do not have a strongly observable trend through the time series. There are large volatilities in estimates of temperature risk premiums, while the smoothed 5-year moving average is relatively stationary. There is a slight upwards trend in the moving average in the late 2000's however this movement is in the opposite direction to the hypothesis and does not last. I test for a linear trend in estimated temperature risk premiums with the following regression.

$$\gamma_t^{temp} = \alpha + \beta^t * t + \varepsilon_t \qquad (1.12)$$

The estimated temperature risk premium γ_t^{temp} is regressed against time t in order to estimate a linear trend coefficient β^t. The constant and error terms are captured by α and ε_t respectively. I present the regression results in Table 1.8.

Table 1.8: Estimated trend coefficient of the temperature risk premium, calculated using the Carhart 4-factor model. Newey-West adjusted p-values with 5-month lags are displayed below estimates in brackets. P-values in bold denote significance at the 10% level.

Temperature risk premium trend

Intercept	Trend coefficient	N
-0.01333	0.00009	347
(0.122)	**(0.047)**	

I find no evidence of an economically significant trend in temperature risk premiums.[26] The monthly temperature risk premium is only increasing by just under 0.0011% each year, and is of the opposite sign compared to the alternative hypothesis; given an increasing climate change awareness, the estimated temperature risk premium should be *decreasing* instead of increasing. The intercept estimate is negative but is insignificant.

I perform another regression to test whether temperature risk premiums have shifted in recent decades. Estimates of the temperature risk premium γ_t^{temp} are regressed against two time dummies, D^{2000} and D^{2010}, which are activated during the 2000's and 2010's respectively. The constant temperature risk premium α and decadal dummy effects β^{2000} and β^{2010} are presented in Table 1.9.

$$\gamma_t^{temp} = \alpha + \beta^{2000} * D_t^{2000} + \beta^{2010} * D_t^{2010} + \varepsilon_t \qquad (1.13)$$

[26] The results of the Durbin-Watson test do not provide evidence of autocorrelation in the temperature premium linear trend regression errors.

Table 1.9: Decadal dummy coefficients of the temperature risk premium, calculated using the Carhart 4-factor model. Newey-West adjusted p-values with 5-month lags are displayed below estimates in brackets. P-values in bold denote significance at the 10% level.

Temperature risk premium decadal effects

Constant	2000's Dummy	2010's Dummy	N
-0.006	0.011	0.019	347
(0.397)	(0.244)	(0.107)	

I find no evidence of negative decadal effects in estimated temperature risk premiums. The dummy coefficients have a directional sign opposite to the hypothesis, and neither are significant.

Results are puzzling, providing no evidence of a negative linear trend for temperature risk premiums over the sample, nor any evidence of average decadal effects in the last two decades.

1.5.3. Portfolio tests

I implement a long-short portfolio strategy and create a tradeable temperature hedge portfolio to test for a priced *Temp* factor. For robustness, I create both equal-weighted and value-weighted portfolios based on return sensitivity to the *Temp* variable. I combine monthly excess returns data for individual equities with the five control risk factor models and *Temp* in the time series. I begin the portfolio creation process by generating temperature beta estimates for individual equities.

$$R_{i,t} = \alpha_{i,t} + \beta_{i,t}^{temp} * Temp_t + \boldsymbol{\beta_{i,t}^{cont}} * \boldsymbol{cont_t} + \varepsilon_{i,t} \qquad (1.14)$$

The beta estimation methodology follows the same process as the first stage regressions in temperature risk premium tests, but instead uses individual equities. I regress individual equity excess returns against *Temp* and control risk factors in the 60-month rolling window first stage regressions, from which coefficients formed with less than 30 prior periods are set as missing. After the beta estimation, the sample is reduced to observations from 1988 - 2016. In each month I generate portfolio breakpoints based on the deciles of lagged NYSE temperature betas. Stocks are then sorted into one of the ten temperature portfolios at the beginning of each month.

I implement a long-short portfolio strategy by subtracting the returns of the lowest temperature decile portfolio from the returns of the highest temperature decile portfolio, labelling the resulting long-short portfolio *Temphedge*. If a premium exists for temperature risk, then the monthly rebalanced *Temphedge* portfolio strategy should generate abnormally negative excess returns on average. I test for abnormal returns using the following regression.

$$R_t = \alpha + \boldsymbol{\beta}^{\boldsymbol{cont}} * \boldsymbol{cont_t} + \varepsilon_t \qquad (1.15)$$

The excess returns of the *Temphedge* portfolio R_t are used in a time series regression against the control risk factors $\boldsymbol{cont_t}$ with Newey-West 5-month lag adjustments. Sensitivity to control risk factors is captured in the vector $\boldsymbol{\beta}^{\boldsymbol{cont}}$, with an estimated intercept α and error terms ε_t. The intercept parameter estimate is interpreted as the *Temphedge* portfolio abnormal returns, which is expected to reflect a priced temperature risk factor. I repeat the portfolio formation process and test for abnormal returns using all five benchmarks models, using both equal and value-weights for robustness.[27] Tables 1.10 and 1.11 present the average returns, abnormal returns and factor sensitivities for the equal and value-weighted portfolios respectively.

[27] Because estimated temperature betas are dependent on the risk factors used in the first stage regressions, portfolio composition will also vary, effectively creating different portfolios for each benchmark model. Each column in the output tables is thus a different portfolio that is sorted on slightly different estimates of temperature sensitivities by separately controlling for one of the five benchmark factor models, from which the tabulated estimates for portfolio sensitivities and abnormal returns are again calculated using the same benchmark. I calculate portfolios weights based on beginning of month information.

Table 1.10: Equal-weighted *Temphedge* portfolio regression results. The average monthly portfolio returns are shown in the first row. The coefficients shown below are the estimated sensitivities of *Temphedge* to control risk factors. Portfolio alphas are shown in the third row. Significantly negative alphas would support the alternative hypothesis. Newey-West p-values generated with 5-month lags are reported in brackets below estimates. P-values in bold denote significance at the 10% level.

Equal-weighted *Temphedge* portfolio results

	CAPM	FF 3	Carhart	FF 5	HXZ
Avg. return	0.063	0.142	0.149	0.099	0.101
	(0.615)	(0.221)	(0.197)	(0.387)	(0.370)
Alpha	0.014	0.192	0.193	0.095	0.240
	(0.919)	**(0.095)**	(0.142)	(0.440)	**(0.053)**
MKT	0.076	0.007	0.002	0.022	-0.045
	(0.156)	(0.772)	(0.931)	(0.466)	(0.168)
SMB		0.011	0.021	-0.077	
		(0.824)	(0.679)	**(0.072)**	
HML		-0.229	-0.198	-0.204	
		(0.000)	**(0.007)**	**(0.000)**	
MOM			0.001		
			(0.985)		
RMW				0.008	
				(0.893)	
CMA				0.184	
				(0.040)	
ME					-0.048
					(0.403)
I/A					-0.163
					(0.064)
ROE					-0.102
					(0.057)
N	347	347	347	347	347
Adj. R^2	0.014	0.102	0.074	0.049	0.030

Table 1.11: Value-weighted *Temphedge* portfolio regression results. The average monthly portfolio returns are shown in the first row. The coefficients shown below are the estimated sensitivities of *Temphedge* to control risk factors. Portfolio alphas are shown in the third row. Significantly negative alphas would support the alternative hypothesis. Newey-West p-values generated with 5-month lags are reported in brackets below estimates. P-values in bold denote significance at the 10% level.

Value-weighted *Temphedge* portfolio results

	CAPM	FF 3	Carhart	FF 5	HXZ
Avg. return	-0.214	-0.079	-0.064	-0.107	-0.030
	(0.306)	(0.652)	(0.709)	(0.515)	(0.852)
Alpha	-0.358	-0.137	-0.209	-0.123	-0.140
	(0.089)	(0.408)	(0.248)	(0.527)	(0.555)
MKT	0.220	0.155	0.168	0.062	0.055
	(0.002)	**(0.002)**	**(0.002)**	(0.282)	(0.360)
SMB		0.007	0.029	-0.044	
		(0.925)	(0.700)	(0.455)	
HML		-0.177	-0.156	-0.172	
		(0.143)	(0.138)	(0.109)	
MOM			0.124		
			(0.145)		
RMW				-0.079	
				(0.385)	
CMA				0.192	
				(0.322)	
ME					-0.001
					(0.986)
I/A					-0.072
					(0.659)
ROE					0.199
					(0.095)
N	347	347	347	347	347
Adj. R^2	0.049	0.059	0.088	0.013	0.012

Average portfolio returns are not significantly negative for either the equal or value-weighted portfolios. Based on the expectation of a negative risk premium, the *Temphedge* portfolios should generate negative alphas on average after controlling for other risk factors. Results show very weak evidence of *Temphedge* abnormal returns. The value-weighted *Temphedge* portfolio strategy has only one significant negative alpha at the 10% level when benchmarked against the CAPM model. Surprisingly, two of the equal-weighted portfolios have positive alphas that are significant at the 10% level. Returns of the equal-weighted portfolios seem to be negatively driven by the *HML* factor, and only three of the value-weighted portfolios have significant loadings on the market premium factor. The *SMB* and *ME* factors do not explain much of the variation in *Temphedge* returns. This is interesting as it indicates that there is no correlation between the size factor and the returns associated with a low frequency temperature risk strategy; *Temphedge* portfolios do not behave like either small-cap or large-cap stocks.

Figure 1.5 illustrates the average excess returns for value-weighted decile temperature portfolios generated with all five benchmark models. The average excess return structure of the decile portfolios does not follow a monotonically negative trend as expected under the hypothesis of a negative temperature risk premium. Though a negative relationship is evident on average, it is volatile. Results overall do not provide any evidence of a priced *Temp* risk factor.

Figure 1.5: Average excess returns of value-weighted temperature beta decile portfolios, separately formed using all five benchmark models. The thick dark line is the average trend. Higher portfolio numbers have greater temperature sensitivity, while lower portfolio numbers have lower temperature sensitivity.

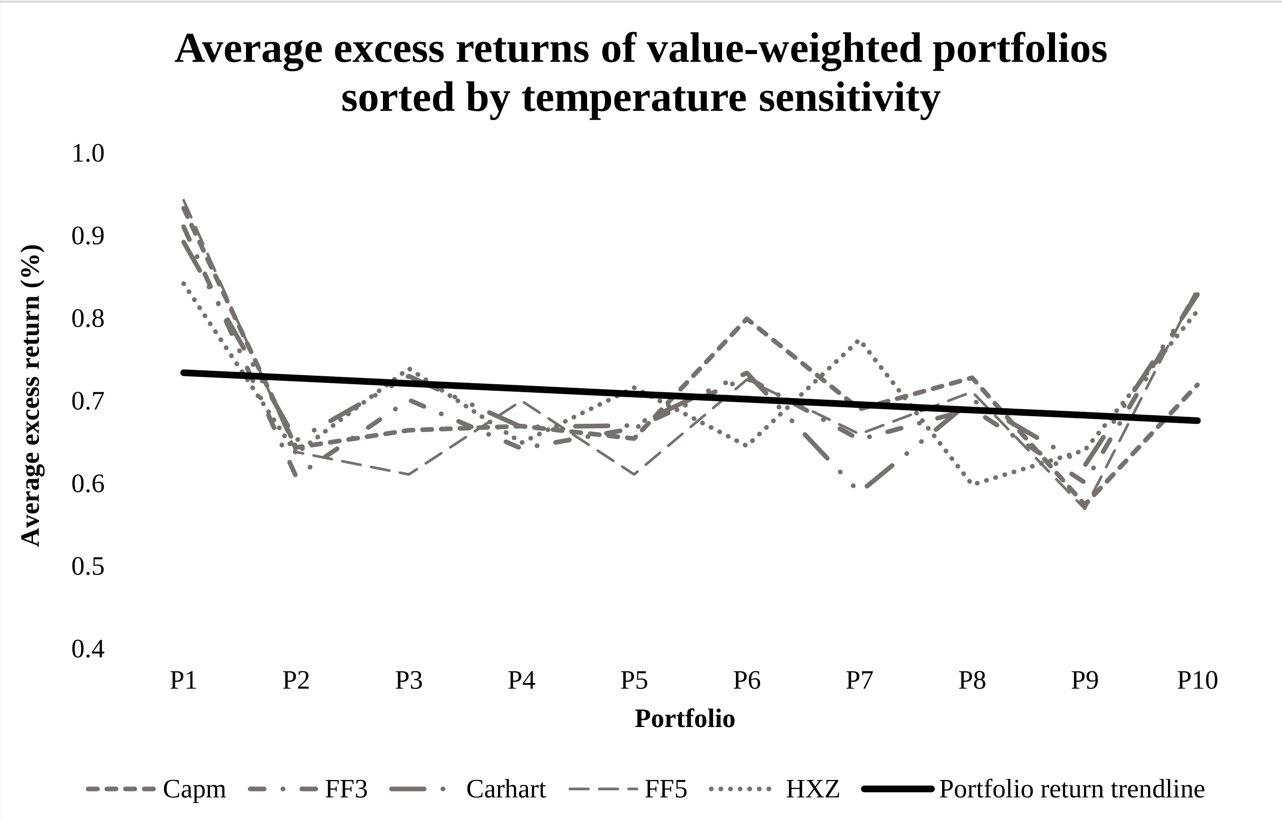

1.6. Additional temperature beta tests

In additional tests, I examine the validity of estimated temperature betas through an event study and by comparing estimates with another firm-specific climate risk variable.

1.6.1. Event study

I use event study methodology to examine the relationship between industry temperature beta estimates and the impacts of environmental regulation. The event study is conducted using the United Nations Framework Convention on Climate Change Paris Agreement, adopted by the U.S. on the 12th of December 2015. Alternative events are available, such as the Kyoto Protocol, the Copenhagen Accord or various physical climate phenomena; however, the Paris Agreement is chosen for its recent occurrence and unexpected outcomes.[28]

The event study serves as an external validity test of the reliability of estimated industry temperature betas. As the Paris Agreement aimed to reduce long-run temperature rise, I test whether the impacts of the event on industry returns are correlated with industry exposure to low frequency temperature risk. The event is expected to have a greater impact on the realised returns of temperature sensitive industries by negatively shocking long-run temperature forecasts, which in turn should shock expected future cash flows.[29]

Industry exposures to low frequency temperature risk are used to form expectations of event abnormal returns. I group industries based on estimated low frequency temperature betas, which I calculate using the following 60-month rolling window regression for each industry.

$$R_{i,t} = \alpha_{i,t} + \beta_{i,t}^{temp} * Temp_t + \beta_{i,t}^{mkt} * MKT_t + \beta_{i,t}^{smb} * SMB_t + \beta_{i,t}^{hml} * HML_t + \beta_{i,t}^{mom} * MOM_t + \varepsilon_{i,t} \quad (1.16)$$

[28] See "Deal done" (2015) for a discussion of the unexpected outcomes of the Paris Agreement.

[29] Temperature betas are measures of the sensitivity of the unexpected component of realised returns to temperature shocks. If the Paris Agreement reduced expectations of future average temperatures, future cash flows and current realised returns must also be impacted in the direction of the temperature beta, assuming ex-ante temperature betas are stable predictors.

Industry monthly excess portfolio returns $R_{i,t}$ are regressed against *Temp* and the Carhart 4-factors. I store the resulting industry $\beta_{i,t}^{temp}$ estimates that fall within an approximate 3-year window prior to the event, ranging from the 1st of January 2013 to the 30th of November 2015. I then average the monthly $\beta_{i,t}^{temp}$ estimates over the 3 years for each industry to generate an ex-ante average temperature exposure. Industries with positive betas are expected to benefit from expectations of increasing temperature, and suffer when temperatures are expected to fall. As the Paris Agreement aimed to reduce long-run temperature rise, I group industries with positive average temperature betas as 'expected losers', and negative average temperature betas as 'expected winners'.

Daily value-weighted industry returns are used to estimate the event impact on U.S. industries. I follow the multivariate regression model (MVRM) as described by Binder (1998) to estimate portfolio abnormal returns.[30] The measurement period is set as the year prior to the event, spanning from the 30th of November 2014, to the 31st of December 2015. A single event dummy is set to equal 1 on the 11th and 14th of December 2015.[31] The following Newey-West regression with is run with 5-day lags for each industry to estimate coefficients for the dummy variable.[32]

$$R_{i,t} = \alpha_{i,t} + \beta_i^{Paris} * Paris_t + \beta_i^{mkt} * MKT_t + \beta_i^{smb} * SMB_t + \beta_i^{hml} * HML_t + \beta_i^{mom} * MOM_t + \eta_{i,t} \qquad (1.17)$$

R_t is the daily excess realised return for a particular industry, α is the estimated constant and ε_t is the estimated error term in the regression. I control for the daily returns of the Carhart 4-factor control

[30] The MVRM methodology involves running a separate regression for each of the portfolios being examined in the event study, and using a dummy variable to capture abnormal returns around the event. Its primary advantages include allowing for abnormal return estimates to vary in sign across portfolios, and avoiding contemporaneous correlation and heteroskedasticity between the abnormal return estimates.

[31] The 12th and 13th of December 2015 fell on a Saturday and Sunday and have no returns data. Activating the dummy variable over the 11th and the 14th allows for some leaked information or delayed reaction impacts to be captured and is a relatively conservative approach.

[32] Following the literature I set the lag equal to $4(T/100)^a$ where T = 274 time periods and a = 4/25 using the quadratic spectral kernel. The output equals 4.70, which I round up to 5.

model. $Paris_t$ is the event dummy, while β^{Paris} is the estimated average daily abnormal return for when the dummy variable is activated. This regression is run for each of the 49 industry portfolios.[33]

I test the correlation between the Paris dummy coefficients and temperature betas across all industries. The Pearson correlation coefficient is 0.13 with a p-value of 0.37, while the Spearman correlation is 0.03 with a p-value of 0.85; both estimates are insignificant.

I present the summarised outcomes of the event study based on ex-ante industry groupings in Table 1.12. Out of all 49 industries, only 24 had event abnormal returns consistent with expectations based on temperature betas, of which only 13 had significant estimates.

Table 1.12: A summary of the expected winners and losers from the Paris Agreement event study, grouped by ex-ante temperature betas. Presented are a count of the number of industries in both expected winner and loser groups, along with a count of the number of industries that have expected event abnormal returns met and a count of the number of expectations met with statistically significant event abnormal return estimates at the 10% level.

Paris Agreement event study summary

Expected winners		Expected losers		Total
Count	26	Count	23	49
Consistent outcomes	14	Consistent outcomes	10	24
Significant and consistent estimate	6	Significant and consistent estimate	7	13

If industry groupings based on temperature beta are unrelated to the event outcome, the cumulative probability for 24 or more successfully predicted outcomes out of 49 is 0.612.[34] This suggests that the observed correct 24 industry predictions are likely due to chance acting alone. Results therefore do not provide evidence of a relationship between industry exposure to temperature

[33] Abnormal return estimates along with average temperature beta estimates are presented for each of the 49 industries in the appendix.

[34] The cumulative probability of observed results under the null hypothesis is calculated by randomly distributing the 49 industries into one of the two groups with equal probability. The cumulative binomial probability of 24 or more successful predictions out of 49 is 0.612. This probability is insignificant in a one-tailed hypothesis test.

and the abnormal returns generated around the Paris Agreement. These findings suggest that ex-ante temperature betas may be poor forecasts of future temperature sensitivity.

Alternatively, non-significant results from using industry temperature exposure as the basis of predictions may indicate that the Paris Agreement had impacts on industry returns that are not directly related to temperature sensitivity. This may be because climate-related industry shocks are channelled through regulatory uncertainties (Wellington & Sauer, 2005) as well as exposure to other climate phenomena. Temperature risk is only a subset of aggregate climate risk and may not provide a full picture, especially if weakly correlated with other types of climate risk. Aside from these explanations, it is also possible that the results of the Paris Agreement were already priced through information leakage or were not altogether credible. Overall results do not show that historical temperature betas are linked to industry abnormal returns generated around the 2015 Paris Agreement.

1.6.2. Climate disclosure tests

I examine whether estimated temperature betas are associated with *ClimateScore*, a firm-specific proxy for aggregate climate risk of which physical risk is a component. Tests estimate whether firms that disclose climate risk in their 10-K filings also have temperature sensitive equity returns, and subsequently measure the effectiveness of temperature betas as a proxy for climate risk. I use a subsample of CRSP equity data based on firms that have climate disclosures.[35] I calculate firm-specific equity temperature sensitivities with 60-month rolling window regressions and control for the Carhart 4-factors. I keep estimated temperature betas for the years 2011 – 2014.

$$R_{i,t} = \alpha_{i,t} + \beta_{i,t}^{temp} * Temp_t + \beta_{i,t}^{mkt} * MKT_t + \beta_{i,t}^{smb} * SMB_t + \beta_{i,t}^{hml} * HML_t + \beta_{i,t}^{mom} * MOM_t + \varepsilon_{i,t} \quad (1.18)$$

[35] Following Berkman et al. (2019), I exclude observations of firms in the financial services industry, firms with negative book values of equity, and firms with missing financial data.

I then average the monthly temperature betas β_t^{temp} for each of the 4 years for each firm. I estimate the Pearson and Spearman correlations between yearly β_y^{temp} and *ClimateScore* as 0.00 and -0.02, with insignificant p-values of 0.96 and 0.18 respectively. I also perform the following yearly frequency, industry fixed effects panel regression with yearly clustered standard errors.

$$ClimateScore_{i,y} = \pi^{temp} * \beta_{i,y}^{temp} + \pi^{size} * Size_{i,y} + \pi^{B/M} * B/M_{i,y} + \eta_{i,y} \quad (1.19)$$

ClimateScore is regressed against yearly average temperature beta estimates in the panel. The regression estimates the relationship between within industry variation in equity temperature betas and self-disclosed aggregate climate risk π^{temp}. I hypothesise a negative relationship between β_y^{temp} and *ClimateScore*. The independent variable *Size* controls for the natural log of firm market capitalisation, while *B/M* controls for firm book-to-market ratios. I include industry fixed effects in the regression using the Fama-French industry 49 classifications. The error terms are captured by $\eta_{i,t}$. I present the results in Table 1.13.

Table 1.13: Estimated intercepts and slope coefficients for the *ClimateScore* regression. I include industry fixed effect dummies in the regression. P-values are based on heteroscedasticity-consistent standard errors with clustering by year. P-values in bold denote significance at the 10% level. There are 4,529 observations in the sample. The regression generates an adjusted R^2 value of 0.487 and a within R^2 value of 0.028.

Firm climate disclosure and temperature beta

Variable	Coefficient	P-value
π^{temp}	0.023	(0.240)
π^{size}	4.219	**(0.000)**
$\pi^{B/M}$	10.874	**(0.020)**
N	4,529	

I find no evidence of a relationship between firm-specific climate disclosures and estimated temperature sensitivities.[36] Results fail to provide evidence that temperature betas are a significant contributing factor in explaining total firm-specific climate sensitivity.

1.7. Discussion

I find no evidence of a priced cross-sectional temperature risk factor. This is inconsistent with the hypothesis developed using consumption and disaster pricing theory.

Econometric explanations for the non-results include a potential lack of explanatory power in *Temp*. I find that alternative time series models of *Temp*, listed in the appendix, also generate insignificant temperature risk premium estimates. Low frequency temperatures may have low correlations with disaster states and other important climate variables such as humidity, wind speed and evaporation, which may in turn cause bias in the predicted physical impacts of temperature change on business activity (Zhang, Zhang, & Chen, 2017). Variability in temperature and other climate variables is also linked to complex climate patterns, such as the El Niño Southern Oscillation and solar cycles. These are not explicitly modelled into *Temp*, which may lead to measurement errors. Various estimation issues also arise when using realised returns as a proxy for expected returns.[37]

Temperature averages may hide shocks in geospatial cross-sectional temperature volatility; while *Temp* is a measure of shocks to low frequency temperatures around the U.S., it does not capture shocks to cross-sectional volatility, which may be more correlated with investor consumption. For robustness, I test whether cross-sectional temperature volatility better explains equity risk premia.[38] Results of the pooled panel and Fama-MacBeth regressions using shocks in temperature volatility

[36] In unreported results, I also find no evidence of a relationship between industry temperature betas and climate disclosure at the industry average level, nor is there evidence of a relationship between firm temperature betas and climate disclosure when both variables are standardized by year and industry.

[37] For a greater discussion on the problems of using realised returns as a proxy for expected returns, see Elton (1999), Brav, Lehavy, & Michaely (2005), or Berkman (2013). One alternative is the use of analyst forecasts as a proxy for expected returns, however this measure comes with its own set of biases.

[38] Donadelli et al. (2019) use time series average temperature volatility, as opposed to the cross-sectional temperature volatility robustness check used in this study.

instead of *Temp* are presented in the appendix, which provide no evidence of a premium for cross-sectional temperature volatility either.

Alternatively, I reconcile the lack of results with the following three qualitative explanations. The first is the very long time horizon in which rare climate disasters are expected to take place. The greatest unmanageable climate disasters described by Nordhaus (2013) are more likely to occur in distant future states, whereas the immediate consequences of low frequency temperature shocks on consumption may be negligible. Dasgupta (2008) argues that the consequences of climate change are on both intragenerational and intergenerational welfare. Dasgupta (2008) further points out that the considerations behind saving for our children or grandchildren, who are the real losers of climate change outcomes, are not the same as saving for personal future consumption. This reasoning could influence capital markets if investors do not value the consumption risk of future generations. Similarly, the temperature risk premium curve may not be flat. Even if temperature and consumption are negatively correlated in the long-term, weak short-term correlations would not justify a monthly premium due to a lack of immediate risk. Investors may not price long-term risk within short-term horizons, similar to a term structure effect. This logic assumes that markets are efficient, however, the market may simply behave irrationally around climate change factors (Liesen, 2015), resulting in an irrational or inefficient pricing of temperature risk.

The second possibility is the diversification potential at the investor and country level. Extending the argument of Copeland & Zhu (2007), if investors can diversify away their exposure to temperature shocks then there should not be a priced temperature risk factor in equilibrium. Diversification can also occur at the country level. If cross-country correlation to climate disasters is less than perfect, global investors can reduce portfolio exposure to temperature through diversification. In an efficient market only systematic risks are compensated with returns; if temperature risk is currently idiosyncratic then highly exposed stocks may not necessarily be compensated.[39]

Lastly, firms have dynamic capabilities. Businesses that can adapt to changing environmental factors benefit from built-in real option values (Trigeorgis, 1993) that restrict the negative outcomes driven by overall temperature rise. For example, Mendelsohn et al. (1994) illustrate how, in rising temperatures, farmers can reallocate their production efforts to differing outputs to avoid large losses. Firm adaptabilities constrain the sensitivities of their equity returns to climate shocks in the long-run. Even if investors currently price long-term risk, their investments may not be exposed to long-term temperature shocks. If the impacts on the manageable activities of Nordhaus (2013) constitute a large proportion of total long-run climate change costs, the total cash flow impact of temperature shocks is reduced. Aggregate firm adaptability mitigates the long-term value impacts of climate change, and thus reduces future costs relative to scenarios with limited firm flexibility. The market may also expect technology to improve at a rate which prevents the full scope of climate-related costs from affecting firms in the future. Like firm adaptabilities, potential technological advances constrain forward-looking temperature betas of firms. As a result of these factors, historical temperature betas may be poor estimates of forward-looking temperature sensitivities.

1.8. Conclusion

Overall, I find no evidence of the existence of a temperature risk factor in U.S. equity markets. Low frequency temperature risk is a subset of total climate risk, which has complex impacts on economic variables. Results do not suggest that exposure to low frequency temperature risk is correlated with higher excess returns in U.S. equity markets.

I transform temperature data to create a proxy for low frequency temperature shocks and calculate the temperature exposures for industry portfolios. Using pooled panel and Fama-MacBeth

[39] Anecdotally this explanation is difficult to justify given the potential wide-scale ramifications of climate disaster, of which there are very few, if any, winners in the long-run.

regressions, I find no evidence to support a hypothesised negative risk premium for a temperature risk factor. Cross-sectional estimates of the temperature risk premium are tested in the time series to see whether premiums have increased over time. Contrary to expectations of increasing investor awareness of climate risk, I find no evidence for either a negative linear trend in estimated temperature risk premiums, or decadal dummy effects. I create portfolios sorted by temperature betas but find no evidence that portfolios with higher temperature loadings are outperformed by portfolios formed with lower temperature loadings. Neither the equal nor value-weighted *Temphedge* long-short portfolios provide sufficient evidence of negative returns on average, or in excess of control risk factors. Finally, I test temperature betas using an event study and a firm-specific climate variable. Surprisingly, results indicate that industry temperature betas do not predict the outcomes of the Paris Agreement, nor are firm level temperature betas strongly correlated to firm exposure to climate risk, suggesting that historic temperature betas may be poor proxies for exposure to temperature risk.

Further study could engage with the climate science literature to incorporate additional environmental factors in asset pricing tests; additional research is necessary to further disentangle the impacts of complex climate systems on financial markets. Future research could also take a global outlook and estimate the effects of temperature shocks in international markets.

Chapter 2

Is Pollution a Sin? A Study on the Institutional Ownership of Polluter Stocks

2.1. Introduction

Due to the impact of industrial pollution on the environment and human health, many investors and institutions now increasingly include non-financial environmental information in their portfolio formation process. According to The Forum for Sustainable and Responsible Investment (2016), at the beginning of 2016 more than one-fifth of professionally managed portfolios in the U.S. were invested according to socially responsible investment (SRI) strategies, equivalent to $8.72 trillion or more.[40] The importance of environmental performance as a subset of overall corporate social responsibility (CSR) is also frequently highlighted in the reports of professional services firms,[41] illustrating the perceived relevance of environmental governance to economic performance.

The literature has primarily focused on the performance of SRI in financial markets,[42] however this study examines specific environmental interactions with institutional ownership. I deviate from topics on SRI and positive environmental performance, and instead focus on the implications of socially *irresponsible* behaviour on investment; I target environmental 'sin' stocks by examining polluting firms. Polluters generate negative externalities in markets, impose costs on society, and may therefore be discriminated against during the investment decision making process. Using data from the Toxic Release Inventory (TRI), I identify the largest relative polluters in the sample and examine the relationship between their toxic releases and aggregate institutional ownership. Prior literature

[40] The Forum for Sustainable and Responsible Investment (2016) finds that for institutional investors, climate change and carbon emissions are the second biggest socially responsible investing criteria following conflict risk, with $2.15 trillion institutional investor funds tied into related assets in 2016.

[41] For example, a report from Deloitte (2013) states that in order to generate value, firms should be "positioning themselves to anticipate the drivers of regulatory and stakeholder expectations", and "evaluate the company's readiness to respond to the implications of environmental performance", specifically in relation to "operations, brand image, compliance structures and even company valuations" (p. 1). Similarly, a more recent report by Ernst & Young (2017) finds that it is "commonly understood that serious reputational and environmental risks can and do surface, and they can have very real impacts on the bottom line"; investors who use environmental, social and governance screens in their investments point to both the "long-term benefits" and "lower investment risk" of these investments (p. 3).

[42] For example, see Hamilton, Jo, & Statman (1993), Geczy, Stambaugh, & Levin (2005), Galema, Plantinga, & Scholtens (2008), Renneboog, Ter Horst, & Zhang (2008) and Derwall & Koedijk (2009).

does not focus on polluting as an explicit criterion in the classification of sin stocks;[43] this research aims to fill that gap.

The main contribution of this study is to examine whether higher firm pollution levels in the U.S. are associated with a reduced proportion of institutional equity ownership. There is plenty of anecdotal evidence of an increasing investor awareness of environmental performance and climate change, illustrated by the rising prevalence of 'green funds' and pollution divestment campaigns. Social pressures for SRI and discrimination against polluters may restrict ownership of polluter stocks for the average institutional investor. Following Hong & Kacperczyk (2009), financial institutions that have diverse constituents, publicly known positions in stocks, or are easily exposed to public scrutiny are more likely to be constrained by social norms. In contrast, individual investors and inside owners can keep their stock positions relatively opaque and are thus less likely to be affected by social norms.

I hypothesise a negative relationship between pollution and institutional ownership, and test for evidence of a social norm against the ownership of polluter stocks. I primarily examine whether institutional investors in aggregate have lower equity ownership of public firms with greater toxic releases after controlling for other firm level characteristics. Due to an expected increasing awareness of the costs of pollution on human health and the environment,[44] I hypothesise a negative trend in the difference between institutional ownership of polluter and non-polluter stocks over time. The main hypothesis also implies that certain types of institutions may disproportionately own polluters, and that polluter stocks should be less followed by sell-side analysts, as analyst services tend to cater to institutional investors (Hong & Kacperczyk, 2009). Lastly, I hypothesise that polluters generate abnormal returns as a result of societal discrimination (Angel & Rivoli, 1997).

[43] In the literature, sin stocks are usually defined as stocks in the alcohol, gambling and tobacco industries; see Salaber (2007), Fabozzi, Ma, & Oliphant (2008), Hong & Kacperczyk (2009), Salaber (2009), Liston (2016), and Blitz & Fabozzi (2017). Some SRI funds also exclude armament producers and nuclear energy.

[44] For example, Flammer (2013) provides empirical evidence of increasing external pressures over time on firms to be environmentally friendly.

Results are consistent with the hypothesis that polluter stocks are associated with reduced institutional ownership. I find that on average, institutional investors have approximately 4.5% reduced equity in firms in the top yearly quintile of polluters after controlling for ownership trends and various firm and stock characteristics. There is a positive trend on overall institutional ownership; however, when interacted with pollution the coefficient becomes negative, revealing that institutions are increasingly reluctant to invest in polluting firms. Upon disaggregating toxic releases by chemical classification, I find that dioxins and dioxin-like compounds are the most significantly associated with reduced institutional ownership.

This study also contributes by examining the relationship between social norms and the differing types of institutional investors. I repeat ownership tests after disaggregating institutional ownership based on the institutional investor classifications of Bushee (2001). Results reveal that all three Bushee institutional investor groups have varied relative reductions in their ownership of polluter stocks, suggesting that the impacts of environmental social norms are heterogeneous among institutions. Institutions characterised by long-term, diversified buy-and-hold strategies have the most reduced ownership of polluters. A similar result is found if ownership is separated by 13F institution type; banks, insurance companies, endowments and pension funds on average have a reduced ownership of polluters, while mutual funds and independent investment advisers do not. Using security analyst data, I also find a reduced level of analyst coverage for polluter stocks.

I test whether the institutions that own polluter stocks are more likely to have shorter investment horizons due to their role as market arbitrageur. Using a firm level quarterly churn variable as a proxy for average investor horizons, results reveal that polluter stocks are indeed disproportionately held by institutions with shorter investment horizons. Finally, I test whether polluter firms earn abnormal returns due to the shunned-stock effect (Angel & Rivoli, 1997). I create a long-short polluter portfolio and test for abnormal returns benchmarked against popular risk factor models, however, I find no evidence of abnormal returns.

2.2. Literature review

I review the literature by first defining societal discrimination within an economic context. I then examine the relationships between social norms and SRI. Lastly, I review studies on sin stocks.

2.2.1. Societal discrimination

In a seminal contribution, Becker (2010) describes how market participants are willing to incur financial costs in order to avoid things they dislike. This discriminatory behaviour is linked to disutility generated through contact or association. Societal discrimination can create social norms, which are standards that constrain society members from contact with the discriminated. Akerlof (1980) and Romer (1984) show that social norms are insensitive to arbitrage; social norms may persist if they cause a loss of reputation for the party engaging in discriminatory behaviour that outweighs the benefits of the same behaviour. In the context of this study, social norms refer to the public pressure on investors to divest from socially discriminated polluting firms.

2.2.2. Socially responsible investing

Pollution is modelled as a negative externality in economics; related investments can be thought of as socially irresponsible. While SRI is an example of investor discrimination based on ethical values, it is also sometimes justified as a portfolio performance enhancer. A large portion of the literature is dedicated to the investment returns of SRI; theoretical and empirical studies have contrasting conclusions on the relationship between SRI and portfolio returns (Galema, Plantinga, & Scholtens, 2008). Hamilton, Jo, & Statman (1993) and Statman (2000) reveal that SRI equity funds do not generate significantly different risk-adjusted abnormal returns relative to conventional funds. Similarly, Derwall & Koedijk (2009) find that SRI fixed income funds perform approximately the same as their peers. Renneboog, Ter Horst, & Zhang (2008) find an insignificant SRI fund alpha in most countries. In contrast, Geczy, Stambaugh, & Levin (2005) show that imposing SRI constraints on investments negatively impacts returns, consistent with the theory of reduced opportunity sets in standard portfolio theory. Von Wallis & Klein (2015) provide a comprehensive literature review, listing the various hypotheses of SRI over, under, and in-line performance, along with empirical evidence for each. Von Wallis & Klein's (2015) review illustrates the lack of academic consensus on the consequences of SRI on performance.

Investment strategies may be influenced by social norms which are often motivated by ethics, values and biases. For example, Ivković & Weisbenner (2005) find that individual investors exhibit a preference for local investments. Different groups of discriminating investors have varying preferences for specific stock characteristics; women place greater weighting on stocks with progressive gender policies, while younger investors avoid firms with poor environmental records (Hood, Nofsinger, & Varma, 2014). Similarly, in a study focusing on the importance of political values as a cultural factor that influence asset allocation, Hong & Kostovetsky (2012) find that institutions with liberal values are more likely to discriminate against investments that are deemed socially irresponsible, such as tobacco, natural resource overexploitation, and weapons. Using investor holdings turnover as a proxy for investor horizons, Starks, Venkat, & Zhu (2017) find that institutional investors with longer-term horizons prefer firms with higher ESG scores. These results indicate that socially responsible investors are not a single homogeneous group, but instead exhibit varied discriminatory behaviour. Societal discrimination has thus been found to affect the investment decisions of not only particular groups of individual investors, but also institutions.

2.2.3. Sin stocks

The study of investment 'sin' is a relatively new branch of the literature in economic discrimination; examining social perceptions of 'bad' activities or participants, the subject area is inversely related to SRI. The financial literature focuses on the performance and behaviour of financial assets that are perceived as sinful, such as the shares of tobacco or gaming firms. Sin stocks are more likely to be scrutinised by society, and attract a disproportionate level of discrimination. For example, Kim & Venkatachalam (2011) find that financial reporting quality is greater for sin firms, possibly to compensate for their poor public perception. Additionally, audit and consulting fees are found to be higher for companies that deviate from social norms (Leventis, Hasan, & Dedoulis, 2013), suggesting that discriminated firms are penalised with various additional costs.

Institutional ownership studies have investigated the drivers of institutional interest in equity investments. In their seminal paper, Gompers & Metrick (2001) find a set of financial variables that explain variation in total institutional ownership of a stock; I use these variables in my models. Hong & Kacperczyk (2009) specifically examine the effect of sin on institutional ownership. Identifying alcohol, gambling and tobacco as sin stocks, they find that both the institutional ownership and analyst coverage of sin stocks is reduced due to social norms that target specific sub-industries. Specifically, the effects of social norms are found to constrain publicly scrutinised investors that include banks, insurance companies and endowments; however, there is no evidence to suggest that arbitrageur institutions are similarly constrained in their investments.[45] Arbitrageurs may be less likely to care about social norms, or be unwilling to sacrifice good investment opportunities to satisfy society. Liu, Lu, & Veenstra (2014) also explore the interactions between sin stocks and investment, finding that the institutional ownership and analyst coverage of sin stocks are positively correlated with the degree of social acceptance of the specific sin. This is of relevance to environmental sinners, since social acceptance of polluters is likely to diminish as environmental issues become more pertinent over time.

Polluter stocks are not included in the set of sin stocks within the studies by Hong & Kacperczyk (2009) or Liu et al. (2014). Unlike firms operating in the alcohol, gambling and tobacco industries, polluters can change their pollution levels over time, generating time-variation in their sinner status. Fernando, Sharfman, & Uysal (2017) consider the effects of ‘greenness’ on institutional ownership. Using KLD data, firms that are identified as both ‘green’ and ‘toxic’ are found to have lower institutional ownership. I similarly consider environmental performance in relation to institutional ownership; however, I avoid using KLD data as the discrete environmental scores have limited

[45] Arbitrageur institutions are loosely defined as those with primary aim of exploiting security mispricing.

variation, and lead to clustered observations. I instead opt for continuous, audited and more objective data on pollution, sourced from the TRI, which allows for more granular analysis.

Kim, Wan, Wang, & Yang (2019) also use TRI data in an institutional ownership study. They argue that a prevalence of local institutional investors leads to reduced pollution at the facility level; pollution abatement efforts are also found to increase firm value when there is an increased proportion of local investors relative to firm facilities. Their research considers an opposing causal driver to my study and is accordingly conducted at the facility level. I argue that investment decisions are more likely to be made at the firm level, and therefore environmental screens are more likely to drive lower investment in polluting firms, rather than local shareholder pressure reducing aggregate firm level pollution. I do, however, test the latter channel in three robustness tests of reverse causality, but find no evidence of institutional ownership driving firm pollution.

The shunned-stock hypothesis assumes that the shortage of demand for sin stocks will impact the behaviour of their prices (Derwall, Koedijk, & Ter Horst, 2011). Angel & Rivoli (1997) extend Merton's (1987) segmented information model to argue that shunned controversial stocks generate higher expected returns in proportion to the level of socially responsible investors in the market. Hong & Kacperczyk (2009) argue that institutional aversion to sin stocks causes their prices to be relatively cheaper, and subsequently generate higher expected returns. Consistent with the hypothesis of Becker (2010), Hong & Kacperczyk's (2009) empirical results support the theory that institutions incur opportunity costs in abstaining from stocks which are shunned by society. Similarly, Salaber (2007), Fabozzi, Ma, & Oliphant (2008), Statman & Glushkov (2009) and Derwall et al. (2011) also find that controversial sin stocks produce abnormally high returns. However, Lobe & Walkshäusl (2016) and Blitz & Fabozzi (2017) contrastingly find no evidence of a risk adjusted alpha for sin. Motivated by the shunned-stock literature and subsequent conflicting empirical findings, I test for evidence of polluter abnormal returns with a long-short portfolio of polluter stocks.

2.3. Hypothesis

Institutions are hypothesised to be constrained by social norms, and are expected to abstain from investing in firms that are discriminated by society as environmental sinners. In contrast, individual investors can hide their holdings with relative ease (Hong & Kacperczyk, 2009), and as a result are able to own more equity of polluter firms. Due to reduced institutional interest in these stocks, I also expect reduced analyst coverage of polluters. I further hypothesise that institutions investing in polluters have shorter investment horizons, and that polluter stocks earn abnormal returns through the shunned-stock effect.

2.4. Data

Data used in this study is sourced from four primary sources: the TRI, Compustat, CRSP and Thomson Reuters institutional holdings. I use IBES for security analyst data.

2.4.1. Pollution and fundamentals data

I obtain data on firm releases of pollutants from the Environmental Protection Agency's (EPA) Toxic Release Inventory database.[46] The TRI contains information on the releases of toxic chemicals in the U.S. that may damage the environment and human health. Toxic releases include carcinogens and persistent bio-accumulative toxic chemicals.[47] Release disclosure via TRI is a mandatory program that covers over 50,000 industrial facilities and 500 different chemicals. Firms are required to disclose annual releases of toxic chemicals to the EPA if they employ 10 or more full-time employees, operate in a pollution prone industry, and handle or manufacture TRI-listed chemicals above threshold levels. The EPA releases the yearly TRI National Analysis dataset during December or the following January. The toxic releases data dates to 1987 and covers industries, including mining, utilities,

[46] The TRI was initially established under the 1986 Emergency Planning and Community Right-to-Know Act of 1986, and later expanded with the Pollution Prevention Act of 1990. The TRI was established in reaction to an industrial disaster in Bhopal, India in December 1984, along with a similar chemical release that occurred in West Virginia, 1985.

[47] The chemicals covered by the TRI program are generally those that are linked with cancer and other chronic health effects, significant acute health effects or significant environmental damages. Greenhouse gas emission is not explicitly covered by the TRI program which instead focuses on toxic chemicals, though is some overlap between the two.

manufacturing, publishing, and hazardous waste.[48] TRI data has been used by regulators, media, and environmental activists (Hamilton, 1995), while the economic and financial literature has also used toxic releases from the database to study CSR activity.[49]

The sum of total facility on-site releases, off-site releases, and transfer of releases to public owned treatment works is stored as *Total Releases*, measured in millions of pounds.[50] I aggregate *Total Releases* by firm and year.

I use the Compustat database to source fundamental accounting variables as at the end of year *t*. I merge firm TRI data with firm fundamentals using the CRSP/Compustat merged database, and create a dummy variable named *Polluterdummy* to identify the largest relative polluters. $Polluterdummy_{i,t}$ is activated if firm *i* is in the top quintile of polluters in year *t* in the TRI database. The primary advantage of using a dummy variable is to avoid imposing any assumptions about the structural relationship between pollution and institutional ownership; instead, my main models simply focus on the characteristics of the worst polluters. Using CRSP, I then merge market data on firm equities with the TRI-Compustat dataset. I only include securities with CRSP share codes of 10 or 11, and store returns in percentage format. To be included in the final sample, a firm must exist on all three of these databases.[51] I drop firms that operate in the financial services industries with one digit SIC codes of 6 from the sample (Hong & Kacperczyk, 2009).[52] This provides a total of 8,954 firm-year observations in an unbalanced panel over the final sample period of 1987 to 2014.

---

[48] One limitation of the TRI is that data is self-reported, however measurement error is mitigated through audits run by the EPA. The data also focuses on manufacturing industries; however, it is not immediately clear as to why estimates can not be extrapolated to other sectors. For greater discussion of limitations, see Kim et al. (2019).

[49] For example, see Dooley & Lerner (1994), Hart & Ahuja (1996), Maxwell, Lyon, & Hackett (2000), and King & Lenox (2001).

[50] I avoid transforming *Total Releases* with a natural log as it imposes a structural relationship between pollution and institutional ownership which has not previously been theorised or identified in the literature. However, I find that all tests that use log *Total Releases* generate consistent results; I present some of these tests in the appendix.

[51] I only sample firms with calendar year-end financial reporting dates to match the timing of variables in the cross-section. I find that results are largely constant if I include firms with earlier financial year-ends.

[52] I find that main results are consistent if firms with book values of less than $10m are excluded.

I illustrate the distribution of polluters in the sample by plotting the yearly median and cumulative distribution of *Total Releases* in Figures 2.1 and 2.2.

Figure 2.1: The time series of the median yearly values of *Total Releases*. *Total Releases* is the total toxics released by a firm in a single year, measured in millions of pounds. I source data from the Toxic Release Inventory Program, URL: *https://www.epa.gov/toxics-release-inventory-tri-program/tri-basic-data-files-calendar-years-1987-2016*.

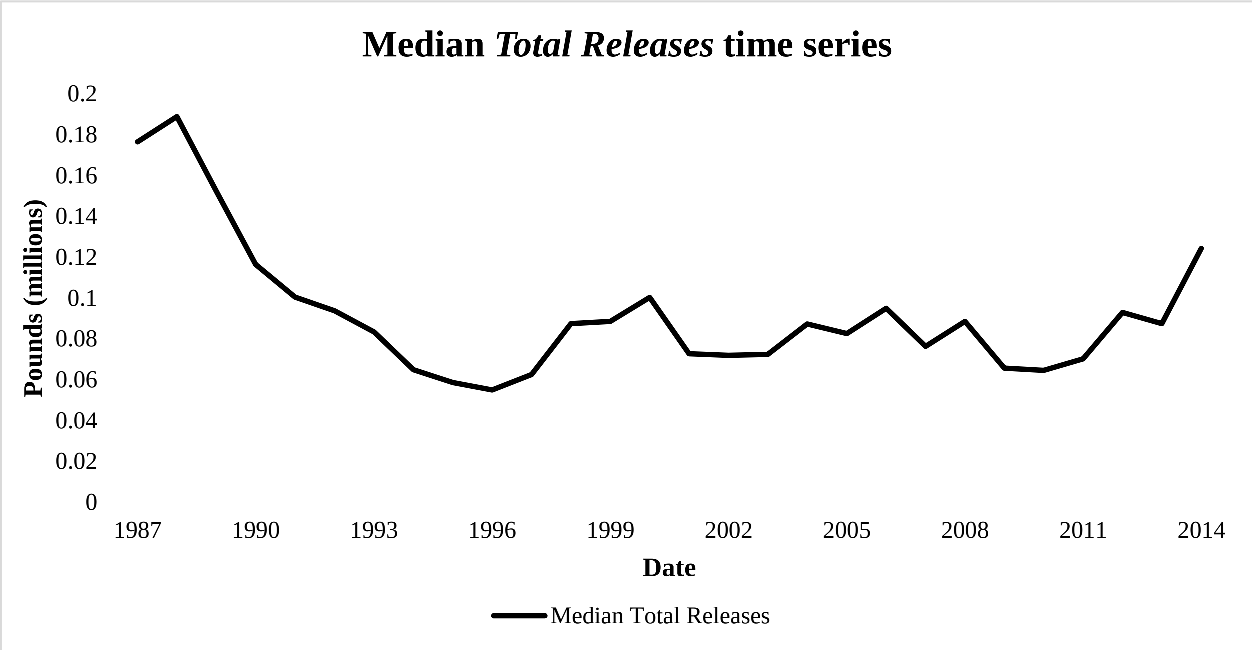

Figure 2.2: Cumulative distribution function of *Total Releases* in millions of pounds. *Total Releases* are represented by the x-axis, while the cumulative proportion of observations that have equal or lower values of those *Total Releases* are represented by the y-axis.

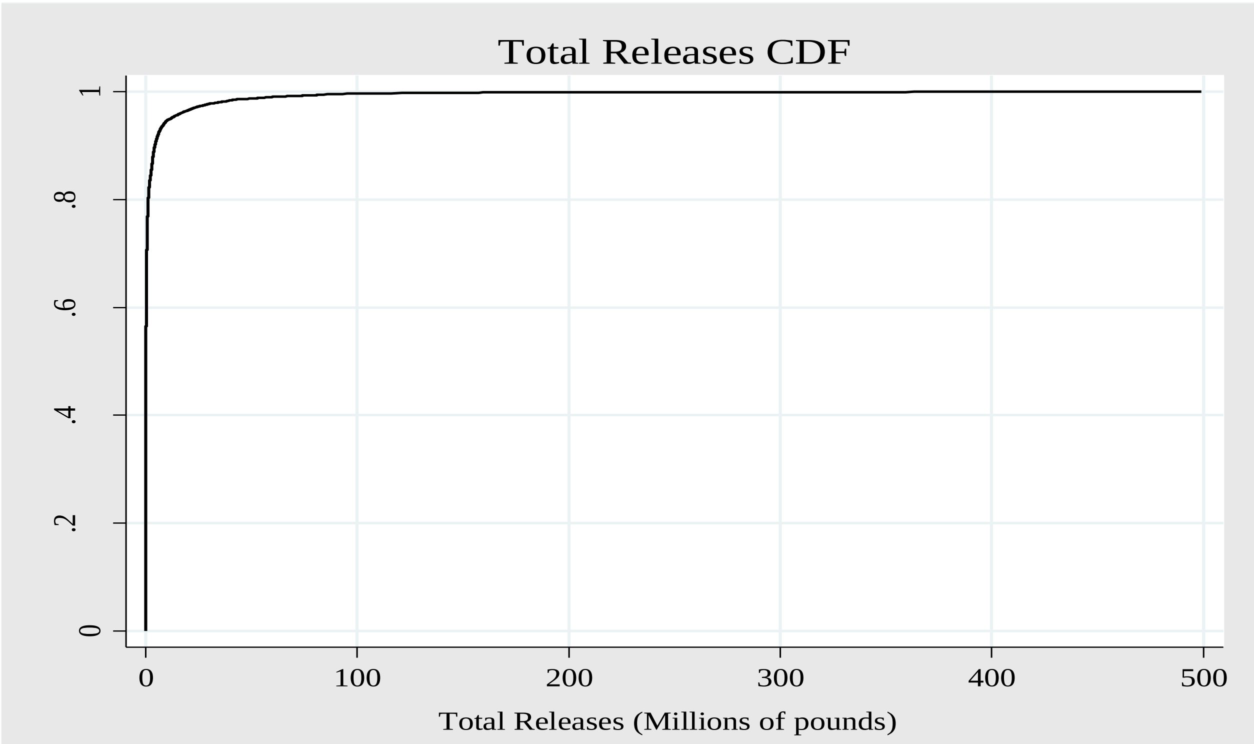

Figure 2.1 reveals that the median toxic releases fall sharply in the early 1990's following the introduction of the TRI program, and then becomes relatively stationary. This is possibly due to increasing environmental regulation, increasing abatement efficiency or a different sample of firms in the data. Figure 2.2 highlights how a minority of polluting firms are accountable for a disproportionately large amount of toxins in the sample. The top quintile of polluters within *Polluterdummy* therefore consists of these extreme firms as well as relatively moderate polluters.

2.4.2. Institutional ownership and analyst coverage data

I obtain data on institutional holdings of equity from the Thomson Reuters Institutional 13F Holdings database.[53] The 13F database contains information about institutional investors with $100 million or more in assets under management. Institutions comprise of banks, insurance companies, mutual funds, investment advisers, and others.[54] $IO_{i,t}$ measures the percentage of firms *i*'s shares outstanding that are owned by an institution at the end of year *t*; despite being a quarterly report, most firms only file timely 13F reports in June and December (Hong & Kacperczyk, 2009). Firms with a missing value of *IO* are assumed to have 0 institutional ownership.[55]

I follow Hong & Kacperczyk (2009) in their choice of control variables used in institutional ownership regressions. $INDBETA_{i,t}$ is the CAPM beta for firm *i*'s industry and controls for industry level market risk.[56] $LOGSIZE_{i,t}$ is calculated by taking the natural logarithm of firm *i's* market capitalisation plus 1 at the end of year *t*, and is a measure of firm size. $LOGBM_{i,t}$ is calculated by taking the natural logarithm of 1 plus firm *i's* annual book value of equity divided by market

[53] The database contains several issues that are highlighted in Geertsema (2014). I follow his methodology in addressing these issues.

[54] 'Investment advisers' includes hedge funds along with an assortment of other institutions. 'Others' includes pension funds, foundations, endowments and universities.

[55] There are only 33 observations in the final sample which have missing values for *IO*. Primary results are unchanged if these observations are excluded.

[56] Industry betas are calculated by regressing each of the Fama-French 49 industry portfolio excess returns (Fama & French, 1997) against market excess returns in 60-month rolling window regressions. The monthly betas of each industry are then averaged by year to convert *INDBETA* to a yearly frequency. I choose industry betas over firm betas to present results that are robust to the controls used in Hong & Kacperczyk (2009), and because some individual firms have less than 60 months of returns data. Industry betas are also less noisy estimates than firm betas. However, I still find that main results are consistent when *INDBETA* is replaced with firm-specific market betas.

capitalisation as at the end of year *t*.[57] $STD_{i,t}$ is the standard deviation of daily returns for the shares of firm *i* for year *t* and measures return volatility. $PRINV_{i,t}$ is the inverse of firm *i's* share price at the end of year *t* and controls for raw price effects. $RET_{i,t}$ is the average monthly return for firm *i's* shares during year *t*. I also use two dummy variables, $NASD_{i,t}$ and $SP500_{i,t}$ which are respectively activated if firm *i* is listed on the Nasdaq or is a constituent of the S&P 500 index during year *t*.

In tests examining analyst coverage, I use a panel dataset of the number of analysts that cover a stock, obtained from IBES.[58] $LOGCOV_{i,t}$ is the natural logarithm of 1 plus the monthly average number of analysts who cover firm *i* during year *t* and have provided a forecast for the next annual earnings announcement.[59] Stocks that are not included in the IBES dataset are assumed to have 0 average analyst coverage. I use the same control variables from institutional ownership tests in analyst coverage regressions. Summary statistics of ownership variables and analyst coverage are presented in Table 2.1.

[57] For firms with negative book values of equity, I set the book-to-market ratio as 0. Results are also consistent if I instead drop these firms from the sample.

[58] The number of analyst forecasts made for a stock in a month is sourced from the IBES summary file, which is desirable as it excludes outlier estimates in the data and represents only realistic forecasts. In unreported results, I find that recreating data using the IBES detail file and including outlier forecasts generates qualitatively similar results as found in the reported analyst forecast tests.

[59] Setting *LOGCOV* equal to average monthly coverage is preferred to using the total number of analysts following a firm; observations with one analyst covering a firm once in the year would otherwise be treated equivalent to those with one analyst covering the firm for the entire year. The monthly average method is also preferred to using total analyst coverage as at year-end due to mitigate heterogeneous seasonality effects.

Table 2.1: Summary statistics for the sample panel data used in institutional ownership and analyst coverage regressions. Variable means are presented for the full sample and for the subset of polluters and non-polluter observations. Polluters are identified with *Polluterdummy*. Explanatory variables are reported below the dependent variables. Significant differences in averages between polluters and non-polluters at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Sample summary statistics

Variable	Full sample	Polluters	Non-polluters	Difference
IO	0.58	0.58	0.58	0.00
LOGCOV	1.72	1.94	1.66	0.29***
INDBETA	1.04	0.89	1.08	-0.19***
LOGSIZE	20.88	21.90	20.63	1.27***
LOGBM	0.42	0.42	0.43	0.00
STD (%)	2.58	2.25	2.67	-0.42***
PRINV	0.08	0.06	0.09	-0.03***
RET (%)	1.31	1.24	1.33	-0.09
NASD	0.17	0.06	0.20	-0.14***
SP500	0.36	0.60	0.30	0.29***
N	8,954	1,809	7,145	

In a simple test of averages, there is no difference between institutional ownership based on polluter status, however polluters have greater analyst coverage. In line with expectations, polluters are larger in size and are more likely to be listed on the S&P 500. Polluter firms also have lower standard deviation of returns and operate in industries with lower market betas. The price of polluter stocks also tends to be higher, however, there is no significant difference between the returns of polluter firms and non-polluters.

2.4.3. Polluter performance data

To test the validity of the shunned-stock hypothesis for polluter stocks, I examine the performance of polluter portfolios. I source stock return data from CRSP. I exclude returns on non-domestic equities. I follow Shumway (1997) in correcting for delisting biases.[60] Benchmark factor models used include the Capital Asset Pricing Model, the Fama-French 3-factors (Fama & French, 1993) and the

[60] If delisting returns in the panel data have a delisting stock code of 500, 520, between 551 and 573 inclusive, 574, 580 or 584, returns are set to -30%; while a missing delisting return with an available delisting code has returns set to -100%.

Carhart 4-factors (Carhart, 1997). I source risk factor data from Kenneth French's data library.[61] Realised stock returns and benchmark risk factors are of monthly frequency in portfolio tests.

2.5. Main results

The main tests in this study examine the relationship between firm pollution, institutional ownership, and analyst coverage in accordance with the societal discrimination hypothesis.

2.5.1. Institutional ownership of polluter stocks

I primarily consider whether firm pollution is negatively associated with institutional ownership of the polluter's stocks. I hypothesise that institutional investors shy away from investing in firms that pollute heavily due to social norms. Also, with increasing environmental concerns among the public and institutions (Flammer, 2013), societal discrimination against pollution is likely to increase over the sample period; I therefore expect decreasing institutional ownership in polluter stocks over time. In a preliminary test, I use a simple yearly average of the *IO* variable for polluters and non-polluters classified by *Polluterdummy*. I also generate a yearly average *IO* for the aggregate sample. I present the time series of yearly averages across groups in Figure 2.3.

[61] Kenneth French data library URL: *http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html*.

Figure 2.3: A time series of the average institutional ownership of firms by *Polluterdummy* classification. The trendlines are generated by taking an average of *IO* for both polluters and non-polluters, as well as the entire aggregated sample.

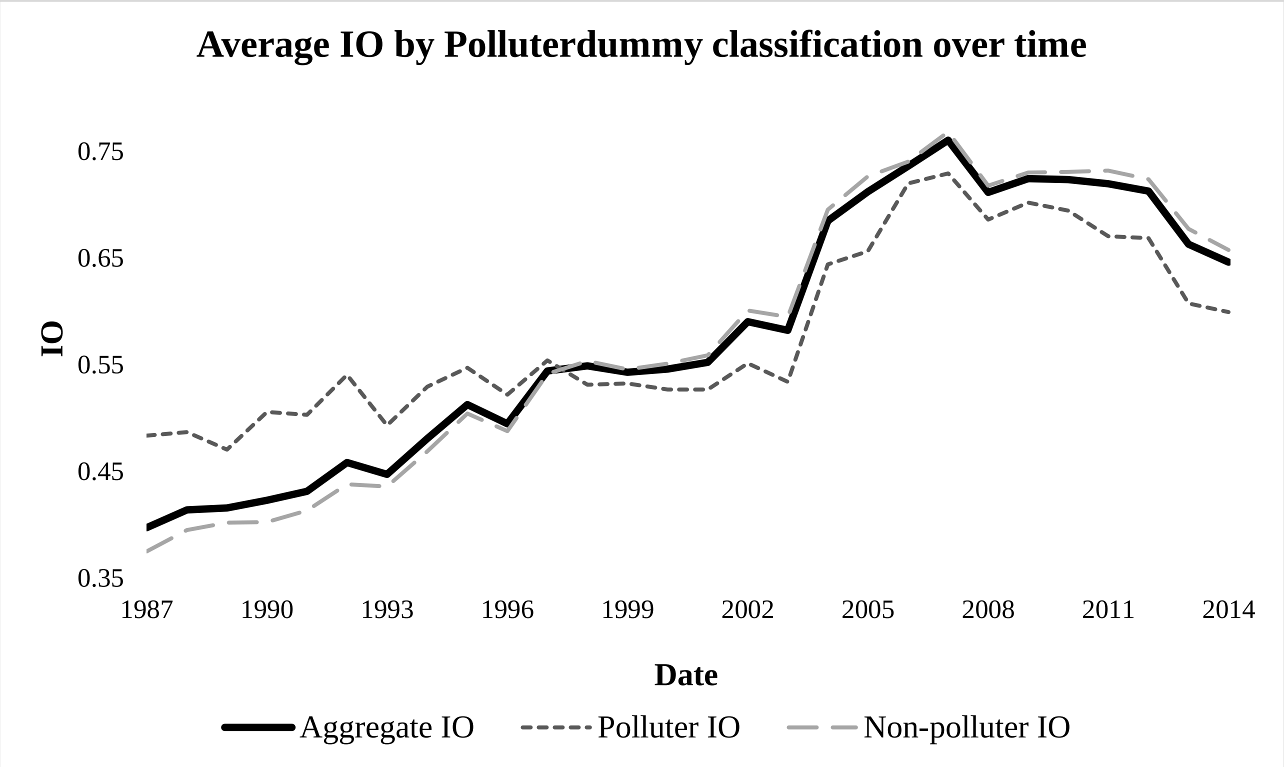

The average *IO* for both the non-polluters and the aggregate sample is increasing over the sample period. These two time series follow each other closely as there are far more non-polluters in the sample that contribute to aggregated sample averages. There is a large increase of ownership leading up to the global financial crisis of 2008, which then levels off following the aftermath. Interestingly, the institutional ownership of polluter stocks starts off higher than non-polluter stocks and is also increasing over the time period, though at a slower rate than that of non-polluters and the aggregate sample. 1998 is the point of intersection between the series, one year after the Kyoto Protocol. The time series are generated without controlling for firm size or stock returns, and hence may be the result of a confounding variable driving differences in the time series.

I run the following fixed effects panel regression to estimate the negative relationship between institutional ownership and *Polluterdummy*, after accounting for ownership control factors.

$$IO_{i,t} = \beta^{poluter} * Polluterdummy_{i,t} + \boldsymbol{\beta^{control}} * \boldsymbol{CONT_{i,t}} + \varepsilon_{i,t} \qquad (2.1)$$

Institutional ownership of a stock is regressed against the *Polluterdummy*, which is equal to 1 if the firm is in the top quintile polluters in the sample in a given year, and is equal to 0 otherwise. ***CONT*** is a vector of control variables that include *LOGSIZE, LOGBM, STD, INDBETA, PRINV, RET, NASD* and *SP500*.[62] The parameter of interest is $\beta^{polluter}$, which is the estimated impact on institutional ownership of being a large polluter. Under the alternative hypothesis, $\beta^{polluter}$ is negative. On average, aggregate *IO* has a positive trend; I account for time heterogeneity in the dependent variable by including yearly fixed effects (Gormley & Matsa, 2013). Hong & Kacperczyk (2009) use the "ultra-conservative" (p. 24) approach of adjusting standard errors by clustering on Fama-French (1997) industry groups. I follow their approach and additionally include clustering by year in a panel regression, thus making significance estimates even more conservative (Petersen, 2009).

I repeat regression *(2.1)* with *Total Releases* as an independent variable in place of *Polluterdummy*. Using *Total Releases* tests for a linear relationship between pollution and institutional ownership.[63] Using *Polluterdummy,* I also implement industry fixed effects in the regression based on the Fama & French (1997) industries; this model examines whether polluters have reduced institutional ownership within their industry groups.

For robustness, I also repeat regression *(2.1)* with a linear trend variable as an alternative control for time-varying heterogeneities. I then interact the linear trend with *Polluterdummy* to test whether institutional reluctance to invest in the largest polluters has changed over time. If societal discrimination of polluter stocks has increased over the sample period, the estimated interaction coefficient should be negative. Though yearly fixed effects provide greater flexibility in the model specifications, using a simple yearly trend is useful in estimating a smoothed average in ownership

[62] I avoid including corporate governance variables in main tests due to concerns of simultaneity with *IO* and a reduced sample size; however, in a robustness test I find a consistent relationship between *IO* and *Polluterdummy* after controlling for KLD governance variables, which include managerial compensation, low governance reporting transparency, total number of governance strengths and total number of governance concerns. I report the results of the robustness test in the appendix. Another potential robustness test could include firm R&D controls, however, R&D data stored in Compustat is largely unrelated to the type of R&D that is relevant to pollution abatement.

[63] *Polluterdummy* identifies the top 20% of polluters by year which is not necessarily the same top 20% of polluters in aggregate. Estimated coefficients for *Polluterdummy* thus have a slightly different interpretation to estimates for the top quintile of *Total Releases* on its own, albeit marginally.

trends along with interactions with *Polluterdummy*. I present the results of all five specifications of regression *(2.1)* in Table 2.2.

Table 2.2: Results of the institutional ownership panel regressions where the dependent variable is *IO*. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 8,954 firm-year observations in the sample for each specification. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Institutional ownership panel regression results

Variable	(1)	(2)	(3)	(4)	(5)
Polluterdummy	-0.0429***		-0.0267	-0.0448***	0.0131
	(-3.03)		(-1.52)	(-3.24)	(0.91)
Total Releases		-0.0003			
		(-1.35)			
t				0.0098***	0.0107***
				(6.39)	(7.02)
*Polluterdummy * t*					-0.0040***
					(-5.60)
INDBETA	0.1052***	0.1110***	0.0310*	0.0902***	0.0868***
	(3.70)	(3.52)	(1.71)	(3.65)	(3.66)
LOGSIZE	0.0428***	0.0411***	0.0489***	0.0464***	0.0459***
	(6.37)	(5.97)	(5.97)	(6.47)	(6.37)
LOGBM	0.0186	0.0129	0.0466**	0.0174	0.0180
	(0.91)	(0.60)	(2.38)	(0.79)	(0.84)
STD	-0.0129*	-0.0131*	-0.0142**	-0.0028	-0.0029
	(-1.71)	(-1.69)	(-2.20)	(-0.39)	(-0.40)
PRINV	-0.1076**	-0.1102**	-0.0855*	-0.1322**	-0.1322**
	(-2.16)	(-2.17)	(-1.78)	(-2.49)	(-2.50)
RET	-0.0013	-0.0014	-0.0006	-0.0028*	-0.0028*
	(-1.22)	(-1.27)	(-0.65)	(-1.91)	(-1.91)
NASD	-0.0554***	-0.0531***	-0.0561***	-0.0575***	-0.0578***
	(-3.96)	(-3.78)	(-3.97)	(-4.15)	(-4.16)
SP500	-0.0246	-0.0265	-0.0290	-0.0315	-0.0318
	(-1.27)	(-1.34)	(-1.44)	(-1.59)	(-1.61)
Fixed effects	Year	Year	Year & Industry	None	None
N	8,954	8,954	8,954	8,954	8,954
Adjusted R^2	0.4068	0.4027	0.4505	0.3737	0.3764

Results are consistent with hypotheses of societal discrimination against polluters and increasing environmental awareness. In columns (1) and (4), the estimated *Polluterdummy* coefficients are significantly negative, revealing that institutions own proportionately less equity in firms that pollute the largest quantities in a year, as expected. Polluting firms are associated with approximately 4.5% less institutional ownership of their stocks on average compared to the institutional ownership of non-polluters within any year. Interestingly, the association between pollution and ownership appears to be non-linear; results of the *Total Releases* regression does not generate a statistically significant coefficient. I explore this non-linearity in the next section.

The use of yearly fixed effects generates a statistically significant coefficient for *Polluterdummy*. However, once industry fixed effects are included, significance disappears. The industry and yearly fixed effects model provides no evidence that polluters have reduced institutional ownership once industry averages are accounted for. This finding suggests that average industry pollution matters, but the within-industry effects of pollution are insignificant. In a similar model reported in the appendix, I use firm fixed effects instead of industry fixed effects. The results of this additional test provide no evidence to suggest that deviation from firm averages in pollution is associated with within-firm variation in institutional ownership. The weak relationships at the within-industry and within-firm levels indicate that institutional reluctance to own polluters is not granular enough to specifically target polluting activity, but rather is a broader social stigma which appears to affect industries and firms that are perceived as polluters.

Consistent with Gompers & Metrick (2001), the linear trend in columns (4) and (5) is positive and statistically significant, indicating that after controlling for firm characteristics, institutions are increasing their proportional holdings of equity by approximately 1% each year on average. However, the polluter-time interaction is significantly negative, revealing that institutions are reducing their holdings of polluters relative to non-polluters by approximately 0.4% each year on average; or equivalently, institutional ownership of polluters is also increasing but at a slower rate than that of non-polluters. Consistent with growing environmental awareness, this interaction effect absorbs significance from the *Polluterdummy* coefficient, and with 28 years in the sample, reverses the positive coefficient too. The interaction estimate is also consistent with Figure 2.3, where initially the ownership of polluter stocks is higher than non-polluters, but then reverses in the late 90's.

Consistent with the literature, industry beta and size appear to have a positive association with institutional ownership, while being listed on the NASDAQ has a negative association. *STD* has a negative coefficient, revealing that volatile stocks are associated with reduced institutional ownership on average; however, the effect is only significant in fixed effects models. The coefficient of *PRINV* suggests that institutions hold more expensive shares. Ownership is also found to be negatively related to the stock's average performance in the past year, but only when fixed effects are excluded. The main results are overall in line with the hypothesised social norms on polluter stocks and provide evidence of a reduced institutional ownership of polluter firms.

In an additional test, I repeat the yearly fixed effects regression *(2.1)* and include *Sindummy*, a dummy variable activated for the list of firms included by Hong & Kacperczyk (2009) in the 'Triumvirate of Sin'. I present the results of this test in the appendix. Estimated coefficients for *Sindummy* are also negative and almost three times as high as those for *Polluterdummy*, indicating that polluter investments are relatively less discriminated against compared to traditional sin industries.

I examine the association between *IO* and *Total Releases* by repeating the fixed effects panel regression *(2.1)* but with 9 polluter decile dummy variables; one for each yearly decile of *Total Releases* relative to other firms in the TRI. Firms that pollute in the lowest decile in a year have no dummy variable activated, while firms that pollute in the 2^{nd} decile (*Total Releases* is between the 10^{th} and 20^{th} percentile) have the corresponding 2^{nd} decile dummy variable activated, and so on. I use the complete set of control variables shown in Table 2.2, and use the yearly fixed effects model. I graph the estimated dummy coefficients from this test in Figure 2.4. This test highlights the average relation between firm yearly rankings of *Total Releases* and institutional ownership.

Figure 2.4: Pollution decile dummy coefficients, estimated from a yearly fixed effects panel regression with *IO* as the dependent variable. Decile 1 has no active dummy, and is therefore the benchmark against which the following dummy coefficients are compared against.

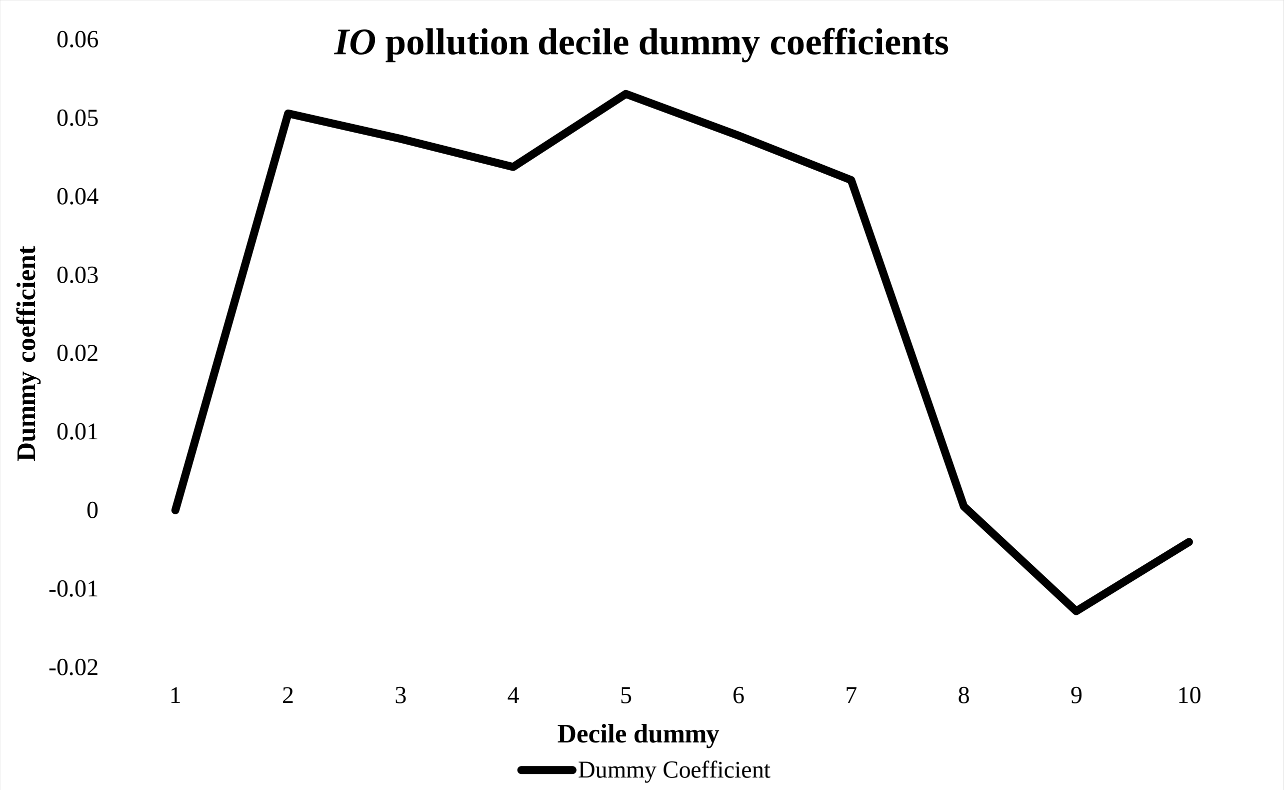

Results reveal a non-monotonic relationship between toxic releases and institutional ownership, explaining the insignificant estimated coefficient for *Total Releases* in Table 2.2 column (2). Figure 2.4 suggests that institutional holdings are more sensitive to pollution once toxic releases exceed a threshold, shown by the sharp decrease in coefficients following decile 7 firms. Institutions appear to prefer firms with moderate levels of toxicity. Figure 2.4 suggests that the greatest polluters are disproportionately discriminated against relative to moderate polluters. Consistent with Fernando et al. (2017), relatively green firms in the lowest decile of pollution also appear to have reduced institutional ownership relative to the middle deciles, suggesting a negative institutional reaction to firm pollution being on either extreme of the spectrum.

Given the presence of extreme polluters as seen in Figure 2.2, I repeat the yearly fixed effects regression *(2.1)* with clustered standard errors, and include two additional dummy variables for the top 5% and 1% of polluters by year. The coefficients of these dummy variables represent the marginal reduction in institutional ownership for more extreme polluter percentiles. The estimated coefficient for the original *Polluterdummy* is -0.037 with significance at the 5% level, and -0.021 and -0.07 for the 95th and 99th percentile polluter dummy variables respectively, both of which are insignificant at the 10% level. Figure 2.4 and these insignificant estimates suggest that the most extreme polluters do not have further reduced institutional ownership compared to firms in the top quintile of pollution. It is possible that the original *Polluterdummy* benchmark closely approximates the average institutional pollution thresholds used in negative screens; firms that pollute in greater percentiles may already be excluded by most institutional investors and thus experience limited marginal decreases in ownership.

As seen in Table 2.2 column (5), the polluter-time interaction is significantly negative, revealing that institutions have reduced their ownership of polluter stocks over time relative to their holdings of non-polluters. However, the interaction is a smoothed slope and does not explicitly show yearly changes in the gap between the institutional ownership of polluters and other firms. I estimate annual differences in the ownership of polluters relative to non-polluters by conducting the yearly fixed effects panel regression *(2.1)* and interacting each of the yearly dummy variables with *Polluterdummy*. The year 1987 has no dummy and is therefore the benchmark to which the fixed effect interaction coefficients are compared against. I plot the interaction coefficients generated by both the linear trend and yearly fixed effects models in a time series, illustrated in Figure 2.5. The estimated coefficients can be interpreted as marginal yearly effects on the difference in institutional ownership between polluters and non-polluters.

Figure 2.5: Estimated interaction coefficients between *Polluterdummy* and time, where the dependent variable is *IO*. Coefficients are estimated using the *Polluterdummy*-time interaction trend in the main regression, as well as *Polluterdummy*-yearly fixed effects interactions. 1987 has no active dummy, and is therefore the benchmark from which the dummy interaction coefficients are compared against. The estimated *Polluterdummy* coefficient is 0.0131 when using the linear trend interaction model and is 0.0339 when using the fixed effects interaction model.

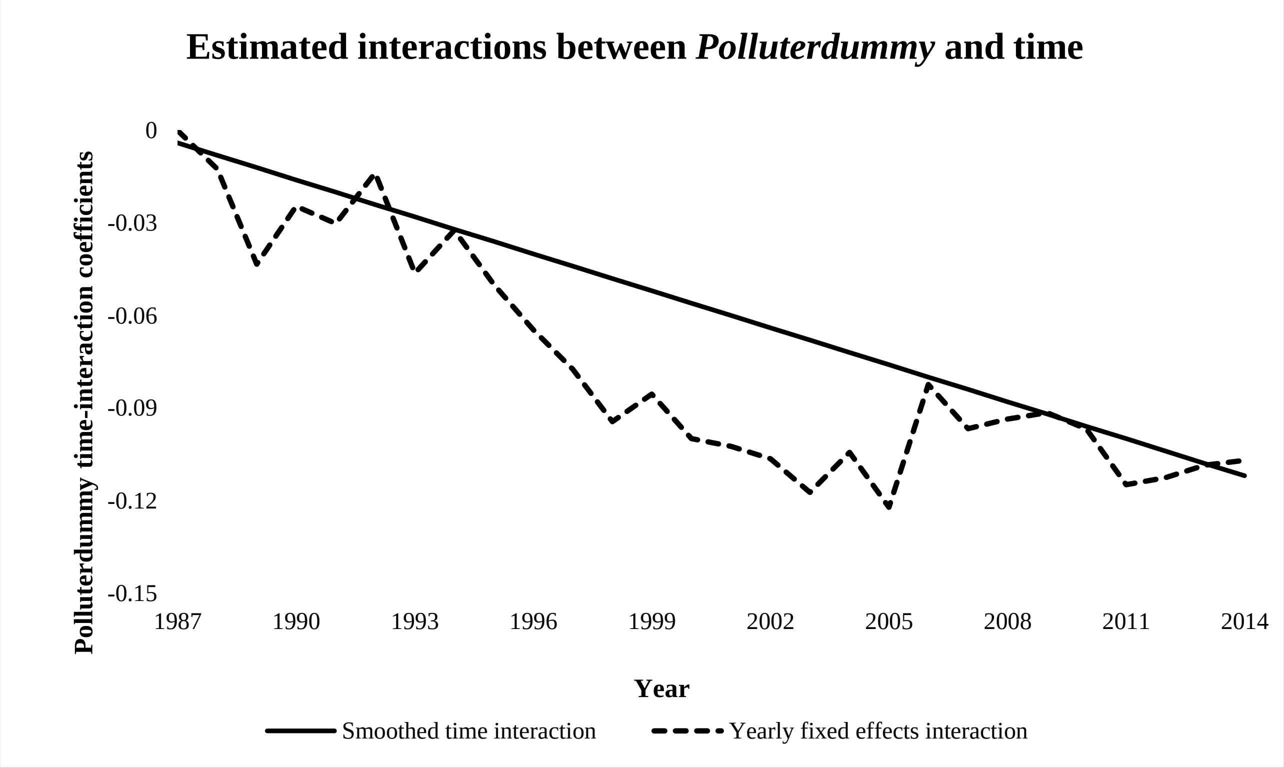

Results reveal that the largest decrease in the interaction coefficients occurred in the period between 1994 and 2003; this is roughly consistent with Figure 2.3. Institutions decreased their holdings of polluter firms relative to non-polluters in the early 1990's, and then maintained the ownership gap after the mid 2000's. From the mid 2000's onwards, the difference in institutional ownership between polluter and non-polluters is more stationary, with the estimated *Polluterdummy* and yearly fixed effects interaction coefficients falling within the approximate band of -0.09 and -0.12. These estimates reveal that after the early broadening of the ownership gap between polluters and non-polluters, there has not been much of a further decline from 2005 onwards. Results are overall consistent with the hypothesis of growing environmental sentiment over the sample.

In Table 2.2, I show that institutional ownership of polluters is reduced relative to non-polluters. However, within the group of polluter firms, some firms may be perceived to generate enough positive economic value to offset the social externalities of pollution. Some firms may also operate in sectors that generate valuable output but cannot avoid polluting in their operations. In contrast, some inefficient polluting firms may pollute at high levels which cannot be justified by their low levels of positive economic output. I therefore test whether *inefficient* polluters, defined relative to their positive economic output, are associated with reduced institutional ownership.

I proxy for polluter efficiency by scaling yearly *Total Releases* by the latest annualised net sales of the firm in year *t*. Dividing *Total Releases* by net sales produces a ratio of negative to positive outputs. An inefficient polluter is defined as having a high ratio, while efficient polluters have a low ratio. I create a new dummy variable labelled *Scaledpolluter*, which is activated if a firm has a pollution to sales ratio that is within the top quintile for a year.[64] *Scaledpolluter* and *Polluterdummy* are positively correlated, with a correlation coefficient of 0.60 significant at the 1% level,[65] indicating that the largest absolute polluters are also likely to be the least efficient polluters. I rerun the yearly fixed effects panel regression *(2.1)* with *Scaledpolluter*; this regression estimates whether institutions hold fewer stocks of firms that are relatively inefficient in their toxic releases. A negative estimated coefficient for *Scaledpolluter* would imply that society discriminates against polluters after considering the economic value that these firms may otherwise generate. I also include *Polluterdummy* and industry fixed effects in two additional regression specifications. Results of the three regressions are presented in Table 2.3.

[64] I set *Scaledpolluter* to 1 for firms that have positive pollution in the numerator but a net sales value of 0 or less in the denominator; there are only two observations for which this is necessary. Alternatively, I find near exactly similar results if these observations are dropped instead.

[65] Pearson and Spearman coefficients are identical when estimating correlations between two dummy variables.

Table 2.3: Results of the polluter efficiency fixed effects panel regressions where the dependent variable is *IO*. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 8,954 firm-year observations in the sample for each specification. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Polluter efficiency panel regression results

Variable	(1)	(2)	(3)
Scaledpolluter	-0.0419**	-0.0261	-0.0152
	(-2.70)	(-1.34)	(-0.81)
Polluterdummy		-0.0254	-0.0176
		(-1.50)	(-0.89)
INDBETA	0.1040***	0.1029***	0.0307*
	(3.72)	(3.73)	(1.69)
LOGSIZE	0.0392***	0.0410***	0.0479***
	(5.60)	(6.20)	(6.03)
LOGBM	0.0134	0.0168	0.0453**
	(0.63)	(0.81)	(2.29)
STD	-0.0130*	-0.0128*	-0.0142**
	(-1.75)	(-1.74)	(-2.21)
PRINV	-0.1105**	-0.1089**	-0.0866*
	(-2.21)	(-2.19)	(-1.80)
RET	-0.0012	-0.0013	-0.0006
	(-1.19)	(-1.19)	(-0.63)
NASD	-0.0570***	-0.0570***	-0.0570***
	(-4.06)	(-4.06)	(-4.08)
SP500	-0.0261	-0.0249	-0.0290
	(-1.32)	(-1.27)	(-1.44)
Fixed effects	Year	Year	Year & Industry
N	8,954	8,954	8,954
Adjusted R^2	0.4070	0.4078	0.4508

Results indicate that the institutional ownership of inefficient polluters is also reduced, by a similar level as that of the greatest absolute polluters. When *Polluterdummy* is included as an explanatory variable in the regression, the estimated coefficients of both *Polluterdummy* and *Scaledpolluter* are negative, however neither has statistical significance. This is likely a result of the high correlation between these two variables; both explanatory variables are competing to explain the same variation in institutional ownership. Though results find a negative relationship between inefficient pollution and institutional ownership, they are unable to differentiate between the effects of absolute pollution and inefficient pollution on ownership. Similar to results shown in Table 2.2, I find no evidence to indicate that either *Polluterdummy* or *Scaledpolluter* are associated with reduced within-industry institutional ownership.

2.5.2. Disaggregated toxic releases test

An advantage of the Toxic Release Inventory is the data granularity; unlike environmental scores such as from KLD, the dataset breaks down the various types of releases by chemical group. By disaggregating *Total Releases* by chemical classification, I test which toxic substances have the greatest negative association with institutional ownership. This test serves to examine the varying effects of toxic chemical groups on institutional ownership.

Total Releases is disaggregated into one of three mutually exclusive chemical groups. The first classification consists of standard TRI chemicals, which comprise of general toxic chemicals such as certain forms of ammonia, aluminium, phosphorus and zinc. These chemicals may significantly damage human health, wildlife, and the external environment. The second category consists of persistent bio-accumulative chemicals (PBT) such as lead or mercury compounds, which accumulate in body tissue over time, cause lasting damage to the environment and are not easily destroyed. The final category of chemicals consists of separately identified persistent toxins, labelled as dioxin and dioxin-like compounds. These are trace level by-products of combustion or industrial processes. Dioxins are extremely toxic, and human exposure mostly occurs through food products.

TRI chemicals are generally released in larger quantities and are relatively less harmful than PBT and dioxin chemicals, which have lower reporting thresholds. PBT and dioxin chemicals have the potential to cause significant environmental and health damage from lower releases (EPA, 1999). In particular, dioxins have the greatest toxicity, and as a result have the lowest reporting thresholds measured in grams released. I hypothesise a negative relationship between institutional ownership and all chemical groups; however, I expect the strongest negative relationship to be with dioxins, and the weakest negative relationship with TRI chemicals.

I run the institutional ownership regression below with *TRI*, *PBT* and *Dioxin* as individual explanatory variables. These are set as either the continuous firm-year releases of their respective chemical classifications, or as dummy variables which are activated for the top yearly quintile of polluters for that chemical group.

$$IO_{i,t} = \beta^{TRI} * TRI_{i,t} + \beta^{PBT} * PBT_{i,t} + \beta^{Dioxin} * Dioxin_{i,t} + \boldsymbol{\beta^{control}} * \boldsymbol{CONT_{i,t}} + \varepsilon_{i,t} \qquad (2.2)$$

Continuous pollution variables are stored in millions of pounds, except for *Dioxin* releases which are stored in pounds. Data on *Dioxin* only begins from the year 2000 onwards and therefore reduces the panel size. The average values for the continuous *TRI*, *PBT* and *Dioxin* releases are 3.47, 0.26 and 0.03 respectively, and the Pearson correlations between the three variables range from 0.02 to 0.54. I use the full set of control variables in the primary ownership tests, and include yearly fixed effects. Results of the regression are presented in Table 2.4.

Table 2.4: Results of the institutional ownership fixed effects panel regression disaggregated by chemical classification, where the dependent variable is *IO*. *TRI*, *PBT* and *Dioxin* measure either the continuous toxic releases of their respective chemical classification, or are dummy variables activated for the top yearly quintile polluters of that subcategory. The continuous versions of *TRI* and *PBT* are measured in millions of pounds, while *Dioxin* is measured in pounds. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 4,778 firm-year observations in the sample. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Toxic releases disaggregated by chemical classification

Variable	Continuous	Dummy
TRI	-0.0002	-0.0305
	(-0.60)	(-1.36)
PBT	-0.0013	0.0097
	(-0.77)	(0.55)
Dioxin	-0.0424***	-0.0598**
	(-2.88)	(-1.99)
INDBETA	0.0968***	0.0786***
	(3.42)	(3.75)
LOGSIZE	0.0425***	0.0425***
	(5.16)	(5.99)
LOGBM	-0.0386	-0.0196
	(-1.61)	(-0.91)
STD	-0.0013	-0.0013
	(-0.14)	(-0.15)
PRINV	-0.1163**	-0.1115**
	(-2.04)	(-2.00)
RET	-0.0032**	-0.0030*
	(-2.05)	(-1.82)
NASD	-0.0756***	-0.0802***
	(-4.03)	(-4.30)
SP500	-0.0804***	-0.0770***
	(-3.04)	(-3.14)
Fixed effects	Year	Year
N	4,778	4,778
Adjusted R^2	0.2649	0.2735

Results reveal a negative association between institutional ownership and releases of all three chemical categories; however, only *Dioxin* is statistically significant. The magnitude of the *Dioxin* coefficient is much larger than that of the other two groups using either continuous or dummy variables. In an unreported test, I find that using the natural log of the continuous pollution variables as independent variables generates similar results. Results indicate that despite having the lowest average releases, the marginal impact of *Dioxin* releases on institutional ownership is significantly greater than that of the other two chemical groups.

2.5.3. Disaggregated ownership tests

While institutions in aggregate may be constrained in owning polluter stocks, some may be less sensitive to social norms. I hypothesise that institutions that face relatively less public scrutiny, have less social accountability, trade more aggressively, attempt to exploit market inefficiencies, and rely less on passive negative investment screens are less likely to be constrained by social norms and have relatively greater polluter stock ownership. These institutions may be less exposed to social pressures or may be willing to accept the trade-off between the potential pecuniary benefits derived from polluter equity and costs of deviating from social norms.

In the following two tests I break down the dependent variable *IO* by institution type; in the first test I separate institutions by their Bushee (2001) classification, while in the second test I disaggregate institutions based on their 13F classes. These tests reveal which institution groups are most associated with reduced ownership of polluter firms relative to their holdings of non-polluters.

I first use Bushee institutional investor classifications to disaggregate institutions into separate groups.[66] Institutions are grouped by Bushee (2001) using a cluster analysis on a set of factors that measure past characteristics of investment behaviour.[67] As described in Bushee (1998), these factors are generated from variables that measure the level of portfolio concentration and the degree of

---

[66] I thank Brian J. Bushee for access to his institutional investor database, which includes data on investor type and a permanent investor unique identifier. The classification cluster analysis is the same as described in Bushee (1998) but with the momentum variables excluded, as in Bushee & Noe (2000) and Bushee (2001). I avoid using the permanent investor type classification and instead opt for the dynamic classification to account for changes in institutional behaviour. Not all investors in the database are given a classification, and therefore some manager holdings observations are excluded from all subsamples. Investor classification data URL: *http://acct.wharton.upenn.edu/faculty/bushee/IIclass.html*.

[67] The classification approach of Bushee (2001) uses principle factor analysis to generate PTURN and BLOCK, which are measures of institutional portfolio turnover and diversification. A k-means cluster analysis is then run on these factor scores to separate institutions into the three Bushee groups. Bushee (2001) describes the mean portfolio characteristics of the three types of institutions as follows. Transient institutions have high portfolio turnover, diversified portfolios and a preference for smaller firms. Dedicated institutions have low turnover and low diversification. Quasi-indexer institutions have low turnover, diversified holdings, and a preference for larger firms.

portfolio turnover. Disaggregation by these factors allows for tests to estimate the association between firm pollution and ownership between different institutional profiles.

Institutions are segregated into three groups. Institutions are classified as either 'dedicated', 'quasi-indexer' or 'transient'. By construction, dedicated institutions have the highest portfolio concentration and low portfolio turnover (Bushee, 1998). Dedicated investors invest large amounts in a small number of firms and have a 'relationship' approach to their investments; dedicated investors are not frequent traders and have long-term stable equity holdings in relatively fewer firms (Bushee & Noe, 2000). Quasi-indexers also have low turnover, but tend to hold large, diversified portfolios, consistent with index-like, buy-and-hold, value strategies (Bushee, 1998). Quasi-indexers are the largest class of institutional investors. Both dedicated and quasi-indexers are characterised by long-term investment horizons. Lastly, transient institutions have the greatest portfolio turnover and exhibit relatively high diversification (Bushee, 1998). Transient institutions trade aggressively on short-term strategies and focus on generating short-term returns (Bushee, 2001). Transient ownership is also associated with future changes in stock price volatility (Bushee & Noe, 2000).

Bushee institution types have varied investment strategies that may impact their holdings of polluter stocks. Because of their short-term strategies and aggressive trading behaviour, I hypothesise that transient institutions are more likely to play the role of arbitrageur in the market and be relatively indifferent to social constraints on polluter ownership. Transient institutions may also be less exposed to public scrutiny as their investments are not held for extended periods of time. In contrast, quasi-indexers are anecdotally more likely to avoid polluter stocks for ethical and reputational reasons, and may use passive negative investment screens that reduce polluter ownership. Due to their long-term, value-driven, buy-and-hold strategies, quasi-indexers are expected to be more sensitive to public scrutiny and social norms that influence the long-run value of their investments. Dedicated institutions also have long-term investment horizons which may similarly influence their preference for non-polluting firms.

Each firm-year observation of *IO* is separated into three observations based on the total holdings of each Bushee institution subcategory. I conduct a separate yearly fixed effects panel regression for each subcategory of institutional ownership by regressing the disaggregated *IO* for each Bushee institution type *x* on *Polluterdummy*, as follows.

$$IO_{i,x,t} = \beta^{polluter} * Polluterdummy_{i,t} + \boldsymbol{\beta^{control}} * \boldsymbol{CONT_{i,t}} + \varepsilon_{i,t} \qquad (2.3)$$

I control for the full set of control variables from regression *(2.1)* in the vector $\boldsymbol{CONT}$, and two-way cluster standard errors by industry and year. I present the results of each of the three regression in Table 2.5.

Table 2.5: Results of the institutional ownership fixed effects panel regressions disaggregated by Bushee institution type, where the dependent variable is *IO* for that institution type. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 8,954 firm-year observations in the sample for each specification. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Institutional ownership disaggregated by Bushee institution type

	Bushee institution type		
Variable	Dedicated	Quasi-indexer	Transient
Polluterdummy	-0.0121***	-0.0601***	-0.0106*
	(-2.89)	(-5.98)	(-1.92)
INDBETA	0.0131***	0.0804***	0.0381***
	(2.80)	(3.34)	(3.94)
LOGSIZE	-0.0030*	0.0145***	0.0080***
	(-1.79)	(2.88)	(3.37)
LOGBM	-0.0003	-0.0107	-0.0117
	(-0.04)	(-0.51)	(-1.25)
STD	-0.0072***	-0.0211***	0.0005
	(-3.02)	(-4.85)	(0.22)
PRINV	-0.0131*	-0.0716*	-0.0379***
	(-1.75)	(-1.86)	(-3.33)
RET	-0.0001	-0.0014	0.0026***
	(-0.44)	(-1.55)	(4.42)
NASD	-0.0062	-0.0258*	-0.0099
	(-1.25)	(-1.83)	(-1.46)
SP500	0.0135*	0.0390*	-0.0145
	(1.85)	(1.81)	(-1.52)
Fixed effects	Year	Year	Year
N	8,954	8,954	8,954
Adjusted R^2	0.2082	0.4237	0.2652

Results reveal that all three Bushee institution types have significantly reduced ownership of polluters. Estimates show that quasi-indexers are the most averse to polluter stocks, followed by dedicated and then transient institutions. Though both dedicated and transient institutions have a similar estimated coefficient for *Polluterdummy*, the estimate is more significant for dedicated institutions. This reveals that transient institutions display relatively more dispersion in their investments in polluter stocks compared to dedicated investors. Quasi-indexers have both the greatest magnitude and statistical significance in their estimated reduced polluter ownership.

Some control variables, such as *LOGSIZE* and *RET*, change significance and sign dependent on the institutional group; this is expected due to the differences in trading strategies between institutional groups. As expected, the adjusted R^2 is much higher for quasi-indexers than transient investors, likely due to the diverse trading strategies used by the latter. Dedicated investors have the lowest adjusted R^2, implying greater variation in their relatively smaller breadth of niche investments.

Following Hong & Kacperczyk (2009) I again disaggregate institutions in the next test, but by 13F class. The 13F dataset classifies institutions into five types; type 1 are banks, type 2 are insurance companies, type 3 are mutual funds, type 4 are independent investment advisers which include hedge funds, and type 5 are other which include universities, endowments, and pension funds. I restrict my Thomson Reuters institutional holdings sample to 1987 - 1997 due to mapping issues in the data post 1997, in which many institutions are incorrectly stored as type 5 (Hong & Kacperczyk, 2009). I classify type 1, 2 and 5 as 'Type A' institutions, and type 3 and 4 as 'Type B' institutions. On average, Type A institutions are more likely to be socially constrained investors with greater visibility and public accountability, while Type B institutions are less likely to be constrained by social norms and may act as arbitrageurs if polluter stocks are ignored by other market participants (Hong & Kacperczyk, 2009). I hypothesise a negative coefficient on *Polluterdummy* for Type A institutions only, due to their greater sensitivity to social pressures. I repeat the disaggregated institutional ownership fixed effects panel regression *(2.3)* for both 13F institution groups. Results of the regressions are presented in Table 2.6.

Table 2.6: Results of the institutional ownership fixed effects panel regressions disaggregated by 13F institution type, where the dependent variable is *IO* for that institution type. 'Type A' consists of banks, insurance firms, pension plans, endowments, universities and employee-ownership plans. 'Type B' consists of mutual funds and independent investment advisers. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. Data is limited to the range 1987 to 1997. There are 3,471 firm-year observations in the sample. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Institutional ownership disaggregated by 13F institution type

	13F institution type	
Variable	Type A	Type B
Polluterdummy	-0.0352***	0.0665**
	(-3.46)	(2.49)
INDBETA	0.0186	0.0016
	(0.75)	(0.03)
LOGSIZE	0.0079*	0.0520***
	(1.65)	(3.99)
LOGBM	-0.0130	0.0775
	(-0.94)	(1.50)
STD	-0.0178***	-0.0005
	(-4.50)	(-0.03)
PRINV	-0.0204	0.2176**
	(-0.95)	(2.19)
RET	-0.0016**	0.0017
	(-2.18)	(0.74)
NASD	-0.0220*	0.0171
	(-1.75)	(0.35)
SP500	0.0622***	-0.1191***
	(3.33)	(-2.66)
Fixed effects	Year	Year
N	3,471	3,471
Adjusted R^2	0.1954	0.0503

The estimated coefficients reveal that on average, Type A institutions have reduced levels of polluter ownership while Type B institutions have relatively more; coefficients are estimated at the 1% and 5% significance levels respectively. Due to the lack of clean data post 1997, the effects of more recent environmental awareness and anti-polluter sentiment are not seen in the results. Despite this, the estimated coefficients themselves are in line with expectations. Type A investors are generally less aggressive in their trading strategy and more prone to public scrutiny, whereas Type B investors are expected to experience less public scrutiny, try arbitrage price inefficiencies, and be indifferent to social norms. Type B investors also have more varied trading strategies within the group, supported by a dramatically lower adjusted R^2 compared to Type A. In this reduced sample, results provide evidence of polluter stocks being ignored by publicly scrutinised institutional investors, and in contrast are invested in relatively more by aggressive institutions.

2.5.4. Analyst coverage

Following the logic of Hong & Kacperczyk (2009), analyst coverage of sin stocks should be reduced due to the relationship between sell-side analysts and institutional investors; if institutions are reluctant to own polluter stocks, there will be reduced demand for coverage of polluters. I test whether analyst coverage is negatively associated with total firm pollution with the following fixed effects panel regression.

$$LOGCOV_{i,t} = \beta^{polluter} * Polluterdummy_{i,t} + \boldsymbol{\beta^{control}} * \boldsymbol{CONT_{i,t}} + \varepsilon_{i,t} \qquad (2.4)$$

LOGCOV is regressed against the full set of control variables from institutional ownership regressions and the *Polluterdummy* variable. I first employ a yearly fixed effects model to test the association between analyst coverage and pollution using *Polluterdummy* and *Total Releases* separately. I also add industry fixed effects in the *Polluterdummy* model to test for within-industry relationships between pollution and institutional ownership. In my final two model specifications, I use an independent trend variable *t* instead of fixed effects, and then include a polluter-time interaction effect between *Polluterdummy* and *t*.

I control for the same variables used in ownership regressions (Hong & Kacperczyk, 2009), represented by the vector ***CONT***. T-stats are calculated using two-way clustered standard errors, with clustering on year and industry. I present the results of all five regression specifications in Table 2.7.

Table 2.7: Results of the analyst coverage panel regressions where the dependent variable is *LOGCOV*. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 8,954 firm-year observations in the sample for each specification. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Analyst coverage panel regression results

Variable	(1)	(2)	(3)	(4)	(5)
Polluterdummy	-0.1412**		-0.1329	-0.1391**	-0.2462*
	(-2.05)		(-1.44)	(-2.08)	(-1.76)
Total Releases		-0.0012			
		(-0.51)			
t				-0.0065	-0.0081**
				(-1.55)	(-2.25)
*Polluterdummy * t*					0.0074
					(1.12)
INDBETA	0.2834***	0.3022***	0.0928	0.3282***	0.3345***
	(6.47)	(6.02)	(1.55)	(6.85)	(6.57)
LOGSIZE	0.3306***	0.3250***	0.3605***	0.3269***	0.3279***
	(10.51)	(9.76)	(12.55)	(10.47)	(10.66)
LOGBM	-0.0677	-0.0863	0.0027	-0.0048	-0.0058
	(-0.61)	(-0.75)	(0.02)	(-0.05)	(-0.05)
STD	0.0220*	0.0212	0.0225**	0.0129	0.0129
	(1.67)	(1.57)	(2.08)	(1.05)	(1.06)
PRINV	-0.0357	-0.0441	0.0100	0.0063	0.0064
	(-0.44)	(-0.53)	(0.14)	(0.09)	(0.09)
RET	-0.0296***	-0.0298***	-0.0297***	-0.0256***	-0.0255***
	(-8.21)	(-8.32)	(-8.59)	(-8.19)	(-8.21)
NASD	0.0624	0.0701	0.0360	0.0582	0.0586
	(0.78)	(0.87)	(0.44)	(0.72)	(0.72)
SP500	0.2576**	0.2515**	0.2128**	0.2827***	0.2833***
	(2.42)	(2.35)	(2.05)	(2.64)	(2.64)
Fixed effects	Year	Year	Year & Industry	None	None
N	8,954	8,954	8,954	8,954	8,954
Adjusted R^2	0.4174	0.4154	0.4562	0.4122	0.4126

Results are consistent with the hypothesised relationship between polluter firms and analyst coverage. I find that after controlling for various firm level variables, firms that pollute in the top quintile in a year have reduced analyst coverage on average. Estimates generated with using the linear trend and yearly fixed effects in columns (1) and (4) are similar and significant at the 5% level.

Like estimates from regression *(2.1)*, the estimated coefficient of *Total Releases* is insignificant, suggesting a non-linear relationship between analyst coverage and pollution levels. The polluter-time interaction coefficient is also insignificant, providing no evidence that analyst coverage for polluters is diverging from that of other stocks over time.

The estimated coefficient of *Polluterdummy* is statistically insignificant when industry fixed effects are included, providing no evidence of a within-industry relationship between polluters and analyst coverage, similar to results from the primary institutional ownership test. I conduct a separate regression in the appendix by replacing the industry fixed effects for firm fixed effects, but again find no evidence to suggest there is a within-firm relationship between pollution and analyst coverage. These findings are consistent with the institutional ownership results and reinforce the notion that societal discrimination does not appear to specifically target pollution levels, but instead appear to penalise polluting industries or firms. As a result, within-industry and within-firm relationships between pollution, institutional ownership and analyst coverage are weak.

Though average *LOGCOV* is increasing with time in the sample, the regression coefficient of the linear trend is negative; this is a result of the marginal effect of time after controlling for other factors such as firm size.[68] The coefficient of *LOGSIZE* is positive and estimated with significance, and along with the coefficients of *INDBETA* and *SP500*, reveals that firms that are large, operate in market-sensitive industries, or are listed on the S&P 500 receive greater analyst coverage. The coefficient of *RET* is also estimated with significance, however, has a negative sign indicating that firms that have more analyst coverage generate lower returns. Overall estimates are consistent with Hong & Kacperczyk (2009) and support the hypothesis that polluters have reduced analyst coverage, albeit with weaker significance than in ownership tests.

[68] The Pearson correlation coefficient between *LOGSIZE* and the trend is 0.30, and is statistically significant at the 1% level. This is not high enough to create severe multicollinearity problems in regressions; however, if *LOGSIZE* is removed from the vector of control variables, the estimated sign of the trend coefficient reverses.

In an additional test, I repeat the yearly fixed effects regression *(2.4)* with *Sindummy*, presented in the appendix. Estimated coefficients for *Sindummy* are more negative than those of *Polluterdummy*, but are estimated without statistical significance.

To illustrate the relationship between pollution and analyst coverage, I recreate Figure 2.4 except with analyst coverage as the dependent variable. Specifically, I regress *LOGCOV* on 9 pollution decile dummies based on yearly rankings of *Total Releases*, using yearly fixed effects and the same control variables from Table 2.7. I graph the estimated polluter decile coefficients in Figure 2.6.

Figure 2.6: Pollution decile dummy coefficients, estimated from a yearly fixed effects panel regression with *LOGCOV* as the dependent variable. Decile 1 has no active dummy, and is therefore the benchmark against which the following dummy coefficients are compared against.

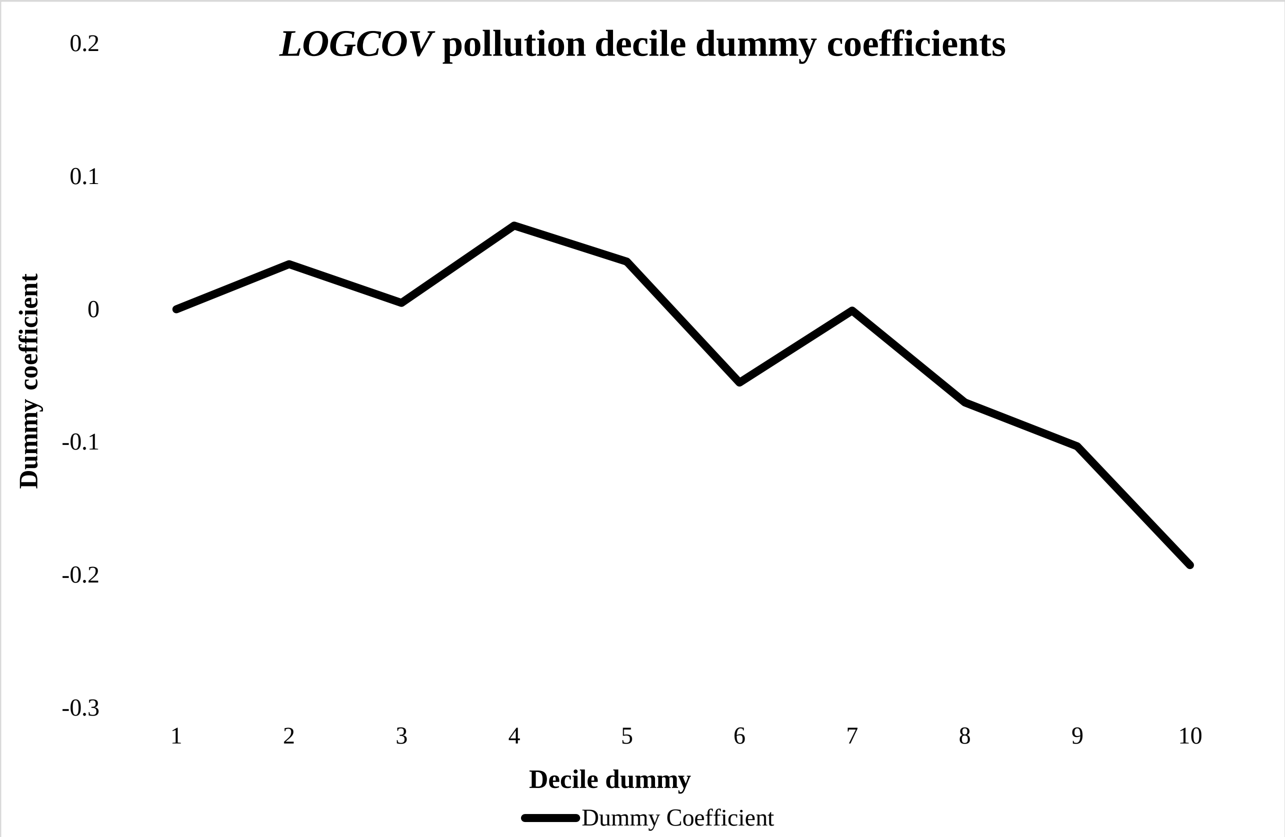

Like Figure 2.4, Figure 2.6 depicts a non-monotonic relationship between the LOGCOV and pollution levels on average. Unlike Figure 2.4, however, Figure 2.6 portrays a relatively flatter relationship, without an abrupt reduction in analyst attention at the higher pollution deciles. This is consistent with a weaker significance of *Polluterdummy* coefficient estimates from Table 2.7 compared to the primary estimates obtained from the main ownership regressions in Table 2.2.

Consistent with tests on *IO*, repeating regression *(2.4)* with 5% and 1% pollution percentile dummies also generates negative but insignificant coefficient estimates, suggesting that extreme polluters in excess of *Polluterdummy* do not have a significant further reduction in analyst coverage.

2.6. Additional tests

In auxiliary tests I examine the institutional churn of polluter stocks, and test a long-short trading strategy derived from the shunned-stock hypothesis. I also search for evidence of reverse causality in the hypothesised relationship between firm pollution and institutional ownership.

2.6.1. Investor churn

Institutional ownership of polluting firms may be influenced specifically by institutional investor horizons. I test whether polluter firms are owned by investors with relatively short-term investment horizons. In previous results I find that quasi-indexers, which are associated with long-term investment horizons, are more reluctant to own polluter stocks compared to transient institutions, which are associated with aggressive short-term trading strategies. However, unlike the disaggregated ownership tests, this additional test serves to explicitly estimate the relationship between institutional investment horizon and polluter stocks. This test also indirectly examines which institutions hold less polluter stocks relative to each other, as opposed to the disaggregated institutional tests which instead examine which investors own less polluter stocks relative to their other holdings.

I hypothesise that institutions with short-term horizons are more likely to exhibit arbitrage strategies for sin stocks, and exploit mispricing from the shunned-stock effect. Due to the potential for tail event regulatory shocks to cash flows, institutions may also be reluctant to hold polluter stocks over long periods, and instead use these stocks in a ‘hot-potato’ momentum strategy. If true, these channels will inflate the quarterly trading churn of polluter stocks, ceteris paribus, which will be disproportionately held by institutions with relatively short-term investment horizons.

I test the short-horizon investor hypothesis by aggregating average investor horizons at the firm level. I use a churn variable to proxy for investor horizons as a function of their trading activity.

Following Gaspar, Massa, & Matos (2005) I first generate an investor churn variable to measure average institutional investment horizons.

$$IChurn_{j,q} = \frac{\sum_{i\in I}\left|Shares_{i,j,q} * Price_{i,q} - hares_{i,j,q-1} * Price_{i,q}\right|}{\sum_{i\in I}\left(Shares_{i,j,q} * Price_{i,q} + Shares_{i,j,q-1} * Price_{i,q-1}\right)/2} \qquad (2.5)$$

IChurn$_{j,q}$ is a weighted average measure of the turnover of institution j at quarter q. *Shares*$_{i,j,q}$ represents institution j's holdings of firm i's shares at the end of quarter q. *Price*$_{i,q}$ represents the price of firm i's shares at the end of quarter t. The ratio is bounded by 0, with a higher ratio indicating a greater turnover of holdings. I then aggregate the *IChurn* ratio at the firm-quarter level with the following equation.

$$Firmchurn_{i,q} = \frac{\sum_{j\in J}\left(IChurn_{j,q} * Shares_{i,j,q}\right)}{\sum_{j\in J} Shares_{i,j,q}} \qquad (2.6)$$

Firmchurn$_{i,q}$ is a weighted average measure of the institutional churn of the shares of firm i at the end of quarter q, with weightings proportional to the number of shares held by institution j as a percentage of total shares of firm i held by institutions. Finally, I take a yearly average of *Firmchurn* for each firm to use as the dependent variable in tests.

$$Firmchurn_{i,t} = \frac{\sum_{q=1}^{4} Firmchurn_{i,q,t}}{4} \qquad (2.7)$$

In a fixed effects panel model, I regress *Firmchurn* on *Polluterdummy* and a set of control variables.

$$Firmchurn_{i,t} = \beta^{polluter} * Polluterdummy_{i,t} + \boldsymbol{\beta^{control}} * \boldsymbol{CONT_{i,t}} + \varepsilon_{\boldsymbol{i,t}} \qquad (2.8)$$

I control for the firm level independent variables used in ownership regressions, represented by the vector ***CONT***. I also include *DIVYIELD*$_{i,t}$ as an additional control variable (Starks et al., 2017), set as the annual dividend yield for firm i during year t.[69] To control for the impacts of yearly business

[69] *DIVYIELD*$_{i,t}$ is generated by dividing the total dividends paid by firm i in year t by the closing share price at the end of the year. Starks et al. (2017) also include *TURNOVER*$_{i,t}$ as an explanatory variable, which is the average monthly stock turnover ratio in year t, calculated by dividing the monthly trading volume of stock i by the shares outstanding at

cycle fluctuations and industry averages, I incorporate yearly and industry fixed effects. I adjust standard errors with two-way clustering on Fama-French industry and year.[70] Results of the regression are presented in Table 2.8.

Table 2.8: Results of the investor horizons fixed effects panel regression, where the dependent variable is *Firmchurn*. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 7,732 firm-year observations in the sample. Compared to main tests, the sample is smaller because only stocks that have institutional ownership greater than 0 and have unique Bushey permanent key identifiers are included. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Investor horizons panel regression results

Variable	(1)	(2)
Polluterdummy	0.0103**	0.0120***
	(2.24)	(2.88)
INDBETA	0.0085	0.0020
	(1.62)	(0.38)
LOGSIZE	0.0019	0.0014
	(0.96)	(0.61)
LOGBM	-0.0151*	-0.0123
	(-1.66)	(-1.47)
STD	0.0108***	0.0094***
	(4.54)	(4.28)
PRINV	-0.0283**	-0.0244**
	(-2.27)	(-2.17)
RET	0.0016**	0.0017***
	(2.58)	(2.77)
DIVYIELD	0.0082***	0.0077***
	(4.25)	(3.85)
NASD	-0.0027	-0.0041
	(-0.50)	(-0.69)
SP500	-0.0365***	-0.0363***
	(-7.85)	(-7.08)
Fixed effects	Year	Year & Industry
N	7,732	7,732
Adjusted R^2	0.8002	0.8048

the end of the month. I omit *TURNOVER* from the set of control variables due to potential simultaneity concerns; however I find a positive coefficient for *Polluterdummy* that is significant at the 10% level if *TURNOVER* is used as the dependent variable instead of *Firmchurn* in the following regression.

[70] I also find that the *Polluterdummy* coefficient is significant if only industry fixed effects are included.

Results for the *Polluterdummy* coefficient estimate is consistent with both a priori expectations and the findings of Starks et al. (2017). Polluter firms have a significantly higher weighted average churn from their institutional investors compared to non-polluters, indicating that institutions investing in polluter stocks generally have shorter investment horizons and higher holdings turnover on average. The estimated *Polluterdummy* coefficient remains significant within both regressions, suggesting that firm pollution is associated with both within-year and within-industry higher ownership from short-horizon investors. Results are consistent with the hypothesised relation between polluter sin stocks and short-term investor horizons.

2.6.2. Polluter portfolio returns

Hong & Kacperczyk (2009) find that a portfolio of sin stocks outperforms comparables. Following the prior findings of reduced institutional ownership and analyst coverage of polluters, similar to traditional sin stocks, I test whether polluter stocks also outperform due to the shunned-stock effect. If arbitrage of mispricing for sin stocks is truly limited due to societal discrimination (Akerlof, 1980), polluter stocks should generate abnormal returns over time.

I test whether polluter stocks outperform in the market on average using a portfolio approach. I create a long-short *Toxicportfolio* that is long the stocks with the highest yearly quintile of toxic releases and short the remaining stocks, and is rebalanced every month.[71] I use both equal-weights and value-weights for robustness.[72]

Inconsistent with expectations, I find that the equal-weighted *Toxicportfolio* generates -0.138% average returns a month with a Newey-West t-stat of -1.25, while the value-weighted *Toxicportfolio* generates -0.115% average returns a month with a Newey-West t-stat of -0.91. The time series of the equal-weighted *Toxicportfolio* monthly returns are illustrated in Figure 2.7.

[71] Portfolio sorting and weights are calculated based on ex-ante 1-month lagged information.

[72] I exclude micro-cap stocks from the equal-weighted portfolio. Micro-caps are defined as the shares of firms with very low market capitalisations of less than $250m (Lins, Servaes, & Tamayo, 2017).

Figure 2.7: The time series of monthly holding period returns of the long-short equal-weighted *Toxicportfolio* over the sample period.

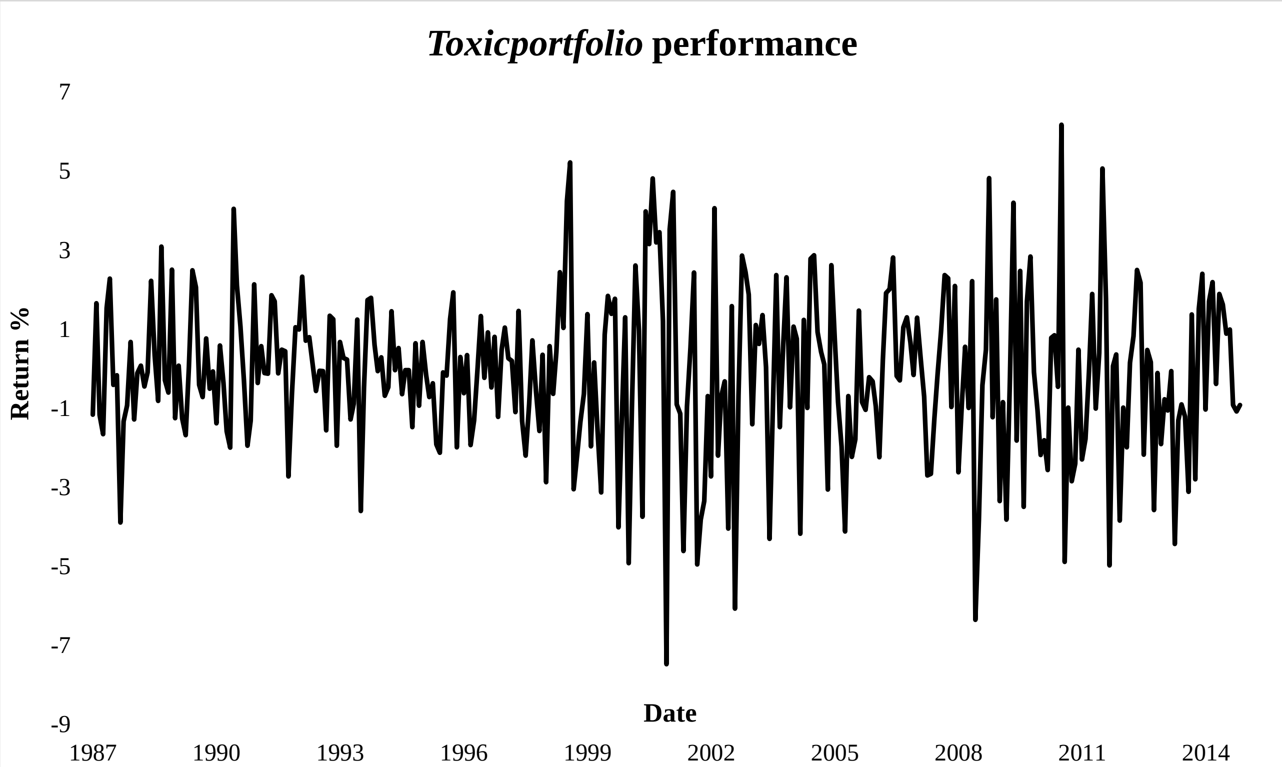

The strategy returns are volatile and show no clear evidence of positive abnormal performance. The equal-weighted *Toxicportfolio* returns become more volatile from 2000 onwards; I find that this is also true for the value-weighted portfolio. This volatility increase may indicate an increased difficulty in the valuation of polluters relative to non-polluters. Alternatively, *Toxicportfolio* volatility may be driven by the volatility of the risk factors the strategy is exposed to.

I regress the returns of *Toxicportfolio* against three popular benchmark risk models in the following time series regression to test for abnormal returns.

$$R_t = \alpha + \boldsymbol{\beta}^{\boldsymbol{control}} * \boldsymbol{CONT_t} + \varepsilon_t \qquad (2.9)$$

The returns of *Toxicportfolio*, R_t, are regressed against three sets of benchmark models which consist of the CAPM, the Fama-French 3-factors and the Carhart 4-factors. The variable of interest is α, which is a measure of abnormal returns generated by the *Toxicportfolio* in excess of the benchmark.

I adjust standard errors using Newey-West corrections for 5-month lags.[73] Results of the portfolio regression are shown in Table 2.9. I repeat the regression for both equal-weighted and value-weighted *Toxicportfolio* returns.

Table 2.9: Regression results for both equal-weighted and value-weighted long-short *Toxicportfolio* returns. I present the abnormal return estimates along with factor sensitivities to the benchmark models. Standard errors are presented in brackets below. Standard errors are adjusted for Newey-West 5-month lags. There are 335 monthly observations. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

***Toxicportfolio* regression results**

Equal-weighted	CAPM	FF 3-factors	Carhart 4-factors
α	-0.073	-0.093	-0.139
	(-0.62)	(-0.88)	(-1.22)
MKT	-0.105**	-0.062	-0.047
	(-2.48)	(-1.54)	(-1.15)
SMB		-0.204***	-0.210***
		(-4.89)	(-5.49)
HML		0.063	0.082
		(1.15)	(1.51)
MOM			0.058
			(1.71)
N	335	335	335
Adjusted R^2	0.0503	0.1667	0.1809

Value-weighted	CAPM	FF 3-factors	Carhart 4-factors
α	0.000	-0.037	-0.065
	(0.00)	(-0.33)	(-0.60)
MKT	-0.188***	-0.160***	-0.150***
	(-5.17)	(-4.00)	(-3.86)
SMB		-0.068	-0.072
		(-1.33)	(-1.36)
HML		0.112	0.124*
		(1.42)	(1.69)
MOM			0.037
			(1.03)
N	335	335	335
Adjusted R^2	0.1076	0.1321	0.1338

Results provide no evidence of polluter outperformance on average. All three benchmark models provide insignificant estimates of abnormal returns for both portfolios. Contrary to the hypothesis,

[73] Following the literature I set the lag equal to $4(T/100)^a$ where T = 335 time periods and a = 4/25 using the quadratic spectral kernel. The output equals 4.85, which I round up to 5.

estimates of portfolio alpha are mostly negative for both equal-weighted and value-weighted portfolios. The equal-weighted portfolio loads negatively and significantly on the market and SMB factors. These negative factor loadings indicate that polluter firms have lower market betas and behave more like large-cap firms. This is both intuitive and consistent with the summary statistics presented in Table 2.1. On average, polluter firms in the sample have lower market sensitivity and are larger in size. The market beta of the equal-weighted portfolio becomes statistically insignificant as additional risk factors are added. The value-weighted portfolio has a stronger negative loading on the market factor that survives as additional risk factors are included, but has an insignificantly negative size loading and a marginally significant value loading when using the Carhart 4-factor model only.

Results are overall inconsistent with the shunned-stock hypothesis. This may be due to weak limits to arbitrage of polluter stocks; as Angel & Rivoli (1997) suggest, in order to observe a material pricing effect there must be a large restriction on discriminated stocks. There is no evidence to suggest that Type B investors from the previous disaggregated ownership tests are constrained in their investment in polluter firms; similarly, transient and dedicated institutions are not as reluctant to hold polluter stocks as quasi-indexers. These types of investors may contribute to the lack of outperformance of polluter stocks through their arbitrage efforts.

2.6.3. Robustness reverse causality tests

I consider potential reverse causality in my model. Reverse causality may pose problems of simultaneity in models that estimate the relationship between institutional ownership and firm pollution. One might argue that it is not firm pollution that drives institutional ownership, but instead institutional ownership that affects toxic releases through institutional oversight and pressures that may encourage a firm to adopt greener policies; however, this is unlikely to invalidate results for the following reasons. There is no obvious rationale for why aggregate institutional ownership, as a proportion of total ownership, should be correlated with abatement pressures. Aggregate institutional pressures on polluters depend on constituent institutions, which may individually promote or discourage pollution based on their own incentives. Furthermore, there may also be pressure from retail or public sector owners to reduce pollution, mitigating reverse causal effects. Due to the high dollar and time costs required to develop efficiency in pollution, potential institutional pressure to reduce toxic releases are also more likely to occur over a longer time-frame compared to a contemporaneous year. *IO* is measured at year-end, whereas *Total Releases* is based on pollution throughout the year, further weakening the likelihood of the ownership variable having a causal relationship with the firm pollution measure used in tests.

In a robustness check of the reverse causality hypothesis, I use a panel VAR (PVAR), a change-on-change analysis, and a natural experiment to test for any evidence of *IO* influencing *Total Releases*.

I first estimate the following PVAR model to test for simultaneity in my main results. To account for non-stationarity, I model the first order differences of all variables. The PVAR simultaneously estimates the effects of lagged changes in institutional ownership on changes in pollution and vice versa, whilst also controlling for exogenous variables and autocorrelations in dependent variable.

$$\boldsymbol{Y_{i,t}} = \boldsymbol{\alpha} + \boldsymbol{\beta_1} * \boldsymbol{Y_{i,t-1}} + \boldsymbol{\beta_2} * \boldsymbol{Y_{i,t-2}} + \boldsymbol{\beta_3} * \boldsymbol{Y_{i,t-3}} + \boldsymbol{\beta_4} * \boldsymbol{X_{i,t}} + \boldsymbol{\varepsilon_t} \qquad (2.10)$$

The matrix $\boldsymbol{Y_{i,t}}$ consists of the variables $\Delta IO_{i,t}$ and $\Delta Total\ Releases_{i,t}$, while the exogenous $\boldsymbol{X_{i,t}}$ consists of the first order differences in control variables used in ownership tests. I drop $\Delta NASD$ from the set of control variables as it does not vary within the sample. I control for 3 lags of dependent variables in the model. If reverse causality is present, coefficients of ΔIO or its lags should be estimated with statistical significance, and the impulse response function (IRF) should show statistically significant effects of shocks in ΔIO on $\Delta Total\ Releases$. I illustrate the results of the PVAR and corresponding IRF's in the following tables and figures.

Table 2.10: PVAR regression results testing for reverse causality in the institutional ownership regressions. I use STATA code for the PVAR developed in Abrigo & Love's (2015) working paper. Dependent variables are ΔIO and Δ*Total Releases*, with 3 lags. The exogeneous Δ*NASD* dummy is omitted from the regression due to no variation in the reduced sample of 5,542 observations. I present PVAR coefficient estimates with t-statistics in brackets below. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

PVAR estimates

Independent variables	Dependent variables	
	ΔIO	Δ*Total Releases*
$\Delta IO_{(-1)}$	-0.1284***	0.3352
	(-5.22)	(0.88)
$\Delta IO_{(-2)}$	0.0150	0.2351
	(0.81)	(0.55)
$\Delta IO_{(-3)}$	0.0168	0.6003
	(0.75)	(1.55)
Δ*Total Releases*$_{(-1)}$	-0.0004**	0.1186
	(-2.00)	(0.55)
Δ*Total Releases*$_{(-2)}$	0.0000	-0.0817
	(0.27)	(-0.54)
Δ*Total Releases*$_{(-3)}$	-0.0001	0.1015
	(-1.44)	(1.53)
Δ*INDBETA*	0.0975***	0.5025
	(6.74)	(0.61)
Δ*LOGSIZE*	0.0838***	-0.0288
	(10.06)	(-0.15)
Δ*LOGBM*	0.0052	-0.1655
	(0.30)	(-0.57)
Δ*STD*	-0.0028*	-0.0268
	(-1.85)	(-0.50)
Δ*PRINV*	0.0084	0.0458
	(0.52)	(0.15)
Δ*RET*	-0.0034***	-0.0012
	(-7.33)	(-0.08)
Δ*SP500*	0.0124	-0.5823
	(0.90)	(-1.37)
N	5,542	

Table 2.11: Impulse response results of the PVAR model in tabulated format.

Impulse response function		
Response variable and forecast horizon	**Impulse variables**	
ΔIO	ΔIO	Δ*Total Releases*
0	1	0
1	-0.1284	-0.0004
2	0.0314	0.0000
3	0.0108	-0.0001
4	-0.0033	0.0000
5	0.0011	0.0000
6	0.0000	0.0000
7	-0.0001	0.0000
8	0.0000	0.0000
9	0.0000	0.0000
10	0.0000	0.0000
Δ*Total Releases*		
0	0	1
1	0.3352	0.1186
2	0.2318	-0.0678
3	0.5808	0.0837
4	0.0179	0.0272
5	-0.0016	-0.0105
6	0.0633	0.0050
7	0.0077	0.0042
8	-0.0038	-0.0010
9	0.0053	0.0000
10	0.0017	0.0005

Figure 2.8: Impulse response function for the PVAR model. The impulse variable is ΔIO while the response variable is Δ*Total Releases*. The dark bands around the impulse response estimate represent 95% confidence intervals generated with bootstrapped standard errors from 1000 random draws.

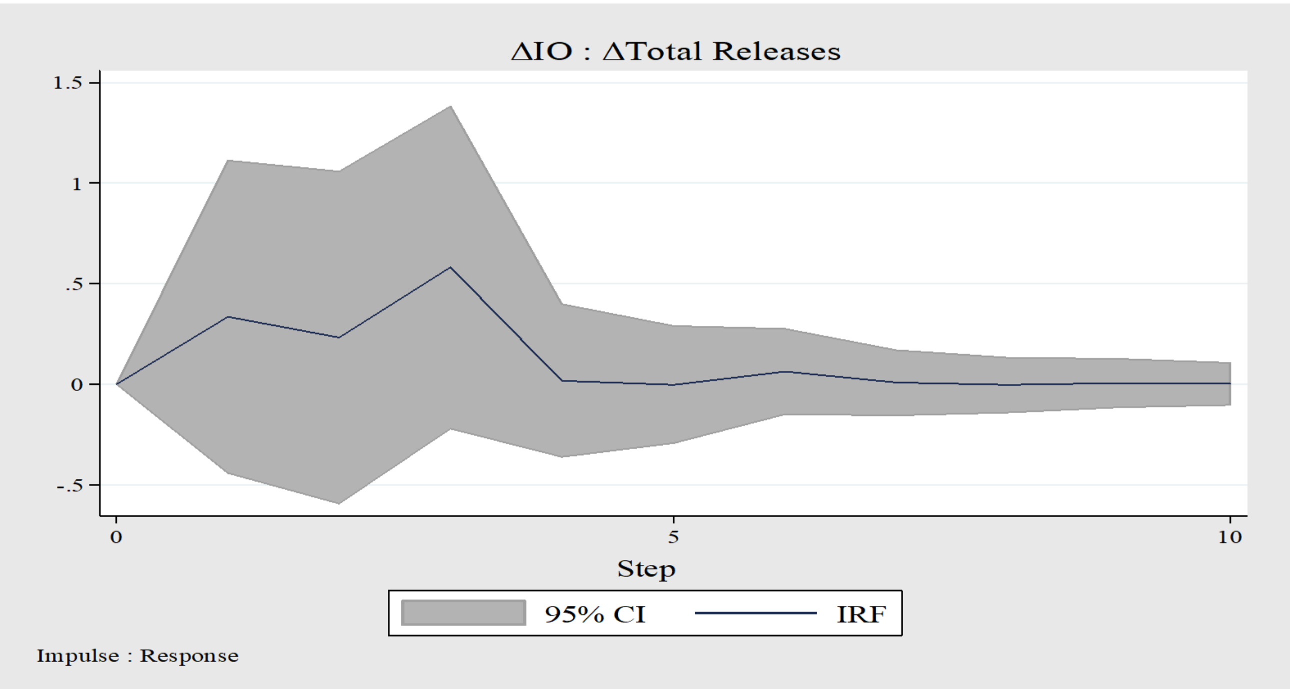

Figure 2.9: Impulse response function for the PVAR model. The impulse variable is Δ*Total Releases* while the response variable is Δ*IO*. The dark bands around the impulse response estimate represent 95% confidence intervals generated with bootstrapped standard errors from 1000 random draws.

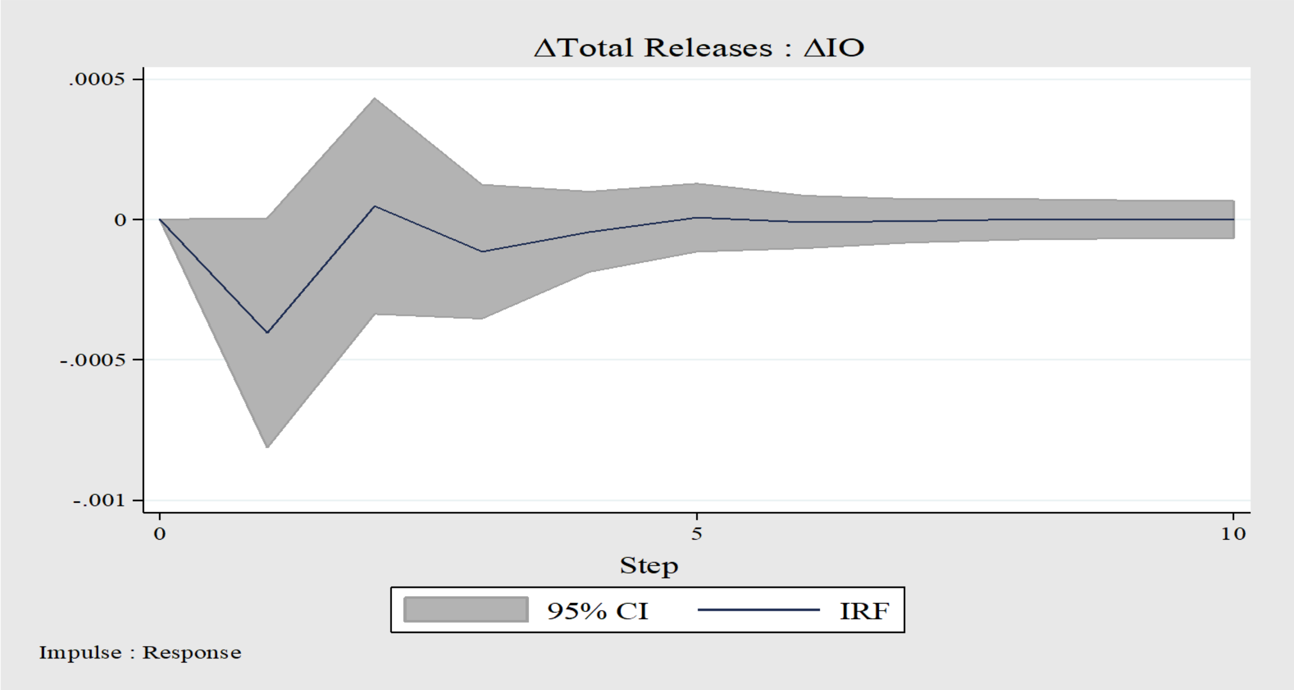

Results provide evidence that Δ*IO* is positively affected by one-year lagged values of Δ*Total Releases*, suggesting the presence of short-run causality consistent with my primary hypothesis. More importantly, there is no evidence of a relation between Δ*Total Releases* and lagged values of Δ*IO*. The t-stat of the estimated effect of one-year lagged Δ*Total Releases* on Δ*IO* is much higher than vice versa. Furthermore, Figure 2.8's IRF graph illustrates that the 95% confidence interval contains 0, and therefore provides no evidence of a causal short-term effect of Δ*IO* shocks on Δ*Total Releases*. For robustness I conduct the same test using only the subsample of firms that have a value of 1 for *Polluterdummy*, and again find no evidence of reverse causality.

I additionally use a change-on-change analysis to test for Granger causality and estimate whether lagged changes in *IO* lead to changes in *Total Releases.*[74] I conduct this test with the following specification.

[74] For example, see Aggarwal, Erel, Ferreira, & Matos (2011), Chhaochharia, Kumar, & Niessen-Ruenzi (2012) or Kim et al. (2019).

$$\Delta Y_{i,t} = \alpha + \beta^{x} * \Delta X_{i,t-1} + \boldsymbol{\beta^{cont}} * \Delta \boldsymbol{CONT_{i,t-1}} + \varepsilon_{i,t} \qquad (2.11)$$

The change in the dependent variable from *t-1* to *t* is regressed on the change in the independent variable and vector of control variables from *t-2* to *t-1*. I run this regression twice, first with Δ*IO* as the dependent variable and Δ*Total Releases* as the independent variable, and then vice versa. I also include yearly fixed effects to control for time-varying heterogeneities. I present results of the two regressions in Table 2.12.

Table 2.12: Change-on-change analysis fixed effect panel regression results. Δ*IO* is regressed against lagged Δ*Total Releases*, and Δ*Total Releases* is regressed against lagged Δ*IO*. I control for lagged changes in prior ownership variables. The exogeneous Δ*NASD* dummy is omitted from the regression due to no variation in the reduced sample of 7,363 firm-year observations. I present coefficient estimates with t-statistics in brackets below. Standard errors are calculated with clustering on industry and year. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Change-on-change analysis

Independent variables	Dependent variables	
	ΔIO_t	$\Delta Total\ Releases_t$
$\Delta Total\ Releases_{t-1}$	-0.0001**	
	(2.59)	
ΔIO_{t-1}		0.1240
		(0.98)
$\Delta INDBETA_{t-1}$	0.0061	0.8711
	(0.38)	(1.01)
$\Delta LOGSIZE_{t-1}$	-0.0032	-0.2216
	(-0.68)	(-0.97)
$\Delta LOGBM_{t-1}$	-0.0072	0.2318*
	(-0.53)	(1.72)
ΔSTD_{t-1}	0.0003	-0.0010
	(0.12)	(-0.01)
$\Delta PRINV_{t-1}$	-0.0389**	-0.4454
	(-2.63)	(-0.76)
ΔRET_{t-1}	0.0010*	0.0063
	(1.98)	(0.52)
$\Delta SP500_{t-1}$	-0.0043	-0.1857
	(-0.48)	(-1.40)
Fixed effects	Year	Year
N	7,363	7,363
Adjusted R^2	0.1155	0.0011

Results again do not provide any evidence of a relationship between lagged changes in ownership and current changes in firm pollution. While the negative effect of lagged changes in pollution on contemporaneous changes in institutional ownership is estimated with significance at the 5% level, estimates of the reciprocal relationship are insignificant. For increased robustness, I also repeat this test with the subsample of firms with a value of 1 for *Polluterdummy*, but find a consistent lack of results.

In a final robustness test, I use a natural experiment which is assumed to cause independent variation in *IO*. Specifically, I examine whether firms that have been added to or removed from the S&P 500 index have significant changes in their pollution. I hypothesise that inclusion or exclusion from the index affects the institutional ownership of the stock, and test for a corresponding effect on *Total Releases* to provide evidence of reverse causality. Filtering by firms that have had a change in the value of *SP500* from the previous year, I generate a sample of 55 firms that have been included in the index in the previous year and 22 firms that have been removed as index constituents in the previous year, for a total of 77 firm-year observations. I then compare the changes in institutional ownership to changes in pollution levels with the following yearly fixed effects panel regression.

$$\Delta Total\ Releases_{i,t} = \beta * \Delta IO_{i,t} + \varepsilon_{i,t} \qquad (2.12)$$

I regress contemporaneous changes in pollution levels on the changes in institutional ownership following this exogeneous shock on the S&P 500 index constituent status. If there is a causal relationship between *IO* and *Total Releases*, I expect a statistically significant negative coefficient estimate. In order to capture time heterogeneities I include yearly fixed effects, and use two-way clustered standard errors on year and industry.[75] I present the results of this regression in Table 2.13, along with the average changes in *IO* and *Total Releases* following inclusion or exclusion from the S&P 500.

[75] I find similar non-results if I include industry fixed effects or drop fixed effects from the panel altogether.

Table 2.13: S&P 500 constituent change analysis. I limit the sample to firms that have been recently included or excluded as constituents of the S&P 500 index in the prior year, grouped as recent inclusions, exclusions or both. I present the mean ΔIO and Δ*Total Releases* for the year following the status change. In the bottom half of the table are regression coefficients generated by regressing Δ*Total Releases* on contemporaneous ΔIO with yearly fixed effects. I present coefficient estimates with t-statistics in brackets below. Standard errors are calculated with clustering on industry and year. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

S&P 500 constituent change analysis

	Recently included	Recently excluded	Total sample
Mean ΔIO_t	0.0394^{**}	-0.0103	0.0252^{*}
	(2.36)	(-0.33)	(1.72)
Mean Δ*Total Releases*$_t$	-0.3160	-0.1741	-0.2754
	(-0.61)	(-0.59)	(-0.73)
ΔIO_t	-6.3445	1.0342	-1.1323
	(-1.25)	(0.69)	(-0.40)
N	55	22	77
Fixed effects	Year	Year	Year

Following shocks to firm S&P 500 constituent status, I find no evidence of any effect of a change in institutional ownership on firm pollution levels. This test provides no evidence of reverse causality between the variables. The estimated coefficient sign is negative for only two out of the three samples, and in the ‘recently excluded’ subsample the sign is positive. All three estimated regression coefficients are insignificant. In an unreported test, I find that replacing the dependent variable with any one or the sum of Δ*Total Releases*$_{t+1,\ t+2,\ t+3}$ also generates insignificant results.

Overall, robustness tests provide no evidence of a relationship flowing from *IO* to *Total Releases*; none of the results highlight institutional ownership affecting firm toxic releases.

2.7. Conclusion

I argue that institutional investors are constrained through social norms, which limits institutional investment in polluter stocks. Social norms are more likely to constrain institutions due to their large public profiles which are more exposed to public scrutiny than individual investors, who are more easily able to keep their positions in sin stocks out of the public eye. Overall, results reveal that polluter stocks are shunned by institutional investors, like the sin stocks of tobacco, alcohol and gambling (Hong & Kacperczyk, 2009). Results also reveal that while institutional equity in firms is increasing over the sample period, ownership of polluter firms is increasing at a significantly slower rate. I attribute these findings to growing environmental sentiment, which include concerns relating to human health and environmental damage. When disaggregated by Bushee (2001) groups, all three institution groups are found to have reduced holdings of polluter stocks at varying levels. Similarly, disaggregation by 13F institution type reveals that institutions more likely to be constrained by society also have reduced polluter ownership, while institutions with relatively more opaque or aggressive strategies have increased ownership. Following findings of reduced institutional ownership, I also find that polluter stocks receive relatively less analyst coverage.

Auxiliary tests reveal that polluter stocks are held by investors with shorter investment horizons, measured by the quarterly churn in their holdings. However, tests do not provide evidence of either underperformance or outperformance of polluter stocks.

The results of this study are consistent with the theory that society shuns environmental sin stocks as a reaction to their costs on social welfare; ownership of polluter stocks generates disutility or costs from association that exceed their benefits.

Studies on the impacts of environmentalism on corporate finance are gaining traction. Further research could incorporate a higher frequency of pollution and investor trading data to estimate causal drivers in investor decision making, or alternatively examine potential window dressing in holdings disclosures. The reduced institutional ownership of polluter firms implies that the remainder of the stocks are held by either retail investors, the public sector or insiders; these ownership channels should also be investigated. Further research could also study how the physical consequences of pollution (i.e. smog, health problems, toxic spills) affect the decision making of investors that suffer the consequences. Finally, the role of regulation is crucial to assessing the performance of both polluter and green stocks; research into environmental-regulatory risk channels and expected stock returns is a developing field of asset pricing with scope for further study.

Chapter 3

Toxic Expectations: Analyst Forecasts and Firm Pollution

3.1. Introduction

The impact of environmental performance on economic value is a developing field in the economic literature.[76] Firms that release toxic emissions are exposed to future regulatory and compliance costs, as well as potential consumer backlash.

Pollution regulation is primarily designed to reduce negative externalities through internalisation mechanics. Abatement and compliance activities are expensive by nature. Abatement often involves large investments and running costs to sustain reduced emissions, while compliance costs may be borne by polluters to accurately measure and monitor emissions. If polluters fail to meet regulations, they may be taxed or fined. Polluters are thus exposed to increasing abatement and compliance costs, which reduce their expected earnings relative to comparable firms. Polluter revenues are also sensitive to consumers switching away from 'dirty production' (Fraj & Martinez, 2007). Consumers may become more concerned about environmental degradation and display preferences towards environmentally friendly products (Do Paço, Raposo, & Leal Filho, 2009). As consumers demand 'green' products, polluters face pressure to reduce emissions in their production processes; failing to do so decreases revenues and earnings. Polluters are thus exposed to earnings losses driven by shifting consumer demand.

Predicting the future consequences of pollution on earnings is a complicated task which may be influenced by behavioural biases present in the forecaster. This research examines security analysts' views on the effects of pollution on firm profitability. Analysts play an important role in financial markets by collecting, processing and forecasting economic information which is ultimately used in setting expectations and pricing assets. Analyst forecasts provide a dimensionality of market expectations from the viewpoint of professional estimators of firm value. Examining pollution from the outlook of analysts provides insight into how industry experts believe toxic emissions interact

[76] For example, see Horváthová (2010), Fisher-Vanden & Thorburn (2011), Guenster, Bauer, Derwall, & Koedijk (2011), Berkman, Jona, & Soderstrom (2019), or El Ghoul, Guedhami, Kim, & Park (2018).

with future profitability, and how accurate these beliefs are. This study has implications on the valuation of pollution costs and polluter securities in the private market, and adds to the environmental debate by assessing whether the analyst-predicted consequences of pollution on earnings are rational.

The contribution of this study is threefold. Firstly, I introduce firm pollution as a variable that is associated with forecast errors through behavioural biases. Polluter firms differ from regular firms as they are exposed to potential regulatory costs and consumer backlash, however these effects may be mis-estimated. I show that investment professionals set systematically pessimistic forecasts for firms with poor environmental performance. Secondly, this research provides a novel example of behavioural factors constraining firm valuation experts by promoting inaccuracy in their forecasts of polluters. This is of relevance to the literature, as it provides additional empirical examples of longstanding behavioural theories. Lastly, I take both a macro and micro approach by examining both the average bias in all forecasts and the bias of individual analysts, and test whether polluter bias is profitable through a trading strategy.

The rational expectations hypothesis assumes that reality only diverges randomly from forecasts. Behavioural finance argues that individuals may deviate from setting rational expectations. As illustrated in various (and sometimes contradictory) behavioural arguments, individuals might be influenced by cognitive constraints. Security analysts often display bounded rationality; the literature has found evidence of systematic analyst irrationalities and subsequent forecast errors.[77] Analyst irrationalities can be explained through behavioural biases, which include analysts overweighting low-probability shocks (Kahneman & Tversky, 1979), analysts being biased through the availability of dramatic and recent information (Tversky & Kahneman, 1974), or analysts failing to generate true conditional expectations (Tversky & Kahneman, 1981). These biases all imply that analysts may

[77] For example, see De Bondt & Thaler (1990), De Bondt & Forbes (1999), Amir & Ganzach (1998), Park & Sabourian (2011), or Fujiwara, Ichiue, Nakazono, & Shigemi (2013).

overshoot in their expectations of firm costs of pollution, resulting in systematic pessimism for polluter firms.[78]

The primary hypothesis of this study examines whether the aggregate forecasts made by analysts for polluting firms are associated with pessimism, as each behavioural channel suggests. I test this hypothesis using firm pollution data from the Toxic Release Inventory (TRI). Results indicate that in aggregate, analyst forecasts have a significantly pessimistic bias for polluters, ceteris paribus. After controlling for other factors associated with forecast bias and accuracy, analyst forecasts tend to undershoot actual polluter annual earnings on average. I find evidence of analyst pessimism for firms that generate high levels of absolute pollution. However, if pollution is scaled by sales, the estimated pessimism loses statistical significance. This suggests that analyst biases are more influenced by information on total pollution, which is more easily noticeable compared to sophisticated information on pollution efficiency. Furthermore, I find that the primary category of chemicals listed in the TRI are the most strongly associated with analyst pessimism. The primary category of chemicals are generally less toxic than the other categories, but are usually released in larger quantities and are again likely to be more noticeable and attract bias.

In a secondary hypothesis, I examine whether forecast errors made by individual analysts are persistent as suggested by the conservatism bias, which proposes that individuals may be slow to update their beliefs. Consistent with the hypothesis, I find that ex-ante identified green (grey) analysts are consistently pessimistic (optimistic) in their forecasts for polluters ex-post. I examine how the bias by analyst type changes for polluter forecasts with shorter horizons; results indicate that both green and grey analysts become more pessimistic closer to the earnings date.

[78] In this study, forecast bias is defined as the directional errors in analyst forecasts when compared to actual realised earnings. Forecast bias is therefore a measure of the signed miscalculation made by an analyst in their predictions of firm earnings. A green analyst is defined as one who is pessimistic about polluters, while a grey analyst is defined as one who is optimistic about polluters.

Lastly, I test whether the pessimistic forecasts for polluters leads to positive returns around earnings announcements. If investors form earnings expectations which are closely tied to analyst forecasts, systematic pessimism for polluting firms should generate systematic positive earnings surprises. Results provide no evidence in favour of this hypothesis; firm pollution is not significantly associated with abnormal returns around either annual or quarterly earnings announcements.

3.2. Literature review

I review the literature by first examining studies on the psychology of security analysts, which include theories on behavioural constraints and cognitive biases in analysts' forecasts for polluter firms. I then review studies on the various sources of analyst bias and identify variables that have been associated with forecast errors. Lastly, I explore prior studies on the relationships between analyst forecast errors, earnings surprises and stock returns.

3.2.1. Security analyst psychology

The psychology of security analysts is important as it comprises of the behavioural factors which drive forecasts and forecast errors. The task of security analysts is often complex, involving the appraisal of various interlinked business activities. Due to the vast amounts of available data and limited processing capacity, decisions may be reduced to simple heuristic procedures that result in bias (Wright, 1980). De Bondt & Thaler (1985) find that investors tend to overreact when incorporating dramatic variables into their forecasts; evidence of stock price reversion is consistent with market overreactions (De Bondt & Thaler, 1987). Analyst forecasts are important inputs for the formation of market expectations and may introduce bias. Furthermore, cognitive biases may persist if individuals are unaware that they are influenced by them (Wright, 1980).

As highlighted under prospect theory (Kahneman & Tversky, 1979), financial forecasts may be biased by attributing disproportionately large probability weights to events that are eye-catching but unlikely to occur, whilst underweighting more probable events. Because low probability events are infrequently encountered, individuals often have little experience in dealing with them (Burns, Chiu,

& Wu, 2010). Burns et al. (2010) give examples of heuristics through which the overweighting of low probability events may occur; these include anchoring on prior ignorance and sorting outcomes into coarse probabilistic groups. Anchoring on prior ignorance occurs when there are a limited number of future outcomes which are not understood well. Forecasters may begin by assigning each outcome approximately equal probabilities of occurring, and then insufficiently adjust the probabilities as they collect more information, and hence be 'anchored' to their initial estimates. Coarse probabilistic grouping occurs when the finer categories of the probability interval, representing unlikely events, are clumped together with more likely outcomes.

The representativeness heuristic occurs when individuals overreact to the 'law of small numbers' and fail to account for true probability distributions (Ritter, 2003). Representativeness is broken down into two subset biases, known as the availability bias and base rate fallacy. The availability bias theory suggests that individuals find it easier to recall and process information around past events that are dramatic, recent or resonate with the individual. As a result, this information is overly weighted when making judgements (Tversky & Kahneman, 1974). Availability of information in recent memory affects forecasts when analysts give too much importance to memorable events. Analysts may also suffer from a cognitive bias known as the base rate fallacy (Tversky & Kahneman, 1981). The base rate fallacy theory suggests that individuals may fail to account for conditional probability when analysing data.

Amir & Ganzach (1998) find that analysts display characteristics that are consistent with the representativeness heuristic. Under these cognitive constraints, analysts choose extreme prediction values that match the extremity of the predictive information, which leads to an overestimation of effect. Amir & Ganzach (1998) argue that analysts are likely to use salient information as an anchor, from which new information is adjusted.

Analysts may also exhibit the conservatism bias, in which individuals are slow to update their beliefs over time despite being proven wrong in the past (Ritter, 2003). Analysts may anchor on their previous forecasts, exhibiting tendencies of belief perseverance. Fujiwara, Ichiue, Nakazono, & Shigemi (2013) provide evidence of analyst stubbornness, such that analyst forecasts are anchored to prior predictions and slow to change with new information.

3.2.2. Analyst forecast bias

While research on security analysts initially focused on the viability of forecasts as a proxy for market expectations compared to time series models, the literature has moved on to assessing sources of analyst bias and inaccuracy (Bradshaw, 2011).

Prior empirical studies have examined whether analysts are effective at analysing information and generating forecasts. The literature finds that security analysts tend to generate forecasts which are systematically optimistic when compared to actual results.[79] Under the hypothesis of analyst rationality, forecast errors should not be predictable. De Bondt & Thaler (1990) find evidence that on average, forecasts are optimistic when compared to actual results; De Bondt & Thaler (1990) attribute this to the bounded rationality of security analysts.

Systematic optimism is attributed to a general analyst underreaction to prior optimistic forecast errors and information on stock prices and fundamentals (Abarbanell, 1991; Mendenhall, 1991; Ali, Klein, & Rosenfeld, 1992). Easterwood & Nutt (1999) find that a source of analyst optimism and inaccuracy is a general overreaction (underreaction) to positive (negative) information. Easterwood & Nutt's (1999) study does not explicitly consider what causes analysts to overreact (underreact) to positive (negative) information. From a behavioural perspective, it may be that analysts are incentivised to be optimistic, exhibiting confirmatory bias around new information to reach their desired conclusion.

McNichols & O'Brien (1997), Lin & McNichols (1998), and Hong & Kubik (2003) suggest that analysts avoid reporting negative forecasts for firms that are investment banking clients of their firm.

[79] For example, see O'Brien (1988), Butler & Lang (1991), Brous (1992), Brous & Kini (1993), Kang, O'Brien, & Sivaramakrishnan (1994), Duru & Reeb (2002), Beckers, Steliaros, & Thomson (2004), Ding, Charoenwong, & Seetoh (2004), Irvine (2004), Marsden, Veeraraghavan, & Ye (2008), So (2013).

Francis & Philbrick (1993) and Lim (2001) argue that analyst optimism results from analyst incentives to maintain access to managerial information. Motivations for systematic analyst optimism also include the incentives to generate commissions through trading clientele (Hayes, 1998; Irvine, 2004; Groysberg, Healy, & Maber, 2011).

O'Brien (1988) finds that the most recent analyst forecasts have the greatest forecast accuracy, consistent with the notion that later forecasts incorporate the most up-to-date information. Barron, Byard, & Liang (2013) provide an alternative explanation, arguing that pessimistic analysts are unlikely to issue early forecasts; early forecasts generated from the same public information tend to be more optimistic due to managerial influence and analyst self-selection. Barron et al. (2013) explain that analysts may accommodate managerial pressures for late, pessimistic forecasts which generate easier hurdles to beat. Also, early analyst pessimism is less profitable; the difficulties in shorting potentially underperforming stock constrain brokerage profits and disincentivise early pessimism.

Bradshaw (2011) argues that the general optimism in forecasts is due to a sample bias resulting from a conflict of interest. Analysts with pessimistic views may be reluctant to issue negative recommendations if they feel that this may damage their relationships with the target firm, and ultimately their own business.[80]

3.2.3. Variables linked to forecast error

This study aims to isolate the effects of pollution on analyst forecasted earnings and irrationality. Over time, the literature has identified specific variables associated with analyst bias and accuracy.

The literature finds that longer forecast horizons are associated with forecast bias and inaccuracy (O'Brien 1988; Richardson, Teoh, & Wysocki, 1999; Burgstahler & Eames, 2003; Eames & Kim, 2012; Barron et al., 2013; Hutira, 2016). Firm size has been documented to influence forecast errors

[80] Following Bradshaw's (2011) logic, the results of this study should not be extrapolated as the effect of pollution on the opinions and errors of *all* analysts, but rather only those who issue forecasts.

(Brown, 1998); larger firms usually have less biased forecasts through extensive coverage, greater scrutiny, stable earnings and more comprehensive disclosure by management (Hutira, 2016).

Alford & Berger (1999) suggest that each analyst contributes new information to the consensus estimate which eventually reduces inaccuracy. Merkley, Michaely, & Pacelli (2017) find complementary empirical evidence suggesting a negative relationship between analyst coverage and forecast bias. Mikhail, Walther, & Willis (1997, 2003) show that analysts with more experience in researching a firm produce more accurate forecasts; experienced analysts better interpret changes in earnings information, contributing to their increased accuracy.

The direction and magnitude of earnings forecasts are also associated with analyst accuracy. Analysts that issue negative forecasts risk their long-term relationship with the management of the firm and may lose this source of information (Hutira, 2016). Das, Levine, & Sivaramakrishnan (1998) argue that analysts may issue strong forecasts to coax information from the management of firms which are difficult to forecast for. Hayes (1998) argues that stocks that are expected to perform well generally have lower forecast errors, as analysts pay greater attention to these stocks. Brown (1998), among others, finds that analysts are systematically optimistic for firms that ex-post report losses. Both Hutira (2016) and Capstaff & Paudyal Rees (1998) find lower forecast accuracy for firms that report lower earnings than the previous year.

Brown (1997) finds that certain industries, including oil and gas extraction and primary metals, have poorer analyst forecast accuracy on average. Similarly, Coën, Desfleurs, & L'Her (2009) provide evidence of industry effects on forecast errors. Ciccone (2005) finds evidence of forecast errors decreasing over time, while Espahbodi, Espahbodi, & Espahbodi (2015) show that regulations aimed at aligning analyst incentives in the early 2000's improved forecast accuracy over the short-term. These studies conclude that forecast errors have structural variation between industries and time.

3.2.4. Earnings surprises and stock returns

Brown & Rozeff (1978) argue that analyst forecasts are a good proxy for market expectations. In an efficient market with rational expectations, earnings expectations should be based on the best available forecasts. To remain employed, analysts must generate value by providing more accurate forecasts than predictive models. Brown & Rozeff (1978) find evidence in favour of this, suggesting that analysts generate value as both acquirers and processors of information. Fried & Givoly (1982) find that analyst forecast errors are more correlated with security price movements than errors from predictive models, suggesting a strong link between analyst and market expectations.

3.3. Hypothesis

Based on the behavioural theories of over-weighting tail event probabilities, the availability bias, and the base rate fallacy, I primarily hypothesise that toxic emissions are associated with a negative bias in analyst forecasts. By sorting on prior ignorance, analysts may expect polluters earnings to fall through regulatory or environmental campaigns, and initially assign a high probability to this state, failing to sufficiently reduce probabilities with new information. By creating coarse probabilistic groups for polluters, the unlikely but disastrous costs for polluters may be grouped with more likely outcomes, leading to underestimated earnings. Due to the availability bias and recent trends in environmental finance,[81] analysts may pay too much attention to historical firm costs of pollution or anti-polluter stigma, and incorrectly extrapolate it into the future. Lastly, as pollution is usually negatively framed, through the base rate fallacy analysts may perceive of polluters as unpopular and heavily regulated firms, failing to notice the true costs of pollution across firms in the contemporaneous economic environment. I avoid testing any singular behavioural theory, and instead test whether the average observed polluter bias is indeed pessimistic as suggested by each channel.

[81] For example, Flammer (2013) shows that various stakeholders have exerted increasing pressures over time on firms to be environmentally responsible.

I additionally hypothesise that the most biased analysts display conservatism, in which they anchor on previous forecasts and generate persistent forecast errors. As a result, individual analysts that have displayed prior bias towards polluters are expected to continue to be biased in the same direction, despite having been proven wrong in their previous predictions. However, when forecasts are made nearer to the earnings date, there is less probability for tail events and more information available. I also test whether biased analysts walk down their extreme forecasts over time.

The final hypothesis is that of positive abnormal returns around polluter earnings announcements. Polluters may generate abnormal returns around earnings announcements if analysts are found to be pessimistically biased in predicting their earnings, as expected under the primary hypothesis. If analyst forecast errors are tied to market earnings surprises, and polluters have systematically pessimistic forecasts, polluters are expected to experience abnormally positive returns around earnings announcements.

3.4. Data and research methodology

I focus my methodology at the analyst-level due to the behavioural motivation of polluter pessimism. The final unbalanced panel dataset is of firm, analyst and month dimensionality.

3.4.1. Polluter data

I use the Toxic Release Inventory database to source data on toxic emissions at the firm-year level. The TRI program is managed by the Environmental Protection Agency (EPA), containing data on the releases of toxins within the contiguous U.S. that are believed to harm the natural environment or human wellbeing. The chemicals covered by the TRI are associated with cancer, chronic and acute health effects, and environmental damages. Firm disclosure to the program is a mandatory and audited program covering over 500 different chemical types and over 50,000 industrial facilities. Firms must disclose their toxic releases if they employ 10 or more full-time employees, operate in specific pollution prone industries, and handle or manufacture listed chemicals above specified thresholds.

Data ranges from 1987 to 2017.[82] TRI emissions data has been used by groups that include regulators, the media, and environmental activists (Hamilton, 1995).[83]

At the firm-year level, I sum on-site, off-site and transferred releases to generate *Toxic Releases*, measured in billions of pounds. I use contemporaneous releases and assume that analysts are aware of the level of a firm's pollution activity. The role of a security analyst often requires in-depth research of target firms; therefore, it is reasonable to assume that analysts are aware of the polluting activity of their research targets. The primary limitation of this assumption is that estimated contemporaneous relationships between *Toxic Releases* and analyst forecast bias are associative and not predictive.[84]

3.4.2. Analyst forecast data

I obtain data on analyst forecasted earnings per share (EPS) from the Institutional Brokers Estimate System (IBES) dataset. IBES stores estimates from sell-side brokerage institutions which often employ multiple analysts. I source data on actual EPS through the IBES actuals file.

In this study, the term forecast date refers to the point in time that an individual analyst announced their predicted EPS for a firm, while earnings date refers to the point in time for which the analyst believes the firm will achieve their predicted EPS. Forecast horizon refers to the difference in time between the forecast and earnings dates. The earnings announcement date is the date on which the firm's management make the realised earnings as at the earnings date public information.

Using the detail historical file on IBES, I source data on all analyst forecasts of EPS for every firm in the database, in monthly frequency. I focus on EPS forecasts for the next financial year-end.

---

[82] One limitation of the TRI is that data is self-reported, however measurement error is mitigated through audits run by the EPA. The data also focuses on manufacturing industries. For greater discussion of limitations, see Kim et al. (2019).

[83] KLD scores provide a popular alternative of firm-level data on environmental performance. I opt for TRI emissions over KLD scores as the latter are discrete, clustered with limited cross-sectional variation, and subjectively measured. TRI emissions also allow for more granular analysis.

[84] An alternative to this assumption is to match *Toxic Releases* to analyst forecasts after the TRI has released data. However, this often occurs around the second calendar year after the pollutants are released, resulting in stale data. However, for robustness, I rerun tests on the primary hypothesis using one-year and two-year lagged *Toxic Releases*; tests generate similar estimates. The one-year lagged pollution coefficient is significant at the 5% level, while the two-year lagged pollution coefficient is only marginally insignificant. Robustness regressions are presented in the appendix.

Some analysts have multiple forecasts for the same firm and earnings date in the same month; I store the last of such forecasts to capture individual analysts' latest revisions for that month.

I collapse the data into an unbalanced panel of firm, analyst, and month dimensions. The panel includes monthly forecast updates if they exist.[85] Forecast observations are interpreted as the monthly expectations of a firm's annualised earnings per share as set by an individual analyst.

3.4.3. Fundamentals and market data

I obtain data on firm fundamentals from Compustat. I use annual fundamentals over quarterly as the former are not subject to intra-year seasonality, and match the periodicity of annual EPS.

I source data on firm institutional ownership from the Thomson Reuters Institutional 13F Holdings database.[86]

I obtain stock market returns data from the CRSP database. I exclude data on non-domestic equities, and follow Shumway (1997) in adjusting for delisting biases. Returns are stored in percentage format.

I source abnormal returns around earnings announcements from Event Study by Wharton Research Data Services (WRDS). WRDS provides risk factor adjusted abnormal returns for equities around specified event dates. Using the U.S. daily event study file, I obtain stock abnormal returns around earnings announcement dates, benchmarked to the CAPM and Carhart 4-factors (Carhart, 1997). Abnormal returns are calculated as the error terms of the estimated benchmark model. Factor exposure is estimated using the previous 365 days of return history from 10 days prior to the earnings announcement. Firms with less than 70 days of return history in this window are excluded.

[85] Some analysts have irregular forecast updates. Monthly observations exist when any analyst has newly announced their forecast for a firm. Months with no new forecasts have no observations. I find similar results if the sample is restricted to monthly forecasts that have at least 2 other competing forecasts by other analysts.

[86] The database contains a number of issues that are highlighted in Geertsema (2014). I follow his methodology in addressing these issues.

3.4.4. Variables and summary statistics

I match data on analyst forecasts, firm fundamentals and firm pollution within the unbalanced panel, and generate the following dependent and independent variables.

I generate *FERROR* as a function of *FEPS* and *Actual*, the forecasted earnings per share and actual earnings per share respectively.

$$FERROR_{i,j,t} = \frac{Actual_{j,h} - FEPS_{i,j,t,h}}{|Actual_{j,h}|} \tag{3.1}$$

FERROR is calculated as the difference between *Actual* and analyst *i*'s *FEPS* for firm *j*, as at forecast date *t*, for earnings date *h*, scaled by *Actual*. *FERROR* captures the signed error of analyst expectations of earnings, relative to ex-post realised earnings. A positive (negative) value indicates a pessimistic (optimistic) outlook of firm earnings. *FERROR* is a measure of directional bias in analyst forecasts of earnings.

Very low absolute values of *Actual* inflate *FERROR*; when *Actual* is 0, *FERROR* cannot be calculated at all. To account for this, I winsorize *FERROR* within the range of ±2. Observations with *Actual* values of 0 are also set at ±2 dependent on whether the forecast was optimistic or pessimistic.[87] The maximum absolute value of 2 for *FERROR* generously limits analyst forecast errors to 200% of actual earnings. The winsorized outlier observations account for approximately 3.7% of the entire sample. I present a histogram of *FERROR* in Figure 3.1.

[87] Observations that have both 0 values for *FEPS* and *Actual* have *FERROR* set to 0 in main tests as well as in winsorization robustness tests.

Figure 3.1: Histogram illustrating the percentage distribution of *FERROR*, the measure of analyst forecast bias. Extreme observations with absolute values greater than 2 have been winsorized in creating this sample. These extreme observations represent approximately 3.7% of the total sample. There are 425,621 observations in the sample.

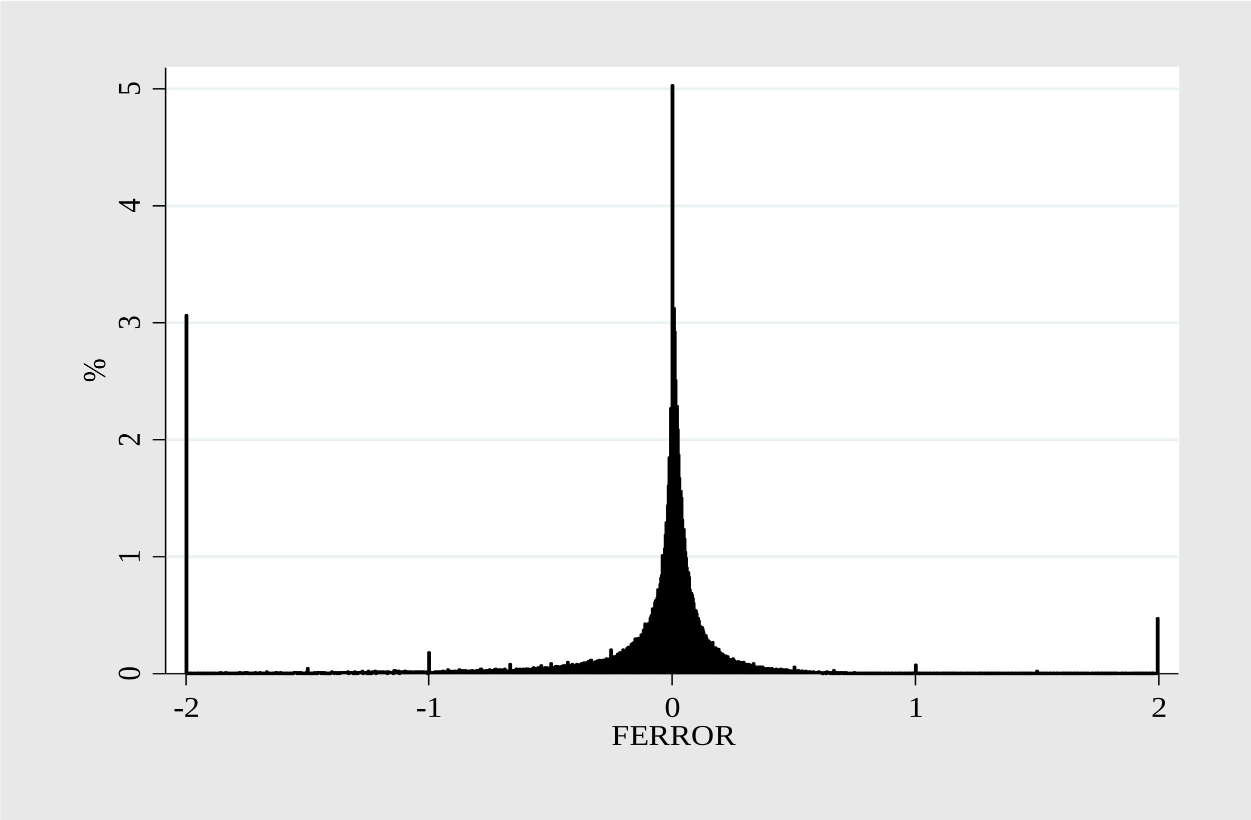

Figure 3.1 reveals that forecast errors are approximately 0 on average. However, there are instances where forecast errors are 200% of actual earnings per share; the winsorization results in frequency spikes at these values. Prior to winsorization, extreme forecast errors are simply diminishing extensions to the tails and are infrequent at individual points. The histogram illustrates slight spikes in *FERROR* frequency at positive and negative values of 0.5, 1, and 1.5.

For robustness, I recreate *FERROR* using alternative deflators in the denominator of equation *(3.1)*. I find that these alternative measures of *FERROR* produce estimates of polluter pessimism similar to primary results. I present the results of these robustness tests in the appendix.

Toxic Releases is the independent variable of interest. *Toxic Releases* is the total amount of pollutants released by a firm in a calendar year, measured in billions of pounds. I present the cumulative distribution of *Toxic Releases* at the firm-year level in Figure 3.2, illustrating the toxic emissions emitted by most firms, along with the maximum emissions observed in the sample.

Figure 3.2: Cumulative distribution function of *Toxic Releases* at the firm-year level. *Toxic Releases* are represented by the x-axis, while the cumulative proportion of observations that have equal or lower values of *Toxic Releases* are represented by the y-axis.

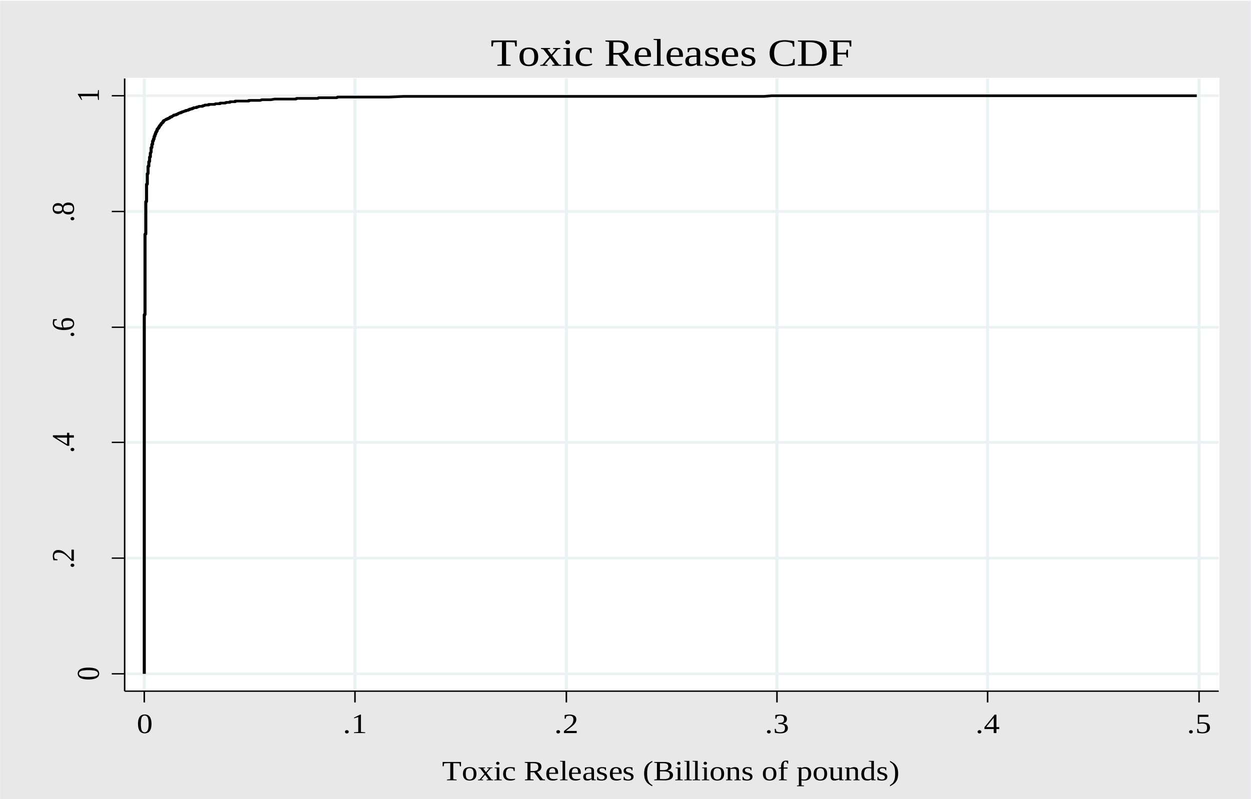

The cumulative distribution graph illustrates that most observations of *Toxic Releases* range from 0 to 0.1 billion pounds; however, there are some observations of emissions as high as 0.5 billion pounds. There are 591 firm-years with 0 *Toxic Releases*; when expanded to the forecast-month level, observations of firms that do not pollute at all are approximately 4.8% of the full sample. The firm-year Pearson autocorrelation coefficients with one-year and two-year lags of *Toxic Releases* are 0.97 and 0.93 respectively, both significant at the 1% level. Correlations support the assumption of analyst awareness of contemporaneous firm pollution; estimates imply that toxic releases are strongly predictable over time.

In primary tests, I control for firm, analyst, and earnings-related independent variables, motivated by prior research as highlighted in the literature review.

Firm controls include *LOGSIZE*, set as the natural log of 1 plus the market capitalisation of the firm as at the beginning of the month. *LOGSIZE* controls for the effects of firm size on forecast errors (Brown, 1998; Hutira, 2016). I control for book-to-market value using *LOGBM*, measured as the natural log of 1 plus the most recent book value of equity divided by market capitalisation, as at month start. I control for firm leverage with *LEV*, set as the most recent book value of debt divided by book value of total assets as at the beginning of the month.[88]

Analyst independent variables include *FPERIOD*, measured as the gap between the forecast date and earnings date in hundreds of days (Richardson et al., 1999; Burgstahler & Eames, 2003; Eames & Kim, 2012; Barron et al., 2013; Hutira, 2016). *COV* is the number of analysts that cover firm *i* in month *t* (Alford & Berger, 1999; Merkley et al., 2017). *SPREAD* is the total number of firms covered by the forecasting analyst in the month. The more firms that are researched by an analyst, the less attention that can be given to an individual firm; *SPREAD* is therefore expected to correlate negatively with accuracy. *EXP* controls for analyst experience (Mikhail et al., 1997, 2003), and is set as the gap between the current forecast date and the first forecast date for the same analyst-firm forecast in IBES in hundreds of days.

FTE controls for the earnings bias found to vary with analyst forecasted total earnings due to greater attention paid to high forecasts (Hayes, 1998), or greater access to managerial information in exchange for generous forecasts (Das et al., 1998; Easterwood & Nutt, 1999; Hutira, 2016). *FTE* is generated as the natural log of 1 plus the forecasted total firm earnings in hundreds of billions of dollars.[89] I include the dummy variable *LOSS*, which is activated if actual earnings per share are

[88] Financial statement book values are lagged 4 months from Compustat year-end as a conservative estimate of the lag between the financial year-end and financial statement release. I exclude observations with negative book values of equity to avoid skewing estimates with distressed firms (Berkman, Dimitrov, Jain, Koch, & Tice, 2009).

[89] Forecasted total firm earnings are measured as *FEPS* multiplied by the adjusted number of shares outstanding. In unreported tests, I avoid using a natural log transformation on *FTE* and find similar results.

negative (Brown, 1998). Following Capstaff & Paudyal Rees (1998) and Hutira (2016), I include the dummy *ECHANGE*, activated for firms reporting lower earnings than the previous year.[90]

I present summary statistics on the final sample of 1987 to 2017 by *Toxic Releases* quartiles in Table 3.1.

Table 3.1: Summary statistics of variables used in tests of analyst bias. There are 425,621 observations in the sample ranging from 1987 to 2017. *Toxic Releases* is measured in billions of pounds. The mean of all variables used in the full regression sample are presented in the top half, disaggregated by quartiles of *Toxic Releases*. Quartile 1 consists of firms the lowest 25% of *Toxic Releases*, while quartile 4 consists of the top 25%. The bottom half is the pairwise Pearson correlation coefficients between variables. Correlations with significance at the 10% are bolded.

Quartile (*N*)	*FERROR*	*Toxic Releases*	*LOGSIZE*	*LOGBM*	*LEV*	*FPERIOD*	*COV*	*SPREAD*	*EXP*	*FTE*	*LOSS*	*ECHANGE*
1 (*106,321*)	-0.08	0.00000	21.92	0.33	0.19	181	9.96	7.72	15.46	0.006	0.07	0.30
2 (*106,488*)	-0.10	0.00004	22.13	0.37	0.19	179	10.45	7.67	16.20	0.010	0.07	0.32
3 (*106,374*)	-0.11	0.00033	22.20	0.36	0.22	178	8.86	7.98	17.39	0.010	0.06	0.31
4 (*106,438*)	-0.13	0.01489	22.69	0.42	0.23	175	9.40	8.68	16.87	0.017	0.07	0.37
FERROR	1											
Toxic Releases	**0.01**	1										
LOGSIZE	**0.16**	**0.08**	1									
LOGBM	**-0.15**	**0.03**	**-0.40**	1								
LEV	**-0.04**	**0.04**	**-0.10**	**0.04**	1							
FPERIOD	**-0.11**	**-0.01**	**0.05**	**-0.04**	**-0.02**	1						
COV	**0.07**	**0.02**	**0.54**	**-0.14**	**-0.12**	**0.03**	1					
SPREAD	**0.01**	**0.01**	**0.06**	**0.05**	**0.02**	**-0.05**	**0.06**	1				
EXP	**0.04**	**0.03**	**0.20**	**-0.05**	**0.01**	**0.01**	**0.12**	**0.08**	1			
FTE	**0.05**	**0.03**	**0.55**	**-0.13**	**-0.14**	**0.02**	**0.27**	**0.05**	**0.08**	1		
LOSS	**-0.33**	**-0.01**	**-0.21**	**0.36**	**0.08**	**0.00**	**-0.02**	**-0.01**	**-0.03**	**-0.11**	1	
ECHANGE	**-0.30**	**0.03**	**-0.09**	**0.29**	**0.03**	0.00	**0.06**	**0.03**	**0.02**	**-0.03**	**0.29**	1

Summary statistics reveal that *Toxic Releases* exponentially increase by pollution quartile. *LOGSIZE* increases monotonically with the groups, suggesting the polluter firms tend to be larger in size. *FERROR* decreases by pollution quartile; this effect may be driven by correlations between *Toxic Releases* and other variables. The overall negative values for *FERROR* across all groups support previous findings that analyst forecasts are optimistic on average (Easterwood & Nutt, 1999).

[90] *FTE* and *LOSS* do not cause look-ahead bias as the dependent variable itself is a function of realised earnings. While these variables may appear to have a mechanical relationship with *FERROR*, this assumes that *FEPS* and *Actual* are completely uncorrelated; rational expectations would suggest that the two are strongly positively correlated. Including *FTE* and *LOSS* in regressions allows for imperfect correlations between the components of *FERROR*. I avoid using earnings quality variables in regressions due to simultaneity concerns between forecast accuracy and earnings management (Embong & Hosseini, 2018).

3.5. Polluter forecast bias

In this section I test for evidence of an average and systematic earnings forecast pessimism associated with polluting firms, as theorised under the primary hypothesis.

3.5.1. Panel regressions

In the unbalanced panel there are unique observations of *FERROR* for each forecast made by analyst *i*, for firm *j*, in month *t*. I test whether analyst forecasts for polluting firms exhibit systematic pessimism using the following fixed effects panel regression.

$$FERROR_{i,j,t} = \beta^{TR} * Toxic\ Releases_{j,t} + \boldsymbol{\beta^{X}} * \boldsymbol{X_{i,j,t}} + \varepsilon_{i,j,t} \qquad (3.2)$$

FERROR is regressed on *Toxic Releases* and a vector of control variables, $\boldsymbol{X}$. The coefficient of interest is β^{TR}, which measures the linear association between the *FERROR* and *Toxic Releases*. I include industry and monthly fixed effects to control for heterogeneities (Gormley & Matsa, 2013) across industry (Brown, 1997; 1998; Coën, Desfleurs, & L'Her, 2009) and time (Ciccone, 2005; Espahbodi, Espahbodi, & Espahbodi, 2015). I two-way cluster standard errors by firm and month (Petersen, 2009).[91] Results of the regressions are presented in Table 3.2.

[91] Results are similar if standard errors are instead clustered by analyst and firm or firm and year dimensions.

Table 3.2: Results of the analyst forecast bias panel regression. The dependent variable is *FERROR*, while the independent variable of interest is *Toxic Releases*. There are 425,621 observations in the sample of data. *Toxic Releases* is measured in billions of pounds. I present regression coefficient estimates with p-values in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Firm toxicity and analyst forecast bias

	(1)	(2)	(3)	(4)
Toxic Releases	1.198***	0.804***	0.795***	0.603***
	(0.000)	(0.001)	(0.001)	(0.004)
LOGSIZE		0.023***	0.030***	0.026***
		(0.000)	(0.000)	(0.000)
LOGBM		-0.202***	-0.197***	0.124***
		(0.000)	(0.000)	(0.001)
LEV		-0.116***	-0.118***	-0.047
		(0.003)	(0.003)	(0.227)
FPERIOD			-0.066***	-0.063***
			(0.000)	(0.000)
COV			-0.003***	0.001
			(0.000)	(0.252)
SPREAD			0.000	0.000
			(0.499)	(0.610)
EXP			0.0002*	0.0004***
			(0.069)	(0.000)
FTE				-0.988**
				(0.014)
LOSS				-0.506***
				(0.000)
ECHANGE				-0.245***
				(0.000)
Fixed effects	Industry & Month	Industry & Month	Industry & Month	Industry & Month
N	425,621	425,621	425,621	425,621
Adjusted R^2	0.0715	0.0878	0.0977	0.2162

Results are consistent with the hypothesised pessimistic polluter forecast bias. Estimated coefficients of *Toxic Releases* reveal that pollution is positively related to *FERROR*, and is thus associated with systematically pessimistic analyst forecasts for polluting firms on average. Significance for the polluter bias coefficient marginally decreases as variables are added to the model, but maintains over 1% significance. Based on the final coefficient from column (4), increasing emissions by one standard deviation leads to an approximate 1.26% increase in the within-industry and month forecast error percentage.[92] I find consistent estimates with 1.1% statistical significance if *Toxic Releases* is transformed with a logarithmic function. Removing either industry or monthly fixed effects generates insignificant estimates, suggesting that the relationship between analyst pessimism and toxic emissions only occurs within-industry and month. I find consistent estimates with 5.1% statistical significance if analyst fixed effects are added to the model.

Estimated control variable coefficients are mostly consistent with the literature. As hypothesised by Brown (1998) and Hutira (2016), larger firms have significantly less optimistic forecasts. The *LOGBM* coefficient retains significance but switches sign once *FTE*, *LOSS* and *ECHANGE* are included, indicating that optimistic forecasts for high book-to-market stocks become pessimistic once other variables are accounted for. Forecasts made closer to the earnings date are less optimistic (Richardson et al., 1999; Burgstahler & Eames, 2003; Eames & Kim, 2012; Barron et al., 2013; Hutira, 2016). Analyst coverage is associated with forecast optimism until *FTE*, *LOSS* and *ECHANGE* are controlled for. Analyst experience is associated with pessimistic and more accurate forecasts (Mikhail et al., 1997, 2003). Lastly, the *FTE*, *LOSS* and *ECHANGE* coefficients are statistically significant, indicating that analysts are optimistic for firms with high expected earnings (Das et al., 1998; Hayes 1998; Hutira, 2016), that eventually report losses (Brown, 1998), or report earnings lower than the previous year (Capstaff & Paudyal Rees, 1998; Hutira, 2016).[93]

The adjusted R^2 reveals that the full model accounts for approximately 22% of the variation in forecast bias. I find that excluding *Toxic Releases* from the full model generates approximately the same adjusted R^2, suggesting that pollution contributes little to the explanatory power of the model.

I conduct four robustness regressions in which *FERROR* is winsorized differently. In the first of these tests, I avoid winsorizing data altogether and exclude observations in which the denominator *Actual* is 0. In the second and third of these tests, I winsorize *FERROR* at the 0.5% and 5% levels on

[92] The standard deviation of *Toxic Releases* for the regression sample is 0.0209.

[93] In a robustness test presented in the appendix, I rerun regression *(3.2)* and control for corporate governance variables. Results show that the *Toxic Releases* coefficient is positive and significant at the 1% level.

both sides respectively. Lastly, I replace *FERROR* with a rank variable of 1 to 10, dependent on the decile ranking of *FERROR* over the entire sample. Using these robustness *FERROR*'s, I rerun regression *(3.2)* with all control variables, fixed effects and two-way clustered standard errors. I present the estimated robustness coefficients for *Toxic Releases* in Table 3.3.

Table 3.3: Robustness tests using alternative winsorization schemes for *FERROR*. I run 4 separate regressions in these tests. The first test avoids winsorizing at all and excludes observations in which the actual earnings per share reported by a firm is 0. The second and third tests symmetrically winsorize *FERROR* at the 99.5% and 0.5% and 95% and 5% levels respectively. The last method creates new variables that correspond to the sample decile ranking of *FERROR*. I present estimated coefficients for *Toxic Releases* below, which are generated after repeating regression *(3.2)* with all control variables and industry and monthly fixed effects. P-values are shown in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Winsorization robustness test regression estimates	
Non-winsorized	1.126
	(0.113)
Winsorized at 0.5% on both sides	1.229*
	(0.081)
Winsorized at 5% on both sides	0.344**
	(0.046)
Decile rankings	3.537***
	(0.000)

Robustness estimates are mostly consistent with main results. Apart from marginally insignificant estimates using non-winsorized *FERROR*, robustness winsorization techniques yield qualitatively similar significant estimates in the same direction, and support primary regression estimates.

As an additional test, I estimate the regression coefficient for *Toxic Releases* when regressed against *FERROR* in subsamples of observations by yearly pollution quintiles. This estimates whether the polluter pessimism association is stronger within specific polluter groups. Behavioural biases are unlikely to operate in a smooth manner across all firms. Analysts may generate differing levels of pessimism through variation in pollution within these groups. As behavioural biases are irrational by definition and a product of subconscious heuristic-driven processes, it is reasonable to expect the strongest biases within extreme polluter quintiles which are the most noticeable.

I sort observations into quintile groups based on firm pollution relative to other observations in the same year. A higher quintile indicates greater yearly pollution. I then run regression *(3.2)* by quintile. I include all control variables from regression *(3.2)*, and control for monthly and industry fixed effects. Significance is estimated using two-way clustered standard errors. All five estimated regression coefficients for *Toxic Releases* are presented in Table 3.4.

Table 3.4: Analyst forecast bias by polluter quintile. I generate quintile subsamples based on the ranking of a firm's *Toxic Releases* relative to other firms in a year. I run regression *(3.2)* for all five subsamples. I present estimated coefficients for *Toxic Releases* below, which are generated after including all control variables from regression *(3.2)* with industry and month fixed effects. P-values are shown in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***. I also present the subsample means and standard deviations of *Toxic Releases* by quintile.

Polluter quintile subsample regressions

Quintile	Estimate	*Toxic Releases* mean	*Toxic Releases* std dev
1	-2610.055	0.00000	0.00000
	(0.618)		
2	-210.760	0.00001	0.00001
	(0.796)		
3	43.048	0.00007	0.00005
	(0.874)		
4	47.408	0.00036	0.00021
	(0.235)		
5	0.556**	0.01439	0.03912
	(0.015)		

Results of the polluter quintile test reveal that a linear relationship between *Toxic Releases* and *FERROR* is only significant for the largest polluting quintile. Most striking is the range of magnitudes of regression coefficients. Lower quintiles of *Toxic Releases* have exponentially lower subsample averages and standard deviations of pollution, creating estimates of greatly varying magnitudes. Results show that for firms in quintile 5, increases in pollution leads to a significant increase in forecast pessimism; however, this does not hold for firms outside the top polluter quintile.

From a behavioural perspective, this result may be explained by analysts being more sensitive to increases in pollution of the most toxic firms. Analysts may not notice or scrutinise the emissions of moderate polluters, but may be more influenced by information on the dramatic toxic emissions of the most extreme polluters. This information is likely to be more attention-grabbing and hence generate the most bias. Pollution for the top quintile is more likely to be salient information, and as Amir & Ganzach (1998) suggest, this is more likely to generate analyst overreaction.

I examine the interaction between polluter pessimism and time to test for a trend in pessimism. This test aims to illustrate whether polluter pessimism has always been prevalent throughout the 31 years included in the sample, or if it is relatively recent phenomenon. Both hypotheses are plausible; pollution has been on the minds of businesses and policymakers much prior to 1987;[94] however, recent sentiment for pollution abatement (Flammer, 2013) may generate additional forecast bias. I repeat regression *(3.2)* with all control variables, industry and yearly fixed effects, and an interaction term between *Toxic Releases* and the yearly dummies. Yearly fixed effects are used in place of monthly fixed effects as the latter leads to noisier interaction estimates, from which a trend is less discernible. Along with the estimated slope for *Toxic Releases*, the panel regression yields 30 interaction coefficients between the yearly dummies and *Toxic Releases*; one for each of the 31 years in the sample except 1987, which has no dummy and is the benchmark to which the estimates are compared against. I present the estimated interaction coefficients in Figure 3.3.

[94] For example, Eisenhower's Air Pollution Control Act of 1955 provided funds for federal research in air pollution, while Nixon's Clean Air Act of 1963 allowed federal programs to monitor and control air pollution.

Figure 3.3: Estimated interaction coefficients between *Toxic Releases* and yearly fixed effects, where the dependent variable is *FERROR*. Coefficients are estimated using *Toxic Releases*-yearly fixed effects interactions. 1987 has no active dummy, and is therefore the benchmark to which the dummy interaction coefficients are compared against. The estimated *Toxic Releases* coefficient is -6.834 in this industry and yearly fixed effects interaction model. The points on the thick black line are the estimated interaction coefficients, while the dotted black line is a fitted trend. The fitted coefficients of the trend are shown in the bottom right corner of the figure.

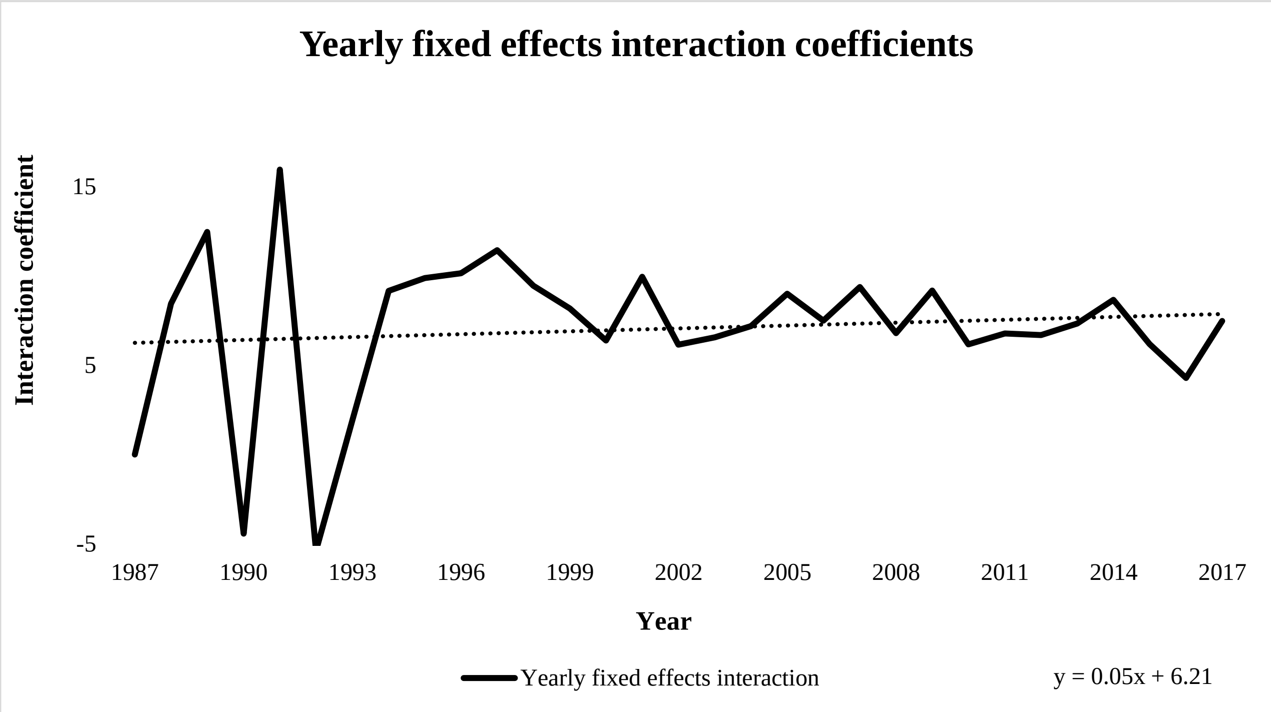

While the interaction estimates appear greater in magnitude than the original coefficient found in regression *(3.2)*, after adjusting for the standalone *Toxic Releases* coefficient of approximately -7, the total pessimism approaches the ballpark of the initial estimate. Though estimated coefficients are volatile early on, they become relatively stable from 1994 onwards. The initial volatility may be due to the launch of the TRI program, in which reporting standards, chemical lists, and federal compliance laws were being updated. There is a slight upwards drift in yearly interaction coefficients; an estimated trend of 0.05 indicates greater pessimism over time, however, this is economically insignificant given the overall scale of *Toxic Releases* and *FERROR*. Estimates of this test imply that analyst pessimism is largely stable throughout the sample.

3.5.2. Scaled pollution

In primary tests, absolute pollution is compared to forecast bias. However, some polluters may generate enough societal benefits to justify their pollution. For example, a pharmaceutical company which creates valuable medical products that are needed by society but also releases substantial toxins may not be perceived as poorly by analysts as a failing metals firm that creates little value but also pollutes heavily. Therefore, I test the association between *inefficient* pollution and forecast pessimism.

I create a measure of pollution inefficiency, labelled *Scaled Releases*, by dividing *Toxic Releases* by the most recent annualised net sales of the firm.[95] When constructing *Scaled Releases*, both *Toxic Releases* and sales are measured in total pounds and dollars respectively. The Pearson correlation coefficient between *Toxic Releases* and *Scaled Releases* at the firm-year level is measured as 0.41 and is significant at the 1% level, indicating absolute polluters are generally also inefficient polluters.

I conduct the following fixed effects panel regression to test the association between pollution inefficiency and forecast bias.

$$FERROR_{i,j,t} = \beta^{SR} * Scaled\ Releases_{j,t} + \boldsymbol{\beta^{X}} * \boldsymbol{X_{i,j,t}} + \varepsilon_{i,j,t} \qquad (3.3)$$

FERROR is regressed on *Scaled Releases* along with a set of control variables. I also rerun the regression and include *Toxic Releases* as an additional independent control variable. I use industry and monthly fixed effects, and cluster standard errors by firm and month. I present the results of the polluter efficiency regression in Table 3.5.

[95] As with other firm-specific fundamental variables, I use a 4-month lag to account for the gap in time between the balance sheet year-end date and the release of the financial data. I exclude observations in which net sales are zero or negative. There are only 20 observations in the entire sample for which this is necessary.

Table 3.5: Results of the polluter efficiency panel regressions. The dependent variable is *FERROR*, while the independent variables of interest is *Scaled Releases*. There are 425,601 observations in the sample. *Scaled Releases* is measured as the ratio of *Toxic Releases* in pounds to the most recently announced net sales of the firm in dollars, while *Toxic Releases* is measured in billions of pounds. I present regression coefficient estimates with p-values in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Polluter efficiency regression results

Variable	(1)	(2)
Scaled Releases	0.623	0.415
	(0.153)	(0.188)
Toxic Releases		0.486**
		(0.028)
LOGSIZE	0.027***	0.027***
	(0.000)	(0.000)
LOGBM	0.127***	0.126***
	(0.001)	(0.001)
LEV	-0.043	-0.046
	(0.275)	(0.243)
FPERIOD	-0.063***	-0.063***
	(0.000)	(0.000)
COV	0.001	0.001
	(0.259)	(0.261)
SPREAD	0.000	0.000
	(0.586)	(0.609)
EXP	0.0004***	0.0004***
	(0.000)	(0.000)
FTE	-0.990**	-0.992**
	(0.015)	(0.014)
LOSS	-0.508***	-0.506***
	(0.000)	(0.000)
ECHANGE	-0.245***	-0.245***
	(0.000)	(0.000)
Fixed effects	Industry & Month	Industry & Month
N	425,601	425,601
Adjusted R^2	0.2162	0.2164

Results of the polluter efficiency regressions reveal that while *Toxic Releases* is significantly associated with increased forecast pessimism, *Scaled Releases* is not. Despite having a positive coefficient like *Toxic Releases*, *Scaled Releases* does not display statistical significance. A lack of significance cannot be attributed to the collinearity with *Toxic Releases*, as *Scaled Releases* is insignificant even when *Toxic Releases* is excluded from the regression. Including *Scaled Releases* also leaves the adjusted R^2 virtually unchanged from primary tests.

Estimates reveal that absolute releases have a stronger relationship with forecast pessimism than scaled releases, implying that analyst bias is not influenced by polluter efficiency, but is by absolute pollution. By nature, analyst behavioural biases are unlikely to be sophisticated, and instead are likely to be affected by naïve and eye-catching information around absolute pollution, consistent with these results. Though *Scaled Releases* is more informative and comparable across firms, forecast bias is more strongly associated with *Toxic Releases*, which is a more immediately salient variable.

3.5.3. Disaggregated toxic releases

Toxic Releases is identified as the sum of all chemical groups that are classified by the EPA as hazardous to human health or to the external environment. This includes over 500 different chemical groups which may individually have differing consequences when released. *Toxic Releases* implicitly assigns equal weights to these chemicals, but can be broken down into the major chemical groups to examine their disaggregated, and potentially unequal, associations with forecast bias.

I disaggregate *Toxic Releases* into three chemical subgroups, *TRI*, *PBT* and *Dioxin*. *TRI* consists of the standard chemicals covered by the Toxic Release Inventory, including certain forms of ammonia, aluminium, phosphorus, zinc, and related by-products. These chemicals are detrimental when exposed to humans or the environment. *PBT*, the second category of chemicals, is named after persistent bio-accumulative chemicals which include lead and mercury compounds. These chemicals are known to have long lasting effects on human health and the environment, and are not easily removed or destroyed. The final chemical group is *Dioxin*, which comprises of trace level quantities of pollutants generated as by-products of combustion and other industrial processes. Despite being released in small quantities, dioxins are highly toxic and cause a variety of health complications including reproductive, developmental, immune system or hormonal damage; dioxins may also cause cancer. As dioxins are released in very small quantities, their contribution to *Toxic Releases* and forecast pessimism is effectively ignored in primary regressions.

PBT and *Dioxin* chemicals are generally more harmful to the environment and human health (EPA, 1999), but are released in significantly lower quantities than *TRI*.

TRI, *PBT* and *Dioxin* measure the yearly releases of their respective chemicals by a firm in continuous terms. Data on pollutants are stored in billions of pounds, except for *Dioxin* which is measured in pounds. I shorten the sample period as data on *Dioxin* only begin from 2000 onwards. At the firm-year level, the Pearson correlation coefficients for *TRI*, *PBT* and *Dioxin* range from 0.03 to 0.55.

I conduct the following regression to test for the relationships between individual chemical groups and analyst forecast bias.

$$FERROR_{i,j,t} = \beta^{TRI} * TRI_{j,t} + \beta^{PBT} * PBT_{j,t} + \beta^{Dioxin} * Dioxin_{j,t} + \boldsymbol{\beta^{X}} * \boldsymbol{X_{i,j,t}} + \varepsilon_{i,j,t} \tag{3.4}$$

I regress *FERROR* on the firm releases of each of *TRI*, *PBT* and *Dioxin*. I include prior control variables and industry and monthly fixed effects. Standard errors are adjusted with two-way clustering by firm and month. I present the results of the test in Table 3.6.

Table 3.6: Results of the disaggregated toxic releases panel regression. The dependent variable is *FERROR*, while the independent variables of interest are *TRI*, *PBT* and *Dioxin*, which sum up to equal *Toxic Releases*. There are 283,707 observations in the sample from 2000 to 2017. *TRI* and *PBT* are measured in billions of pounds, while *Dioxin* is measured in pounds. I present regression coefficient estimates with p-values in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Disaggregated toxic releases regression

Variable	Coefficient
TRI	0.661***
	(0.003)
PBT	0.897
	(0.695)
Dioxin	0.004
	(0.126)
LOGSIZE	0.019***
	(0.000)
LOGBM	0.132***
	(0.005)
LEV	-0.011
	(0.807)
FPERIOD	-0.053***
	(0.000)
COV	0.002***
	(0.008)
SPREAD	0.000
	(0.575)
EXP	0.0003***
	(0.003)
FTE	-0.718*
	(0.063)
LOSS	-0.435***
	(0.000)
ECHANGE	-0.203***
	(0.000)
Fixed effects	Industry & Month
N	283,707
Adjusted R^2	0.1795

Results of the disaggregated toxic releases test suggest that analysts are more pessimistic for firms with greater levels of *TRI* than the other two chemical groups. While all three chemical group variables have positive coefficients, only *TRI* is estimated with statistical significance. Despite being the most released pollutants, *TRI* chemicals are usually relatively less harmful than the other two categories. Again, it may be that *TRI* chemicals are the most noticed by analysts as they are generally released in greater quantities, and generate the most pessimism as a result of being attention-grabbing.[96]

3.6. Analyst-specific tests

Behavioural biases revolve around the tendencies of individual agents which are likely heterogenous. In this section I test the biases of individual analysts towards polluters.[97]

3.6.1. Analyst conservatism

I test whether prior analyst-specific biases on polluters are associated with similar future bias. I hypothesise that green and grey analysts display persistence in their bias towards polluters, and generate forecasts that are consistently biased in the same direction.

I identify individual analysts as green or grey dependent on their relative attitudes towards polluting firms. This procedure involves identifying polluting firms, identifying analyst biases towards polluters, and finally testing for a relationship between ex-ante and ex-post biases.

I start by sorting firms into yearly quintiles of *Toxic Releases*. Quintile 5 represents the highest yearly ranked polluters, while quintile 1 represents the lowest.

Analysts can only be identified as having been biased in their forecasts once realised firm earnings have been announced; therefore, I estimate observed biases as at earnings announcements. I calculate the average observed bias displayed by an analyst for a given firm and earnings announcement as the simple mean bias in all forecasts made over *t* months by analyst *i*, for firm *j*, for earnings announcement month *h*, as follows.

$$Observed\ FERROR_{i,j,h} = \frac{\sum_{t=1}^{T} FERROR_{i,j,t,h}}{T} \qquad (3.5)$$

[96] I find that transforming the disaggregated pollution variables with a natural log function generates insignificant estimates for all three coefficients.

[97] Analysts may individually vary in exhibiting behavioural biases. While in aggregate, analyst forecast bias is found to be negatively associated with firm pollution, individual analysts may exhibit forecast optimism through various factors, which include opposing behavioural biases, personal incentives or political beliefs.

Observed FERROR is the realised average bias of all forecasts made by an analyst for a specific firm-earnings announcement month. I then take the average of *Observed FERROR* made by analyst *i* for all firms in polluter quintile *q* over the last year.

$$Q\ FERROR_{i,q,h} = \frac{\sum_{y=h-11}^{h} Observed\ FERROR_{i,j,q,h}}{n} \quad for\ n\ observations\ in\ the\ rolling\ polluter\ quintile\ window \tag{3.6}$$

This process generates a rolling window of the observed average bias of individual analysts by polluter quintile for every month *h* in which there was an earnings announcement for a covered firm. For all months in between earnings announcements by firms covered by analyst *i*, I set the value of *Q FERROR* as its last known analyst-specific value.

I rank individual analysts based on their average bias towards firms in quintile 5. Only analysts that produce forecasts for the top polluter quintile are eligible to be defined as green or grey. *Q FERROR* for quintile 5 polluting firms is labelled *Polluter FERROR*.

$$Polluter\ FERROR_{i,h} = Q\ FERROR_{i,q,h}\ \ if\ q = 5 \tag{3.7}$$

Analysts with the highest 20% values of *Polluter FERROR* in a month are labelled as *Green*, the lowest 20% are labelled *Grey*, and those in the middle 60% or that do not forecast for quintile 5 polluters are labelled *Neutral*. *Green* and *Grey* analysts are identified with two separate dummy variables. Of the 7,428 analysts in the entire sample, there are 1,439 (1,243) analysts who are considered *Green* (*Grey*) analysts in at least one month.

This dynamic analyst identification strategy sorts analysts using the most recent realisation of their bias. The primary advantage of this dynamic structure is that analyst identifiers are updated every month following earnings announcements; analysts that display new evidence of bias around polluters can be classified appropriately. A disadvantage of this strategy is that the *Green* and *Grey* identifiers may be volatile.

I test whether *Green* and *Grey* analysts display persistent biases when forecasting for extreme polluters with the following fixed effects panel regressions.

$$FERROR_{i,j,t} = \beta^{GreenQ5} * (Green_{i,t-1} * Q5_{j,t}) + \beta^{GreenQ1} * (Green_{i,t-1} * Q1_{j,t}) + \beta^{GreyQ5} * (Grey_{i,t-1} * Q5_{j,t}) + \beta^{GreyQ1} * (Grey_{i,t-1} * Q1_{j,t}) + \boldsymbol{\beta^{X}} * \boldsymbol{X_{i,j,t}} + \varepsilon_{i,j,t} \qquad (3.8)$$

I regress *FERROR* on *Green* and *Grey*, along with *Q5* and *Q1*, representing the top and bottom quintile of polluting firms respectively. I lag the analyst identifier dummies to capture analyst identifications ex-ante. If an analyst was identified as *Green* (*Grey*) in the previous month, all observed forecasts in the following month made by the analyst have a value of 1 for the *Green* (*Grey*) dummy, and otherwise 0. The variables of interest are the interaction terms between the analyst identifiers and the polluter quintiles. The hypothesis suggests that the *Green* (*Grey*) dummy should have a positive (negative) interaction coefficient with *Q5*. I include interactions with polluter *Q1* to test whether estimates reverse for the least polluting firms. Forecasts made by *Neutral* analysts are the benchmark to which the standalone dummy and interaction coefficients are compared against.

The vector ***X*** consists of the standalone dummy variables, as well as all previously included control variables. These include *LOGSIZE*, *LOGBM*, *LEV*, *FPERIOD*, *COV*, *SPREAD*, *EXP*, *FTE*, *LOSS*, *ECHANGE*. I also include industry and monthly fixed effects. Standard errors are adjusted using two-way clustering by firm and month. Estimated coefficients for the analyst identifier and polluter quintile dummy variables and their interaction coefficients are presented in Table 3.7.

Table 3.7: Analyst-specific bias persistence regression. Analysts are identified as *Green* or *Grey* dependent on their relative bias for polluting firms. Firms identified with the *Q5* and *Q1* dummy variables are the top and bottom yearly quintiles of polluters respectively. Analyst identifier and polluter quintile interactions are reported below. I control for all independent variables used in primary regressions and use industry and monthly fixed effects. I present regression coefficient estimates with p-values in brackets below. Coefficient estimates for control variables are omitted for brevity. There are 383,621 observations in the panel. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***. I present Wald test p-values for each subsample to test the equality of interaction coefficients.

Analyst persistence regression results

Variable	Coefficient
Green	-0.004
	(0.709)
Grey	0.006
	(0.610)
Q5	-0.009
	(0.528)
Q1	0.004
	(0.727)
*Green * Q5*	0.029*
	(0.084)
*Green * Q1*	0.002
	(0.900)
*Grey * Q5*	-0.048**
	(0.041)
*Grey * Q1*	0.005
	(0.808)
Fixed effects	Industry & Month
N	383,621
Adjusted R^2	0.2147

Wald test p-values

	*Green * Q5*	*Grey * Q1*
*Green * Q1*	0.166	0.914
*Grey * Q5*	0.005***	0.024**

Results are consistent with the analyst conservatism hypothesis, revealing that analysts are persistent in their bias towards top polluting firms. Relative to forecasts made by *Neutral* analysts, ex-ante identified pessimistic and optimistic analysts do not display a statistically significant bias for the mid quintiles of polluters, as indicated by the standalone coefficients for *Green* and *Grey*. Similarly, *Neutral* analysts do not display additional bias for the top or bottom polluter quintiles relative to the mid quintiles, as indicated by the standalone coefficients for *Q5* and *Q1*. However, after accounting for the benchmark and standalone dummy coefficients, analysts previously identified as *Green* (*Grey*) continue to issue pessimistic (optimistic) forecasts for the top quintile of polluters, as revealed by the significantly estimated interaction coefficients between analyst identifiers and *Q5*. Results show that being identified as a *Green* (*Grey*) analyst in the previous month is associated with a 2.9% (4.8%) increase in forecast pessimism (optimism) for forecasts made for polluters in the current month. Estimates provide no evidence of a reverse relationship for *Q1*.

The difference in interaction effects between *Green* and *Grey* analysts for *Q5* is approximately 6.7% after adjusting for the standalone analyst dummies. The Wald test finds a significant difference between two interaction coefficients at the 10% level.

I consider the forecast bias persistence of *Green* and *Grey* analysts for more extreme polluters. I rerun regression *(3.8)* but substitute *Q1* and *Q5* for *D10*, a dummy which is activated for the top yearly decile of polluters. *Green* and *Grey* are interacted with *D10*. The estimated interaction coefficients are expected to be of the same sign as those estimated using *Q5*, but of greater magnitude. I present the results of this test in Table 3.8.

Table 3.8: Analyst-specific bias persistence regression for the top decile of polluters. Analysts are identified as *Green* or *Grey* dependent on their relative bias for firms identified as polluters. Firms identified with the *D10* dummy variable are the top yearly decile of polluters. Analyst identifier and polluter decile interactions are reported below. I control for all independent variables used in primary regressions and use industry and monthly fixed effects. I present regression coefficient estimates with p-values in brackets below. Coefficient estimates for control variables are omitted for brevity. There are 383,621 observations in the panel. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Analyst persistence regression for top decile of polluters

Variable	Coefficient
Green	-0.006
	(0.455)
Grey	-0.003
	(0.795)
D10	-0.003
	(0.872)
*Green * D10*	0.060^{***}
	(0.001)
*Grey * D10*	-0.043
	(0.118)
Fixed effects	Industry & Month
N	383,621
Adjusted R^2	0.2147

As expected, the estimated interaction coefficient for *Green* analysts and *D10* is of the same direction as the interaction with *Q5*, but of a greater magnitude. This is not true for *Grey* analysts, who surprisingly have slightly lower estimates of forecast pessimism with marginal insignificance. It is possible that only for these extreme polluting firms, some *Grey* analysts revise their beliefs somewhat, whereas *Green* analysts do not. I find similar results if *D10* is substituted with a top 5% polluter dummy variable.

Interestingly, the estimated coefficient for *D10* itself is not statistically significant, implying that *Neutral* analysts forecasting for the greatest decile of polluters are not more pessimistic relative to the remaining 90% of firms; this also true for the *Q5* coefficient in the previous test. This suggests that the aggregate pessimism found in the initial tests is generated by *Green* analysts, offset somewhat by *Grey* analysts, and largely unrelated to forecasts made by *Neutral* analysts.

Results of this section suggest that analysts that are the most biased towards polluters display traits consistent with the conservatism bias.[98] Benchmarked to other forecasts, *Green* (*Grey*) analysts continue to display pessimism (optimism) for polluters, ceteris paribus, despite having been proven wrong in the past.

3.6.2. Polluter bias by forecast horizon

In this section I test whether individual analyst biases towards polluters change over the forecast horizon. I hypothesise that biased analysts ‘walk-down’ their polluter bias nearer to the earnings date. As time goes by, analysts may be pressured by the economic reality around polluter earnings. Nearer to the earnings date, the chance of unpredictable or extreme events occurring within the timeframe diminishes, and hence the expected cost to polluters is lower. Also, closer to the earnings date, more information is available which may promote cognitive forecasts that are less dependent on heuristics or emotion. Finally, there is likely to be greater investor attention to forecasts closer to earnings dates, incentivising analyst accuracy.

I first examine the analyst bias walk-down hypothesis with a univariate test. Using the ex-ante *Green*, *Grey* and *Neutral* identifiers for analysts, I disaggregate the full sample into subsamples by analyst identifier and the forecast horizon groupings of 1 – 90, 91 – 180, 181 – 270 and 271 – 360 days inclusive. I present the average value of *FERROR* for the top yearly quintile of polluters by analyst identifier subsample in Figure 3.4.

[98] As pollution is heavily clustered by industry, tests may be picking up an industry effect. However, I argue that because industry groups themselves are based on the common activities of the firms that operate within them, it is these activities that drive results instead of the industry group itself, which is little more than a title. In other words, the pollution effect may drive the industry effect, as opposed to the other way around.

Figure 3.4: Average values of *FERROR* for the top yearly quintile of polluters, by analyst identifier and forecast horizon subsamples. Analysts are grouped by ex-ante indicators as *Green*, *Neutral* or *Grey*, as described in the analyst conservatism section. All forecasts made by these analysts are further grouped into 4 subsamples based on the difference in days from the forecast date and the earnings date. Average values of *FERROR* are then calculated and presented for these 12 subsamples.

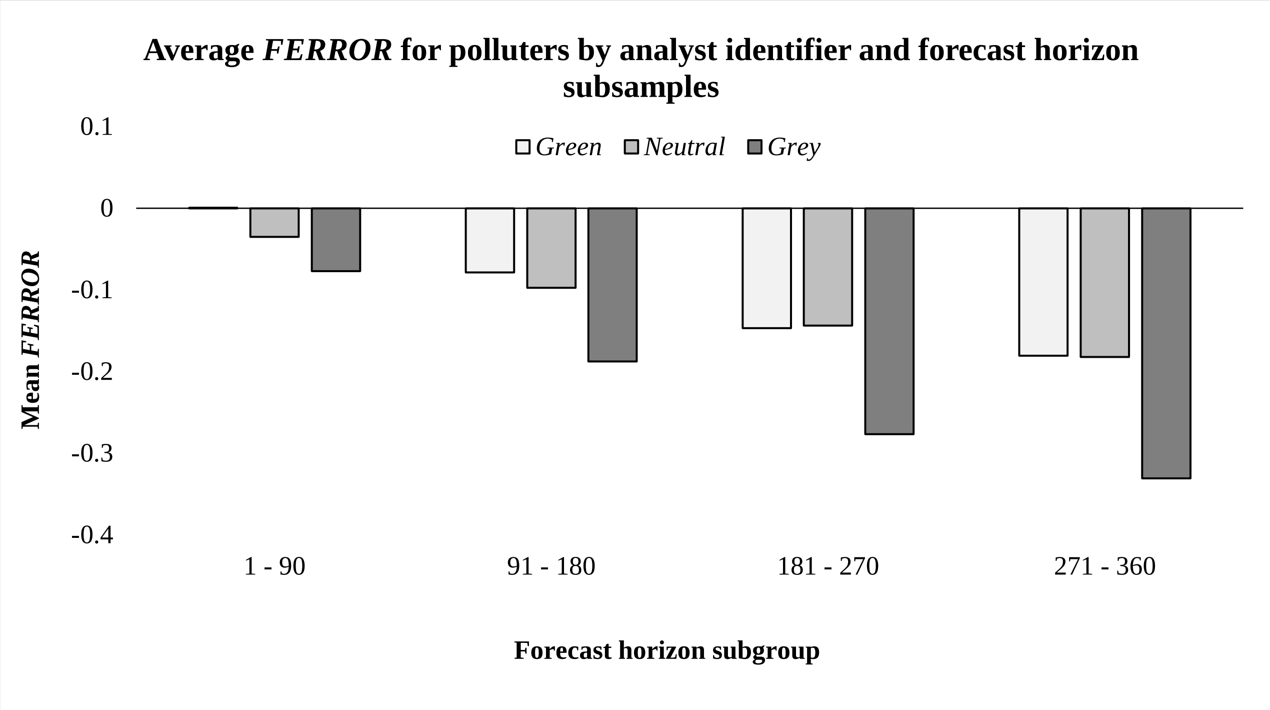

Figure 3.4 illustrates that all three analyst types initially have optimistic forecasts for the top polluters on average, which is then walked down as the forecast horizon diminishes. As expected, *Grey* analysts have the most optimistic polluter forecasts, followed by *Neutral* and then *Green* analysts. Figure 3.4 reveals that in the last 90 days before the earnings date, all analyst types have reduced optimism, with the average *Green* forecasts becoming marginally pessimistic. While Figure 3.4 provides some evidence of a forecast optimism walk-down by *Grey* analysts, it does not indicate that *Green* analysts walk-down their pessimism as hypothesised; I test whether this is the case after controlling for other factors in a regression.

I focus on the forecasts made by analysts for the top quintile of polluters. I conduct the following panel regression on this subsample of firms.

$$FERROR_{i,j,t,q=5} = \beta^{Green} * (Green_{i,j,t-1} * FPERIOD_{i,j,t}) + \beta^{Grey} * (Grey_{i,j,t-1} * FPERIOD_{i,j,t}) + \boldsymbol{\beta^{X}} * \boldsymbol{X_{i,j,t}} + \varepsilon_{i,j,t} \quad (3.9)$$

For the top quintile of polluters, I regress *FERROR* on the interaction between the *Green* and *Grey* analyst identifier dummy variables and *FPERIOD*. In another regression, I alternatively replace *FPERIOD* with 3 dummy variables, activated for forecasts made within the 91 – 180, 181 – 270, and 271 – 360 day forecast horizons, and interact each of dummy with *Green* and *Grey*. These dummy variables are labelled *FGROUP2*, *FGROUP3* and *FGROUP4* respectively. The benchmark group against which standalone dummies and interaction effects are compared to are forecasts made by *Neutral* analysts for polluting firms in the 90 days prior to the earnings date. I include all standalone interaction components in both regressions. I control for all independent variables from primary regressions,[99] and use industry and monthly fixed effects. Standard errors are two-way clustered by firm and month.

As per the hypothesis, the interaction coefficient between *Green* (*Grey*) and *FPERIOD* is expected to be positive (negative), indicating that with a greater forecast horizon, *Green* (*Grey*) analysts are more pessimistic (optimistic). Similarly, the interaction coefficients between *Green* (*Grey*) and the forecast horizon dummy variables should be positive (negative) and monotonically increasing (decreasing). I present the estimated interaction coefficients in Table 3.9.

[99] The second regression model does not include the independent variable *FPERIOD*.

Table 3.9: Results of the analyst bias by forecast horizon test for the subsample of polluters. I run two panel regressions, in which the analyst identifiers *Green* and *Grey* are interacted with the forecast horizon variable. In the first regression, the forecast horizon variable is *FPERIOD*, measured as the gap between the forecast and earnings date in hundreds of days. The second regression has 3 dummy variables, activated if the earnings forecast is made within 91 – 180, 181 – 270, and 271 – 360 day period prior to the earnings date, labelled *FGROUP2*, *FGROUP3*, *FGROUP4* respectively. Only forecasts for firms in the top yearly quintile of polluters are included in the sample. I control for all independent variables used in previous regressions, except for *FPERIOD* in column (2), and use industry and monthly fixed effects. I present interaction and standalone variable coefficients estimates with p-values in brackets below. Coefficient estimates for control variables are omitted for brevity. There are 100,834 observations in the panel. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Analyst bias and forecast horizon interaction results

Variable	(1)	(2)
Green	0.048**	0.046***
	(0.015)	(0.003)
Grey	0.042*	0.027
	(0.068)	(0.158)
FPERIOD	-0.060***	
	(0.000)	
*Green * FPERIOD*	-0.016*	
	(0.071)	
*Grey * FPERIOD*	-0.041***	
	(0.002)	
FGROUP2		-0.065**
		(0.013)
FGROUP3		-0.107***
		(0.000)
FGROUP4		-0.171***
		(0.000)
*Green * FGROUP2*		-0.026*
		(0.089)
*Green * FGROUP3*		-0.043**
		(0.045)
*Green * FGROUP4*		-0.045*
		(0.061)
*Grey * FGROUP2*		-0.047**
		(0.032)
*Grey * FGROUP3*		-0.083**
		(0.011)
*Grey * FGROUP4*		-0.112***
		(0.002)
Fixed effects	Industry & Month	Industry & Month
N	100,834	100,834
Adjusted R^2	0.2285	0.2288

Results support the walk-down hypothesis for *Grey* analysts, but are in the opposite direction to the hypothesis for *Green* analysts. The continuous *FPERIOD* regression results in column (1) shows *Grey* analysts become more pessimistic towards polluters closer to the earnings date, and are more pessimistic than *Neutral* analysts immediately prior to the earnings date. Surprisingly, *Green* analysts also become more pessimistic as *FPERIOD* decreases, and are more pessimistic than *Neutral* analysts immediately prior to the earnings date. The same is observed for estimates from the dummy variable regression. Based on the standalone regression coefficients in column (2), *Green* (*Grey*) analysts are significantly (insignificantly) more pessimistic towards polluters compared to *Neutral* analysts in the last 90 days prior to the earnings date. In line with the hypothesis, *Grey* interaction coefficients with the forecast horizon dummy variables are negative and rising in magnitude monotonically for horizon subgroups further away from the earnings date, indicating increased pessimism for early forecasts. The same is also true, albeit in smaller magnitude, for *Green* analysts, which is inconsistent with the hypothesis.

The negative *FPERIOD* coefficient along with the standalone dummy monotonic negative trend indicates that *Neutral* analysts become more pessimistic in later forecasts. Both *FPERIOD* and the forecast horizon dummies provide evidence of an average walk-down of analyst forecasts; however, the additional interactions with the *Green* and *Grey* dummies indicate that these analysts walk-down their forecasts for polluters more so than the average *Neutral* analyst.

In a robustness test, I repeat regressions (*3.9*) with firm-earnings date fixed effects. This test controls for the average effect of time invariant factors at the firm-earnings date level. I exclude control variables that have little or no variation at the firm-earnings date, consisting of *LOGSIZE*, *LOGMB*, *LOSS* and *ECHANGE*. The results of this test, which are presented in the appendix, are similar to Table 3.9.

I also examine analyst bias towards polluters in their last forecasts for a firm-earnings date. Using the same panel as in regressions (*3.9*), I introduce the *Last Forecast* dummy, activated if the observation is the last forecast made by an analyst for a firm-earnings date. For robustness, I separately consider unconditional last forecasts made by an analyst, as well as last forecasts conditionally made within the final 90 days prior to the upcoming earnings date. I interact *Last Forecast* with the *Green* and *Grey* analyst dummies. In accordance with the hypothesis, I expect the interaction between *Last Forecast* and *Green* (*Grey*) to generate a negative (positive) coefficient. I control for industry and monthly fixed effects, and two-way cluster standard errors by firm and month. I present the results of the regression in Table 3.10.

Table 3.10: Results of the final analyst forecast bias test. I run panel regressions in which the analyst identifiers *Green* and *Grey* are interacted *Last Forecast*, a dummy activated for the final forecast made by an analyst for a specific firm-earnings date. *Last Forecast* is separately activated for the absolute final forecasts made by analysts prior to the earnings date, and conditional for final forecasts that are made within a 90 day forecast horizon prior to the earnings date, with results of regressions run with either variable in columns (1) and (2) respectively. Only forecasts for firms in the top yearly quintile of polluters are included in the sample. I control for all independent variables used in previous regressions and use industry and monthly fixed effects. I present interaction and standalone variable coefficients estimates with p-values in brackets below. Coefficient estimates for control variables are omitted for brevity. There are 100,834 observations in the panel. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Analyst bias and final forecast interaction results

Variable	Unconditional last forecast	Last forecast within a 90 day forecast horizon
Green	0.015	0.015
	(0.373)	(0.399)
Grey	-0.043**	-0.044**
	(0.046)	(0.030)
Last Forecast	0.053***	0.067***
	(0.000)	(0.000)
*Green * Last Forecast*	0.016	0.024
	(0.264)	(0.198)
*Grey * Last Forecast*	0.048**	0.077***
	(0.033)	(0.002)
Fixed effects	Industry & Month	Industry & Month
N	100,834	100,834
Adjusted R^2	0.2264	0.2264

Results are consistent with the previous test. Estimated coefficients indicate that on average, the final forecasts for *Neutral* analysts are relatively pessimistic, consistent with analysts walking down their forecasts closer to earnings dates (Barron et al., 2013). As indicated by the analyst identifier and last forecast interactions, both *Green* and *Grey* analysts are more pessimistic compared to *Neutral* analysts; however, only the *Grey* interaction coefficient is statistically significant. Results generated on *Last Forecast* being conditionally activated for forecasts made in the last 90 days prior to the earnings date are similar but stronger in magnitude and statistical significance. Results indicate that in their final forecasts for a firm, *Grey* analysts walk-down their previous optimism for polluter firms, but *Green* analysts do not walk-down their prior pessimism.

Overall, tests indicate that benchmarked to *Neutral* analysts, *Grey* analysts further walk down their initial optimism, but *Green* analysts do not. In fact, some estimates suggest that *Green* analysts become increasingly pessimistic closer towards the earnings date. This is puzzling and in contrast to the polluter bias walk-down hypothesis. It is possible that *Green* analysts are influenced by *Grey* and *Neutral* analysts, and exhibit herding by also walking down their already pessimistic forecasts.

3.7. Polluter earnings surprises

I examine the return predictability of polluters around earnings announcements. If investor earnings expectations are tied to aggregate analyst forecasts, systematically pessimistic polluter forecasts are expected to generate systematically positive earnings surprises, ceteris paribus. Positive earnings surprises increase firm value through improved business outlooks, and above expected equity reserves due to higher than expected earnings. Polluters are therefore hypothesised to generate positive abnormal returns around earnings announcements as their equity prices move towards their intrinsic value. I test for this using both annual and quarterly earnings announcements.

I first examine the abnormal returns of polluter firms around earnings announcements. All firm-event observations in the sample are separated into polluter quartiles based yearly rankings of *Toxic Releases*. Abnormal returns are calculated as the excess of expected returns from the Carhart 4-factor model (Carhart, 1997). In the following figures, I plot both the average daily and cumulative abnormal returns of the quartile groups around their annual and quarterly earnings announcements, with a window of *t – 10* to *t + 10*.

Figure 3.5: Average daily abnormal returns of polluter quartiles over a window of t – 10 to t + 10 around firm annual earnings announcements. Abnormal returns are measured as the excess realised returns from the Carhart 4-factor model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 9,377 observations in the sample. Daily abnormal returns are presented in percentage format.

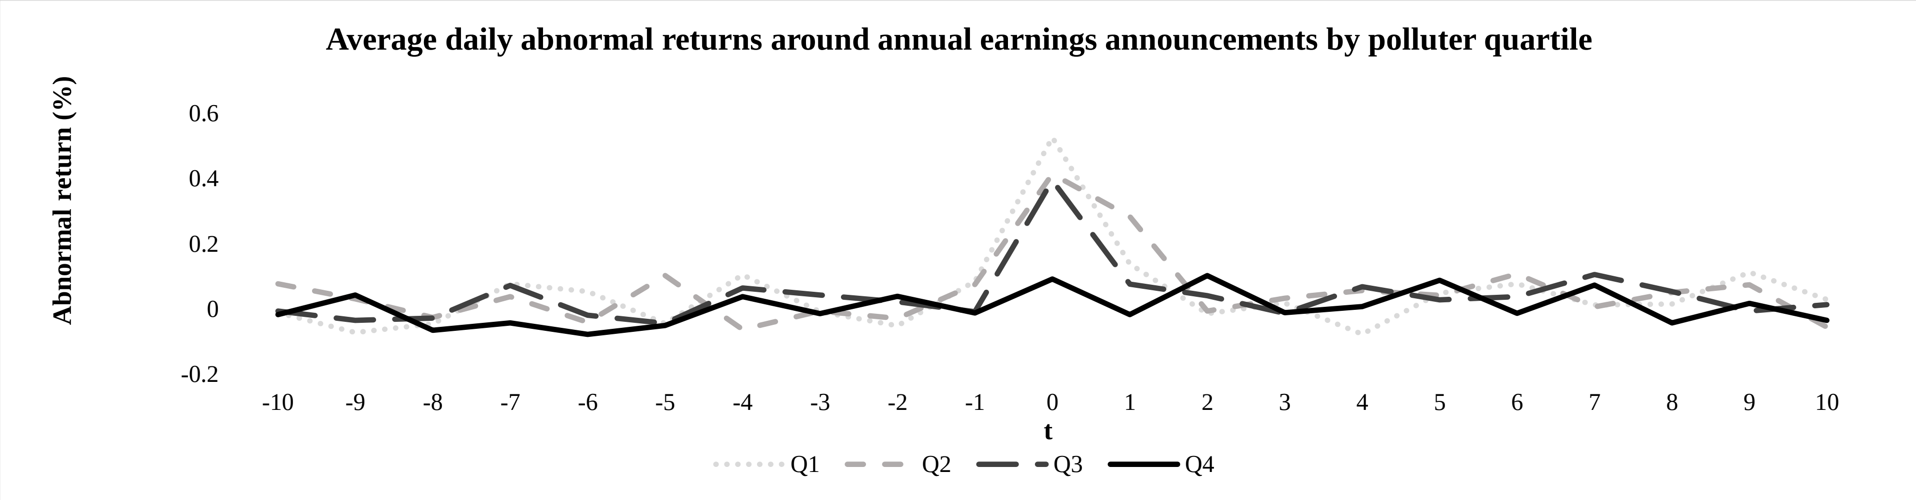

Figure 3.6: Average cumulative abnormal returns of polluter quartiles over a window of *t – 10* to *t + 10* around firm annual earnings announcements. Abnormal returns are measured as the excess realised returns from the Carhart 4-factor model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 9,377 observations in the sample. Cumulative returns are compounded daily abnormal returns and are presented in percentage format.

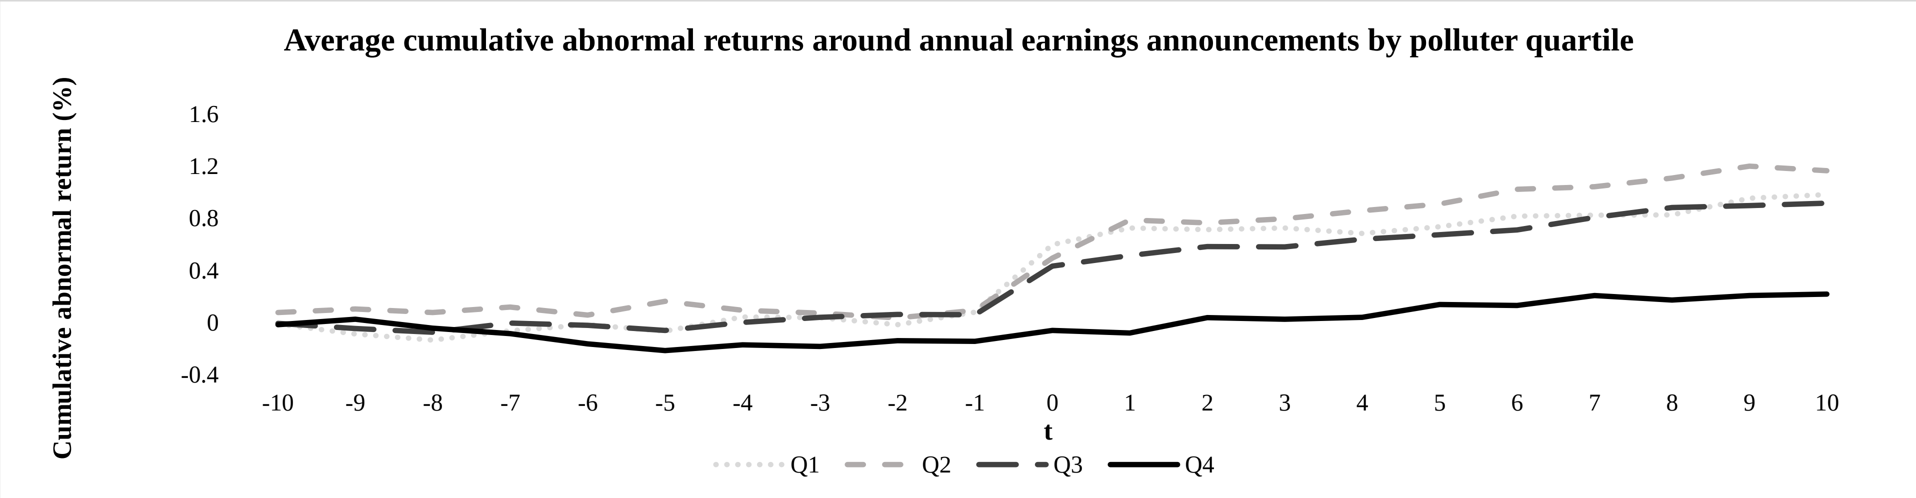

Figure 3.7: Average daily abnormal returns of polluter quartiles over a window of t – 10 to t + 10 around firm quarterly earnings announcements. Abnormal returns are measured as the excess realised returns from the Carhart 4-factor model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 39,399 observations in the sample. Daily abnormal returns are presented in percentage format.

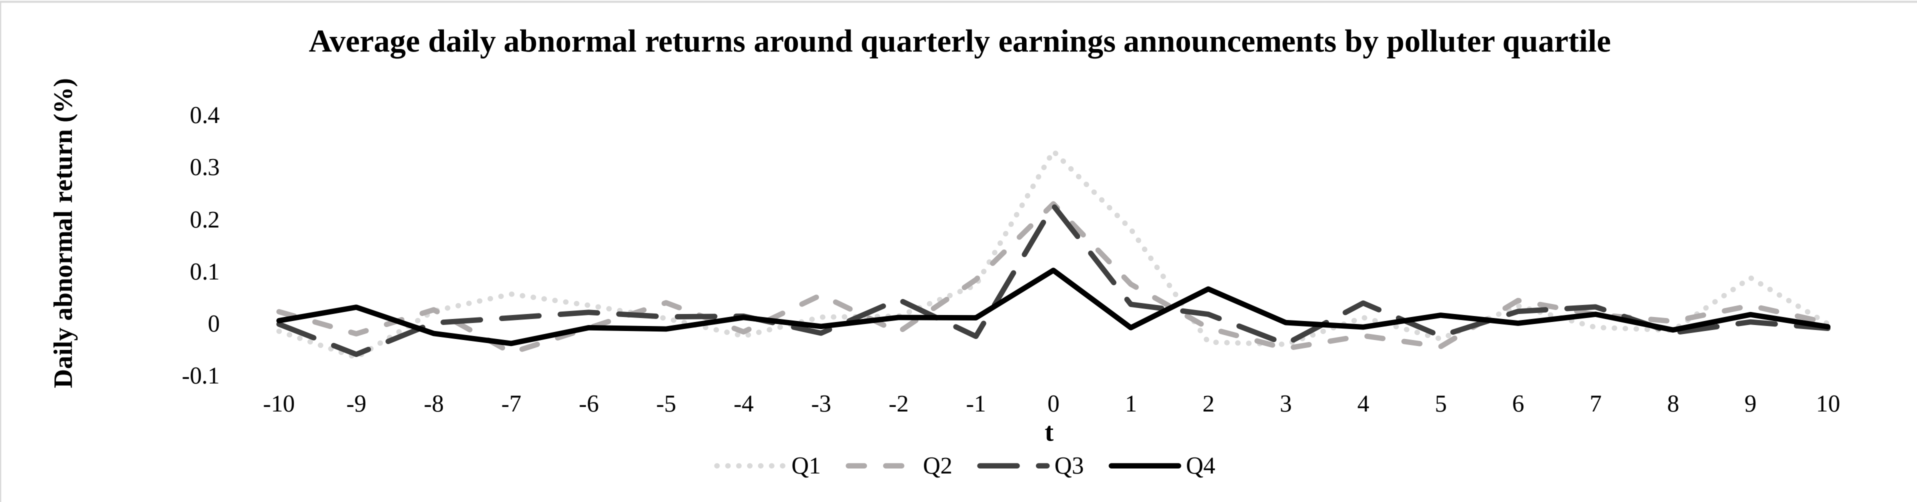

Figure 3.8: Average cumulative abnormal returns of polluter quartiles over a window of *t – 10* to *t + 10* around firm quarterly earnings announcements. Abnormal returns are measured as the excess realised returns from the Carhart 4-factor model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 39,399 observations in the sample. Cumulative returns are compounded daily abnormal returns and are presented in percentage format.

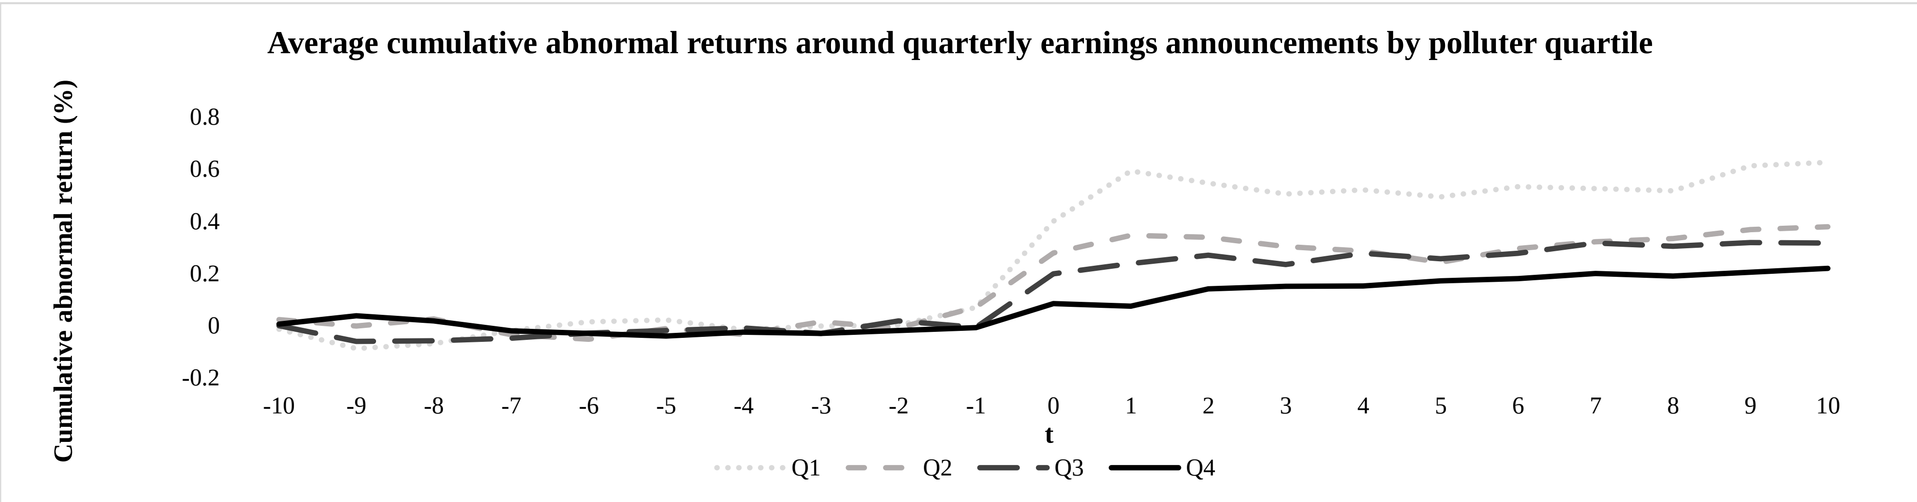

The relationship between the polluter quartiles and excess abnormal returns appears to be negative, in contrast to the hypothesis. Figures 3.5 and 3.7 reveal that for both annual and quarterly earnings announcements, there is a near-perfect negative monotonic relationship between pollution quartile and average returns. This relationship is reinforced by the cumulative abnormal return functions in Figures 3.6 and 3.8. For robustness, I repeat this test using abnormal returns benchmarked by the CAPM and find similar results; these figures are presented in the appendix.[100]

I test the polluter abnormal returns hypothesis within a panel regression after controlling for other variables.[101]

I follow Berkman et al. (2009) in creating the dependent variable *EXRET*, which is the cumulative abnormal returns generated over a window of $t - 1$ to $t + 1$ around either the annual or quarterly earnings announcement date on IBES. Berkman & Truong (2009) reveal that after-hours earnings announcements, which occur after trading has closed, have increased by a considerable amount in recent years, motivating the 3-day window. *EXRET* thus captures some leaked or delayed market reactions to earnings announcements.

I follow Berkman et al. (2009) in their choice of control variables. Independent variables include *LOGSIZE*, which controls for firm size effects on returns around earnings announcements. I control for *LOGBM*, which accounts for the effect of book to market value on returns around earnings announcements (Levis & Liodakis, 2001). Jegadeesh & Titman (1993) reveal that recent past winners earn higher announcement period returns; therefore, I control for the monthly geometric average stock returns from up to the last 3 months with *MOM*. *LEV* controls for firm leverage.

In an extension of Miller's hypothesis (Miller, 1977), Berkman et al. (2009) find that short sale constraints generate lower returns when combined with dispersion in investor opinions around

[100] In unreported robustness checks, I generate abnormal returns as based on the Fama-French 3 and 5-factor models and find similar results.

[101] In a similar test Berkman et al. (2009) use weighted Fama-MacBeth regressions, however Petersen (2009) shows that two-way clustered standard errors are more efficient when both firm and time conditional error clustering is present.

earnings announcements. I therefore control for the interaction of these two factors using the same proxy variables used by Berkman et al. (2009). *IO* is a proxy for short sale constraints, measured as the proportion of institutional ownership of equity as last reported in the 13F database. If *IO* is missing, I replace it with its most recent available value within the last 6 months, and otherwise 0. *DISP* is a proxy for difference of opinion in stock value. *DISP* is calculated as the standard deviation of individual analysts' latest forecasts in the 45 days prior to the earnings date. *IO * DISP* is the interaction term between short sale constraints and dispersion of opinion.

I use the following panel regression to test for polluter return predictability around earnings announcements.

$$EXRET_{j,t} = \alpha + \beta^{TR} * Toxic\ Releases_{j,t} + \boldsymbol{\beta^{X}} * \boldsymbol{X_{j,t}} + \varepsilon_{j,t} \quad (3.10)$$

EXRET is regressed on *Toxic Releases* and the vector of control variables $\boldsymbol{X}$.[102] I two-way cluster standard errors by firm and month. I repeat the regression for both annual and quarterly earnings announcements. I present results in Table 3.11.

[102] Independent variables are measured ex-ante of the dependent variable. Because *EXRET* is a cumulative abnormal return over three days, daily measured control variables are based on ex-ante information two days prior to the earnings announcement date as stated by IBES, due to the inclusion of $t-1$ in the return window. These variables include *LOGSIZE* and *LOGBM*. *MOM*, *LEV* and *IO* are measured as at the beginning of the month in which t falls.

Table 3.11: Earnings announcement abnormal returns and firm pollution. The dependent variable used in regressions is *EXRET*, while the independent variable of interest is *Toxic Releases*. *EXRET* measures cumulative abnormal returns from one day prior to the earnings announcement to one day after. Control variables are measured ex-ante of the return variables. *Toxic Releases* is measured in billions of pounds. Columns (1) and (2) use annual earnings announcements, while columns (3) and (4) use quarterly earnings announcements. I present regression coefficient estimates with p-values in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Return predictability of polluters around earnings announcements

	Annual	Annual	Quarterly	Quarterly
Toxic Releases	-2.349	-1.576	-0.602	-0.120
	(0.145)	(0.202)	(0.653)	(0.931)
LOGSIZE	-0.198***	-0.216***	-0.122***	-0.136***
	(0.000)	(0.000)	(0.000)	(0.000)
LOGBM	-0.734	-0.404	0.232	0.204
	(0.825)	(0.502)	(0.509)	(0.585)
MOM	-0.007	0.006	-0.006	-0.005
	(0.713)	(0.740)	(0.541)	(0.664)
LEV	-0.130	-0.211	0.126	-0.027
	(0.825)	(0.729)	(0.674)	(0.933)
IO		-0.089		0.174
		(0.817)		(0.464)
DISP		-0.009		0.011
		(0.652)		(0.532)
*IO * DISP*		-0.049		-0.064
		(0.658)		(0.535)
Constant	4.752***	5.278***	2.782***	3.003***
	(0.000)	(0.000)	(0.000)	(0.000)
N	9,378	8,763	39,408	36,221
Adjusted R^2	0.0030	0.0033	0.0015	0.0016

Results provide no evidence of an association between pollution and earnings announcement returns. The estimated coefficients for *Toxic Releases* are statistically insignificant for both quarterly and annual earnings announcements. Results are inconsistent with the hypothesis that polluters earn abnormally positive returns around earnings announcements through systematic analyst pessimism.[103]

[103] In primary results, pessimism in analyst forecasts is found for within-industry and month forecasts through fixed effects; therefore, it may be argued that within-industry and month polluter pessimism is a more relevant predictor of earnings announcement abnormal returns. In unreported robustness tests, I rerun earnings announcement regressions using industry and monthly fixed effects but again find statistically insignificant estimates.

Control variables are mostly estimated with insignificance. As in Berkman et al. (2009), *LOGSIZE* appears to significantly explain variation in earnings announcement returns, the other variables are not found to be significant at the 10% level.

It could be argued that analyst pessimism for polluters disappears closer to the earnings date; however, the previous section finds that it instead increases. I discuss some alternative reasons as to why polluter firms may not earn abnormal returns around their earnings announcements. One possibility is that markets may have already identified analyst bias and adjusted their expectations after incorporating systematic analyst pessimism for polluting firms. As a result, there are no positive earnings surprises as market expectations are already adjusted upwards relative to analysts' predictions. This theory suggests that investors are aware of, but not influenced by, analyst behavioural biases and rationally incorporate them into prices. Alternatively, earnings surprises may be largely independent of analyst forecasts such that investors are indifferent to analyst expectations altogether; however, this is unlikely given the literature on the correlation between analyst forecast errors and market earnings surprises.[104] As a middle ground to these theories, investors may only be aware of individual analysts that exhibit the most polluter pessimism, and ignore their forecasts. Finally, pollution may be correlated with an unidentified confounding variable which reduces abnormal returns around earnings announcements.

3.8. Conclusion

I examine the relationship between firm pollution and analyst forecast biases. I primarily hypothesise that various cognitive biases exhibited by analysts lead to systematically pessimistic forecasts for polluting firms. Analysts may overweight the probabilities of tail events which dramatically shock the profits of polluters but are unlikely to occur, as illustrated in behavioural finance theories. Analysts may also overly extrapolate the expected costs of pollution due to prior

[104] For example, see O'Brien (1988), Doyle, Lundholm, & Soliman (2006) and Livnat & Mendenhall (2006).

dramatic events in recent memory, and incorrectly set conditional expectations of the consequences of pollution on earnings.

I find that pollution is significantly associated with an aggregate analyst pessimism; earnings forecasts for polluters systematically undershoot actual earnings on average. Tests show that pessimism associated with pollution is the strongest within the top quintile subsample of polluters, and has not significantly changed in magnitude over the sample period. Results do not show evidence of forecast bias when pollution is scaled by the net sales of a firm, indicating that polluter efficiency is not associated with pessimism. Of the pollutants captured by the Toxic Release Inventory, those in the general TRI category generate the most pessimism. These findings are consistent with analysts being influenced by eye-catching, and somewhat naïve, information around pollution.

Tests provide evidence of a persistence in forecasts made by biased analysts. Analysts that are identified as pessimistic or optimistic towards polluters based on ex-ante information continue to display these biases, despite being proven wrong in recent earnings announcements. I also find that optimistic analysts walk-down their polluter optimism closer to the earnings date, while pessimistic analysts increase their pessimism closer to the earnings date.

I hypothesise that the systematic analyst forecast pessimism generates positive earnings surprises for polluting firms. Given the systematic polluter pessimism from analysts, actual earnings should be higher than analyst forecasts on average, and should generate positive earnings surprises. Contrary to this theory, results provide no evidence of polluting firms generating positive abnormal returns around their earnings announcements.

Further research could examine the relevance of these findings on a global scale, in which analysts, consumers and regulatory bodies may display different attitudes to polluting firms. Studies might further integrate chemical analysis to investigate analyst reactions to individual toxins. Additionally, the spatial location of polluters could be considered, as some firms may operate in areas where pollution is of more concern. Analysts may exhibit differing levels of bias dependent on their own beliefs; further study could focus on correlations between analyst specific characteristics and forecast biases. Finally, research could examine the effects of other dimensions of corporate social responsibility on analyst forecasts; these variables could include firm policies on climate change, social and governance policies, or firm involvement with international politics and geopolitical conflict.

Conclusion

This thesis consists of three chapters that examine separate themes of environmental finance. The first chapter focuses on climate change risk in U.S. security markets, while the second and third chapters examine relationships between firm pollution, institutional ownership and security analyst forecast bias.

The first chapter in this thesis tests for evidence of a priced low frequency temperature risk factor. Using U.S. data, I estimate low frequency temperature shocks and hypothesise the existence of a priced temperature risk factor under the assumptions of the classical consumption-based asset pricing framework. Using equity market data, I conduct a pooled panel regression, Fama-MacBeth regressions, and portfolio tests, however none of the empirical asset pricing methodologies provide any evidence of a low frequency temperature risk premium. Furthermore, I find that estimated temperature betas do not correlate with the abnormal returns generated around the Paris Agreement of 2015, nor do they correlate with measures of aggregate climate exposures derived from self-reported climate risks.

Chapter two examines institutional ownership of the equity of U.S. polluters. Theories on social norms suggest that institutional investors may be reluctant to own securities associated with discriminated and controversial firms; I hypothesise that polluters belong to this set of companies. Results provide evidence in favour of this hypothesis, showing that institutional investors hold proportionately less polluter equities, ceteris paribus. I find evidence of a negative trend in institutional ownership of polluter stocks, and that institutional investors with aggressive trading strategies and short-term investment horizons disproportionately own polluter stocks. I also find evidence of reduced security analyst coverage of polluter stocks. Tests find no evidence of polluter stock abnormal returns, in accordance with the shunned-stock hypothesis.

The final chapter examines security analyst biases when forecasting the earnings of polluters. Under a behavioural finance framework, I hypothesise that security analysts are systematically pessimistic when forecasting the next annual earnings of polluters. I find evidence supporting this hypothesis. Tests show that while forecast pessimism is associated with total firm pollution, there is no significant relationship with pollution scaled by firm sales. Furthermore, forecast pessimism is most strongly associated with releases of the standard chemicals in the TRI database, compared to the more toxic bio-accumulative and dioxin chemical types. On an individual analyst level, I find evidence supporting the conservatism bias. Analysts that are ex-ante identified as pessimistic (optimistic) towards polluters again exhibit pessimism (optimism) in future forecasts for polluters. Results indicate that pessimistic, optimistic and neutral analysts all become increasingly pessimistic towards polluting firms nearer to the earnings date. Lastly, despite evidence of an aggregate pessimism in forecasts towards polluters, I find no evidence of positive earnings surprises generating abnormal returns once polluter earnings are announced.

Appendices

Appendix A: Chapter 1

Table A.1: A list of alternative models used to generate *Temp* in robustness tests. I find that these measures of *Temp* also produce estimates of the low frequency temperature risk premium which are statistically and economically insignificant when used to recreate the main results. As an example, in the final column I present pooled panel regression estimates of temperature risk premiums using these alternative measures of *Temp*, based on the Carhart 4-factor model. P-values are based on two-way clustered standard errors and are shown in brackets below estimates. P-values in bold denote significance at the 10% level.

Model	**Description**	**Alternative estimated Temperature risk premium**
$\Delta MA_t = MA_t - MA_{t-1}$	*Temp* is set as the first order difference in 60-month moving temperature averages, as in Bansal et al. (2016).	0.000 (0.948)
$MA_t = \alpha + \beta * MA_{t-1} + \varepsilon$	*Temp* is set as the residual term from an AR(1) model of 60-month moving average temperatures.	0.000 (0.954)
$\Delta MA_t = \alpha + \beta * MA_{t-1} + \beta * \Delta MA_{t-1} + \varepsilon$	*Temp* is set as the residual term from a model which allows for feedback to contemporaneous changes in moving average temperatures from both lagged levels and changes in moving average temperatures.	-0.001 (0.870)
$MA_t = \alpha + \beta * t + \varepsilon$	*Temp* is set as the residual term from a model which allows for a deterministic linear trend in moving average temperatures.	-0.036 (0.103)
$MA_t = \alpha + \beta * t + \beta * e^t + \varepsilon$	*Temp* is set as the residual term from a model which allows for a deterministic linear and exponential trend in moving average temperatures.	-0.018 (0.406)

Figure A.1: Average Fama-French 49 industry portfolio temperature betas. Average temperature betas are estimated using a 60-month rolling window regression controlling for the Carhart 4-factor model, which are averaged in the time series over the entire sample of 1988 to 2016 for each of the 49 industry portfolios.

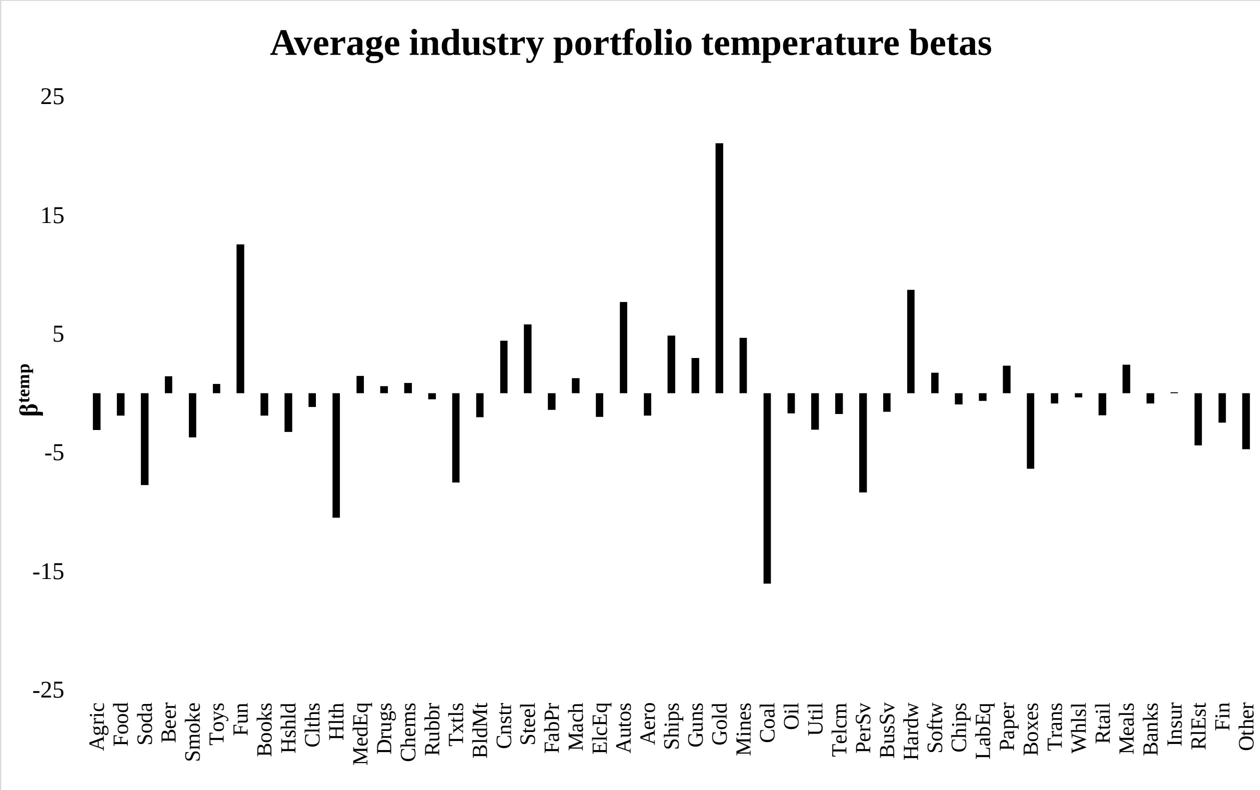

Table A.2: Results of the Paris Climate Agreement on the predicted loser portfolios of the Fama-French 49 industries. Industry portfolios are sorted into the expected loser category based on positive estimated temperature betas measured using the Carhart 4-factor model. Industry average temperature betas are shown along with daily abnormal returns that are captured by the *Paris* dummy coefficient. P-values are Newey-West adjusted for 5-day lags and are shown in brackets below estimations. P-values in bold denote significance at the 10% level.

Paris agreement event study: expected losers

Industry	β^{temp}	Dummy	Industry	β^{temp}	Dummy
Agric	15.516	-0.465	Util	0.349	0.114
		(0.000)			(0.577)
Beer	5.982	-0.189	Telcm	6.024	-0.519
		(0.499)			**(0.000)**
Smoke	8.847	0.191	BusSv	3.241	-0.129
		(0.164)			**(0.001)**
Toys	2.387	1.600	Hardw	0.417	0.102
		(0.000)			(0.529)
Hshld	1.297	0.303	Softw	2.157	0.278
		(0.126)			**(0.051)**
Hlth	2.013	-1.163	Paper	0.550	0.090
		(0.017)			(0.282)
Drugs	3.692	0.148	Trans	0.718	-0.241
		(0.128)			(0.123)
Chems	1.758	-1.163	Meals	1.105	-0.399
		(0.000)			**(0.020)**
FabPr	5.509	0.744	Banks	6.702	0.278
		(0.197)			**(0.000)**
Ships	6.501	0.333	Fin	4.196	-0.855
		(0.025)			**(0.000)**
Gold	7.483	-1.239	Other	0.608	0.265
		(0.174)			(0.104)
Mines	1.640	0.539			
		(0.028)			

Table A.3: Results of the Paris Climate Agreement on the predicted winner portfolios of the Fama-French 49 industries. Industry portfolios are sorted into the expected winner category based on negative estimated temperature betas measured using the Carhart 4-factor model. Industry average temperature betas are shown along with daily abnormal returns that are captured by the *Paris* dummy coefficient. P-values are Newey-West adjusted for 5-day lags and are shown in brackets below estimations. P-values in bold denote significance at the 10% level.

Paris agreement event study: expected winners

Industry	β^{temp}	Dummy	Industry	β^{temp}	Dummy
Food	-6.496	-0.041	Autos	-4.579	-0.410
		(0.689)			**(0.020)**
Soda	-11.806	0.090	Aero	-6.541	-0.086
		(0.306)			(0.814)
Fun	-5.054	-0.151	Guns	-8.303	-0.167
		(0.354)			(0.705)
Books	-11.143	-0.695	Coal	-48.094	-0.941
		(0.000)			(0.319)
Clths	-5.490	0.257	Oil	-6.810	0.618
		(0.151)			(0.363)
MedEq	-1.583	0.232	PerSv	-14.503	-0.307
		(0.196)			**(0.000)**
Rubbr	-1.812	0.353	Chips	-12.692	-0.317
		(0.013)			(0.334)
Txtls	-15.687	0.148	LabEq	-1.315	0.101
		(0.422)			**(0.052)**
BldMt	-1.507	0.412	Boxes	-8.644	-0.198
		(0.033)			(0.786)
Cnstr	-5.178	0.067	Whlsl	-5.282	0.225
		(0.828)			(0.254)
Steel	-11.403	0.404	Rtail	-1.800	0.267
		(0.299)			**(0.059)**
Mach	-5.870	0.512	Insur	-5.124	-0.172
		(0.004)			(0.586)
ElcEQ	-3.953	0.663	RlEst	-2.000	-0.585
		(0.000)			**(0.006)**

Table A.4: Pooled panel robustness tests using shocks to low frequency cross-sectional temperature volatility in the U.S. as an alternative proxy for temperature risk. The robustness variable, *Tempvol*, is created in an almost identical method as *Temp*. First, I generate 60-month moving averages of temperatures for each of the 48 contiguous states, and then estimate the monthly cross-sectional standard deviation. I take the first order difference in the standard deviations in local temperatures and regress them on lagged first order differences. The residuals of this regression are stored as *Tempvol*. I then recreate the pooled panel regressions, as set out in the main tests, but substitute *Temp* with *Tempvol*. P-values are based on two-way clustered standard errors and are shown in brackets below estimates. P-values in bold denote significance at the 10% level.

Temperature volatility robustness pooled panel regression

	CAPM	FF 3	Carhart	FF 5	HXZ
Constant	0.582	0.568	0.752	0.880	0.743
	(0.037)	**(0.022)**	**(0.001)**	**(0.002)**	**(0.008)**
Tempvol	0.002	0.002	0.002	0.003	0.002
	(0.516)	(0.245)	(0.368)	(0.247)	(0.272)
MKT	0.146	0.124	-0.063	-0.194	-0.033
	(0.648)	(0.702)	(0.841)	(0.581)	(0.920)
SMB		0.128	0.235	0.200	
		(0.430)	(0.120)	(0.168)	
HML		0.101	0.006	0.053	
		(0.618)	(0.976)	(0.790)	
MOM			0.114		
			(0.732)		
RMW				0.109	
				(0.446)	
CMA				-0.062	
				(0.630)	
ME					0.092
					(0.586)
I/A					0.150
					(0.266)
ROE					0.001
					(0.996)
N	17,003	17,003	17,003	17,003	17,003

Table A.5: Second stage Fama-MacBeth robustness tests using shocks to low frequency cross-sectional temperature volatility in the U.S. as an alternative proxy for temperature risk. The robustness variable, *Tempvol*, is created in an almost identical method as *Temp*. First, I generate 60-month moving averages of temperatures for each of the 48 contiguous states, and then estimate the monthly cross-sectional standard deviation. I take the first order difference in the standard deviations in local temperatures and regress them on lagged first order differences. The residuals of this regression are stored as *Tempvol*. I then recreate the Fama-MacBeth regressions, as set out in the main tests, but substitute *Temp* with *Tempvol*. P-values are based on two-way clustered standard errors and are shown in brackets below estimates. P-values in bold denote significance at the 10% level.

Temperature volatility robustness Fama-MacBeth regression

	CAPM	FF 3	Carhart	FF 5	HXZ
Constant	0.389	0.305	0.254	0.448	0.354
	(0.173)	(0.189)	(0.110)	**(0.046)**	(0.142)
Tempvol	-0.001	-0.001	-0.001	-0.002	-0.001
	(0.741)	(0.687)	(0.449)	(0.257)	(0.480)
MKT	0.295	0.337	0.410	0.223	0.299
	(0.388)	(0.266)	(0.197)	(0.443)	(0.331)
SMB		0.016	0.053	-0.016	
		(0.909)	(0.709)	(0.914)	
HML		0.273	0.229	0.219	
		(0.155)	(0.229)	(0.233)	
MOM			0.494		
			(**0.070**)		
RMW				0.071	
				(0.567)	
CMA				0.052	
				(0.731)	
ME					0.079
					(0.596)
I/A					0.106
					(0.494)
ROE					0.184
					(0.250)
N	347	347	347	347	347

Appendix B: Chapter 2

Table B.1: Robustness test results, where regression *(2.1)* is repeated with additional corporate governance control variables based on KLD ratings. Governance control variables include limited managerial compensation (cgov_str_a), excessive managerial compensation (cgov_con_b), investment in other meritable companies (cgov_str_c), strong reporting quality (cgov_str_d), weak reporting quality (cgov_con_h), and the total number of governance strengths and concerns (cgov_str_num and cgov_con_num respectively). More information on these variables can be found through the WRDS MSCI ESG KLD STATS variable description page. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 1,115 firm-year observations in the sample for each specification. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Robustness test of institutional ownership panel regressions with corporate governance control variables

Variable	(1)	(2)	(3)
Polluterdummy	-0.0523**	-0.0319	-0.0513**
	(-2.24)	(-0.77)	(-2.25)
t			0.0013
			(0.21)
INDBETA	0.0869***	0.0163	0.0901***
	(4.37)	(1.05)	(4.35)
LOGSIZE	-0.0096	-0.0055	-0.0095
	(-0.83)	(-0.45)	(-0.81)
LOGBM	0.0662	0.1260	0.0606
	(1.28)	(1.94)	(1.12)
STD	-0.0109	-0.0125	-0.0142**
	(-1.18)	(-1.25)	(-2.10)
PRINV	-0.2859***	-0.3431***	-0.2742***
	(-4.10)	(-3.87)	(-3.44)
RET	0.0013	0.0021	0.0012
	(0.60)	(0.74)	(1.55)
NASD	-0.0497**	-0.0589**	-0.0489**
	(-2.33)	(-2.30)	(-2.31)
SP500	0.0014	0.0115	-0.0017
	(0.06)	(0.40)	(-0.07)
cgov_str_a	-0.0663	-0.0762	-0.0664
	(-1.38)	(-1.27)	(-1.37)
cgov_con_b	0.1011***	0.0831***	0.1021***
	(4.21)	(2.86)	(4.34)
cgov_str_c	-0.0814***	-0.1266***	-0.0823***
	(-3.31)	(-4.38)	(-3.16)
cgov_str_d	-0.0624	-0.0580	-0.0644
	(-0.99)	(-0.81)	(-1.02)
cgov_con_h	0.0937***	0.0990***	0.0876***
	(3.95)	(3.87)	(3.61)
cgov_str_num	0.0059	0.0105	0.0068
	(0.14)	(0.19)	(0.16)
cgov_con_num	-0.0298	-0.0249	-0.0290
	(-1.53)	(-1.34)	(-1.47)
Fixed effects	Year	Year & Industry	None
N	1,115	1,115	1,115
Adjusted R^2	0.1488	0.2116	0.1596

Table B.2: Results of the institutional ownership and analyst coverage firm fixed effects panel regressions where the dependent variables are *IO* and *LOGCOV*, reported in columns 1 and 2 respectively. Regressions *(2.1)* and *(2.4)* are conducted using firm fixed effects. The first column of coefficients represents estimates generated from regression *(2.1)* with *IO* as the dependent variable, while the second column presents estimates generated from regression *(2.4)* with *LOGCOV* as the dependent variable. The independent dummy variable *NASD* is excluded as it has no within-firm variation in the sample. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 8,954 firm-year observations in both samples. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Institutional ownership and analyst coverage regressions with firm fixed effects

Variable	*IO*	*LOGCOV*
Polluterdummy	0.0093	0.0634
	(0.65)	(0.80)
INDBETA	0.0275*	0.0968
	(1.91)	(1.44)
LOGSIZE	0.0585***	0.3252***
	(9.95)	(11.54)
LOGBM	0.0480**	0.1275
	(2.55)	(1.53)
STD	-0.0112***	0.0178*
	(-3.08)	(1.71)
PRINV	-0.0112	0.0875**
	(-0.41)	(2.44)
RET	-0.0009	-0.0262***
	(-1.07)	(-9.76)
SP500	-0.0262	-0.1048
	(-1.57)	(-1.04)
Fixed effects	Year & Firm	Year & Firm
N	8,954	8,954
Adjusted R^2	0.7868	0.7944

Table B.3: Results of yearly fixed effects panel regressions *(2.1)* and *(2.4)* with *Sindummy*, a dummy variable activated for securities identified by Hong & Kacperczyk (2009) as sin stocks. Sin firms are identified on their SIC and NAICS codes. Sin firms have SIC codes of 2100-2199, 2080-2085, and/or NAICS codes 7132, 71312, 713210, 71329, 713290, 72112, 721120. There are 86 observations in the sample where *Sindummy* is activated. I include but do not report the full list of independent variables used in regressions *(2.1)* and *(2.4)*. The first column of coefficients represents estimates generated from regression *(2.1)* with *IO* as the dependent variable, while the second column presents estimates generated from regression *(2.4)* with *LOGCOV* as the dependent variable. I present regression coefficient estimates for *Polluterdummy* and *Sindummy* with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 8,954 firm-year observations in both samples. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Institutional ownership and analyst coverage regressions with *Sindummy*

Variable	*IO*	*LOGCOV*
Polluterdummy	-0.0445***	-0.1510**
	(-3.18)	(-2.17)
Sindummy	-0.1291***	-0.7386
	(-5.33)	(-1.29)
Fixed effects	Year	Year
N	8,954	8,954
Adjusted R^2	0.4094	0.4215

Table B.4: Results of the institutional ownership and analyst coverage panel regressions where the dependent variables are *IO* and *LOGCOV*, reported in columns 1 and 2 respectively. Regressions are conducted based on equations *(2.1)* and *(2.4)* with *Log Total Releases* as a continuous pollution variable, measured as the natural log of *Total Releases* in absolute pounds plus 1. The first column of coefficients represents estimates generated from regression *(2.1)* with *IO* as the dependent variable, while the second column presents estimates generated from regression *(2.4)* with *LOGCOV* as the dependent variable. I present regression coefficient estimates with t-statistics in brackets below. Standard errors are adjusted with two-way clustering on industry and year. There are 8,954 firm-year observations in both samples. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Institutional ownership and analyst coverage regressions with *Log Total Releases*

Variable	*IO*	*LOGCOV*
Log Total Releases	-0.0014	-0.0102
	(-0.92)	(-1.34)
INDBETA	0.1105***	0.2946***
	(3.61)	(6.46)
LOGSIZE	0.0414***	0.3286***
	(6.22)	(10.31)
LOGBM	0.0138	-0.0764
	(0.66)	(-0.68)
STD	-0.0132*	0.0219
	(-1.69)	(1.64)
PRINV	-0.1104**	-0.0449
	(-2.18)	(-0.54)
RET	-0.0014	-0.0298***
	(-1.27)	(-8.27)
NASD	-0.0547***	0.0575
	(-3.86)	(0.67)
SP500	-0.0262	0.2566**
	(-1.30)	(2.36)
Fixed effects	Year	Year
N	8,954	8,954
Adjusted R^2	0.4028	0.4164

Figure B.1: Robustness impulse response function for the PVAR model *(2.10)* using a log transformed pollution variable. The impulse variable is Δ*IO* while the response variable is Δ*Log Total Releases*. *Log Total Releases* is measured as the natural log of *Total Releases* in absolute pounds plus 1. The dark bands around the impulse response estimate represent 95% confidence intervals generated with bootstrapped standard errors from 1000 random draws.

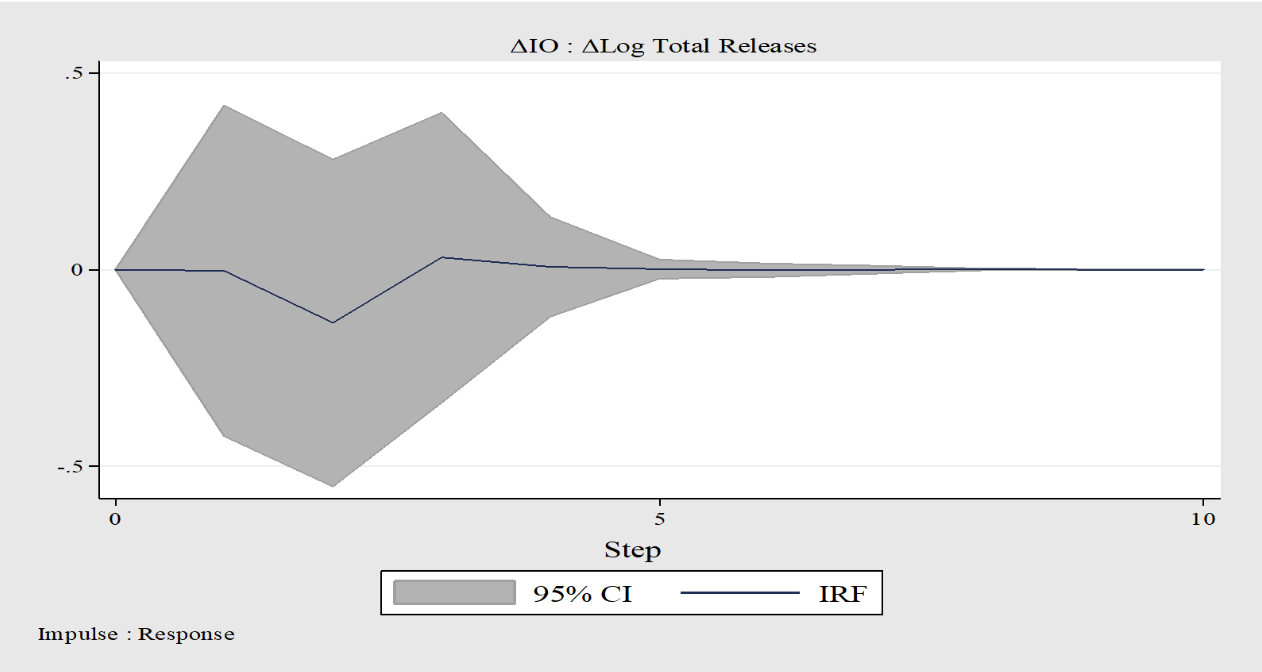

Appendix C: Chapter 3

Table C.1: Results of the robustness analyst forecast bias regression *(3.2)* using one-year and two-year lagged *Toxic Releases*. The dependent variable is *FERROR*, while the independent variable of interest is *Toxic Releases*. *Toxic Releases* is measured in billions of pounds. I present regression coefficient estimates with p-values in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Firm toxicity and analyst forecast bias robustness regressions with lagged *Toxic Releases*

	One-year lagged *Toxic Releases*	Two-year lagged *Toxic Releases*
Toxic Releases	0.575^{**}	0.428
	(0.019)	(0.106)
LOGSIZE	0.025^{***}	0.024^{***}
	(0.000)	(0.000)
LOGBM	0.122^{***}	0.107^{**}
	(0.002)	(0.011)
LEV	-0.064	-0.058
	(0.124)	(0.177)
FPERIOD	-0.063^{***}	-0.060^{***}
	(0.000)	(0.000)
COV	0.001	0.001
	(0.247)	(0.402)
SPREAD	0.000	0.000
	(0.530)	(0.542)
EXP	0.0004^{*}	0.0004^{*}
	(0.000)	(0.000)
FTE	-0.920^{**}	-0.856^{**}
	(0.017)	(0.019)
LOSS	-0.472^{***}	-0.458^{***}
	(0.000)	(0.000)
ECHANGE	-0.240^{***}	-0.233^{***}
	(0.000)	(0.000)
Fixed effects	Industry & Month	Industry & Month
N	398,171	375,104
Adjusted R^2	0.2061	0.2018

Table C.2: Results of robustness tests of primary results using alternative deflators in the denominator for *FERROR* construction. I rerun regression *(3.2)* using alternative methods of constructing *FERROR*, as shown in the first column with specification details given in the second column. I estimate the coefficients of *Toxic Releases* from regression 3 using industry and monthly fixed effects, the full set of control variables, and cluster standard errors using two-way clustering by firm and month. There are 425,621 observations in each regression. Estimates are presented in the final column, with p-values shown in brackets below. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

***FERROR* specification**	**Description**	**Alternative coefficient for *Toxic Releases***
$\frac{Actual_{i,j,h} - FEPS_{i,j,t,h}}{\lvert FEPS_{i,j,t,h} \rvert}$	*FERROR* is deflated by the absolute value of the earnings per share as forecasted by the analyst. I winsorize this specification of *FERROR* at the 2.5% and 97.5% levels and adjust observations with a 0 value denominator as in main tests.	0.435*** (0.000)
$\frac{Actual_{i,j,h} - FEPS_{i,j,t,h}}{\lvert BVAPS_{j,h-1} \rvert}$	*FERROR* is deflated by the most recent book value of assets per share of the target firm. I winsorize this specification of *FERROR* at the 2.5% and 97.5% levels.	0.025*** (0.002)
$\frac{Actual_{i,j,h} - FEPS_{i,j,t,h}}{\lvert REVPS_{j,h-1} \rvert}$	*FERROR* is deflated by the most recent revenue per share of the target firm. I winsorize this specification of *FERROR* at the 2.5% and 97.5% levels.	0.037*** (0.000)
$\frac{Actual_{i,j,h} - FEPS_{i,j,t,h}}{\lvert P_{j,t} \rvert}$	*FERROR* is deflated by the price per share of the target firm as at the beginning of the month of the forecast. I winsorize this specification of *FERROR* at the 2.5% and 97.5% levels.	0.028** (0.020)

Table C.3: Robustness test results, where regression *(3.2)* is repeated with additional corporate governance control variables based on KLD ratings. The dependent variable is *FERROR*, while the independent variable of interest is *Toxic Releases*. Governance control variables include limited managerial compensation (cgov_str_a), excessive managerial compensation (cgov_con_b), investment in other meritable companies (cgov_str_c), strong reporting quality (cgov_str_d), weak reporting quality (cgov_con_h), and the total number of governance strengths and concerns (cgov_str_num and cgov_con_num respectively). More information on these variables can be found through the WRDS MSCI ESG KLD STATS variable description page. There are 66,649 observations in the sample. *Toxic Releases* is measured in billions of pounds. I include but do not report the full list of independent variables used in regression *(3.2)*. I present regression coefficient estimates with p-values in brackets below. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Robustness test of analyst bias panel regressions with corporate governance control variables

Variable	
Toxic Releases	2.537***
	(0.003)
cgov_str_a	-0.006
	(0.908)
cgov_con_b	-0.031
	(0.257)
cgov_str_c	0.130***
	(0.000)
cgov_str_d	0.062*
	(0.077)
cgov_con_h	0.064*
	(0.073)
cgov_str_num	-0.005
	(0.857)
cgov_con_num	0.034
	(0.179)
Fixed effects	Industry & Month
N	66,649
Adjusted R^2	0.1952

Table C.4: Results of the firm-earnings date fixed effects test for analyst bias by forecast horizon for the subsample of polluters. I run two panel regressions with dependent variable *FERROR*, in which the analyst identifiers *Green* and *Grey* are interacted with the forecast horizon variable. In the first regression, the forecast horizon variable is *FPERIOD*, measured as the gap between the forecast and earnings date in hundreds of days. The second regression replaces *FPERIOD* with 3 dummy variables, activated if the earnings forecast is made within 91 – 180, 181 – 270, and 271 – 360 day period prior to the earnings date, labelled *FGROUP2*, *FGROUP3*, *FGROUP4* respectively. Only forecasts for firms in the top yearly quintile of polluters are included in the subsample of observations. I use firm-earnings date and monthly fixed effects. I exclude *LOGSIZE*, *LOGMB*, *LOSS* and *ECHANGE* from the list of control variables. I present interaction and standalone variable coefficients estimates with p-values in brackets below. Coefficient estimates for control variables are omitted for brevity. There are 100,812 observations in the panel. Standard errors are adjusted with two-way clustering on firm and month. Significance at the 10% level is denoted with *, at the 5% level with ** and at the 1% level with ***.

Analyst bias and forecast horizon interaction results with firm earnings-date fixed effects

Variable	(1)	(2)
Green	0.047***	0.041***
	(0.003)	(0.004)
Grey	0.071***	0.055***
	(0.002)	(0.005)
FPERIOD	-0.057**	
	(0.010)	
*GREEN * FPERIOD*	-0.020**	
	(0.013)	
*GREY * FPERIOD*	-0.042***	
	(0.002)	
FGROUP2		-0.004
		(0.888)
FGROUP3		0.013
		(0.666)
FGROUP4		0.001
		(0.973)
*GREEN * FGROUP2*		-0.021
		(0.152)
*GREEN * FGROUP3*		-0.051**
		(0.011)
*GREEN * FGROUP4*		-0.047**
		(0.021)
*GREY * FGROUP2*		-0.050**
		(0.020)
*GREY * FGROUP3*		-0.087***
		(0.003)
*GREY * FGROUP4*		-0.112***
		(0.002)
Fixed effects	Firm earnings-date & Month	Firm earnings-date & Month
N	100,812	100,812
Adjusted R^2	0.6410	0.6409

Figure C.1: Average daily abnormal returns of firms in the polluter quartiles over a window of t – 10 to t + 10 around firm annual earnings announcements. Abnormal returns are measured as the excess realised returns from the CAPM model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 9,378 observations in the sample. Daily abnormal returns are presented in percentage format.

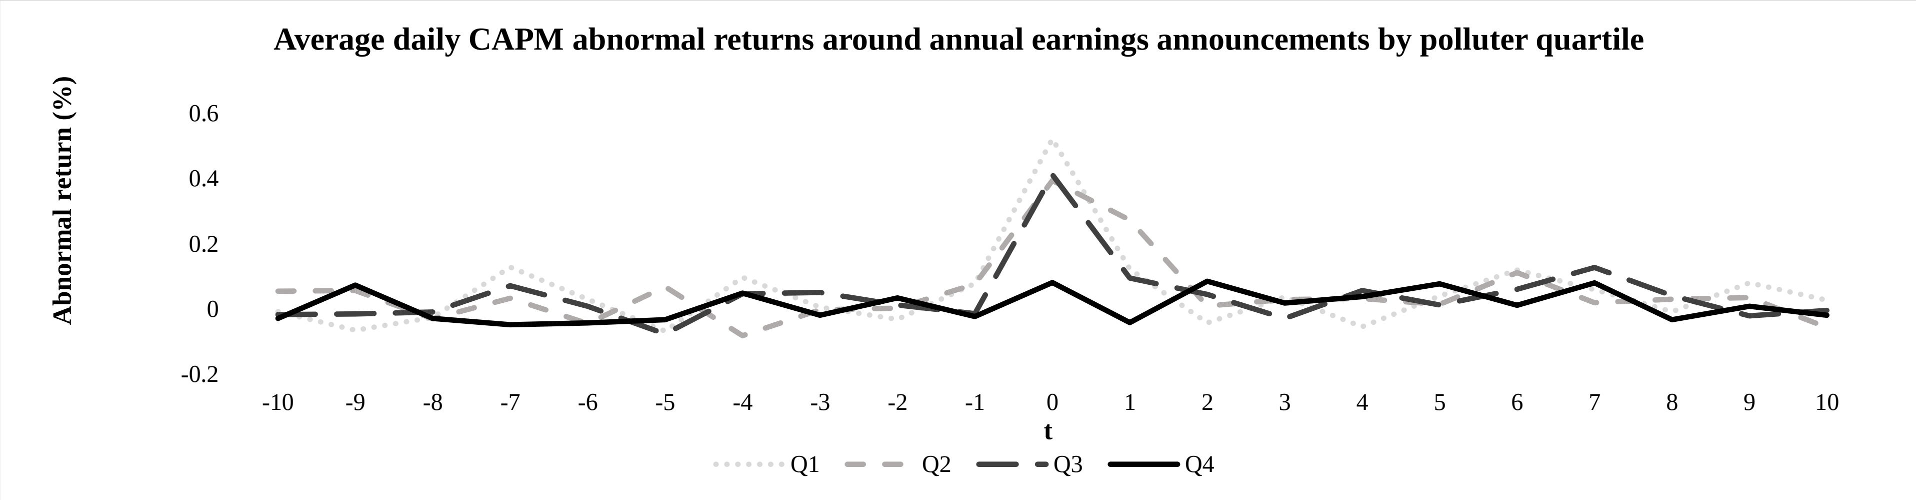

Figure C.2: Average cumulative abnormal returns of firms in the polluter quartiles over a window of *t – 10* to *t + 10* around firm annual earnings announcements. Abnormal returns are measured as the excess realised returns from the CAPM model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 9,378 observations in the sample. Cumulative returns are compounded daily abnormal returns and are presented in percentage format.

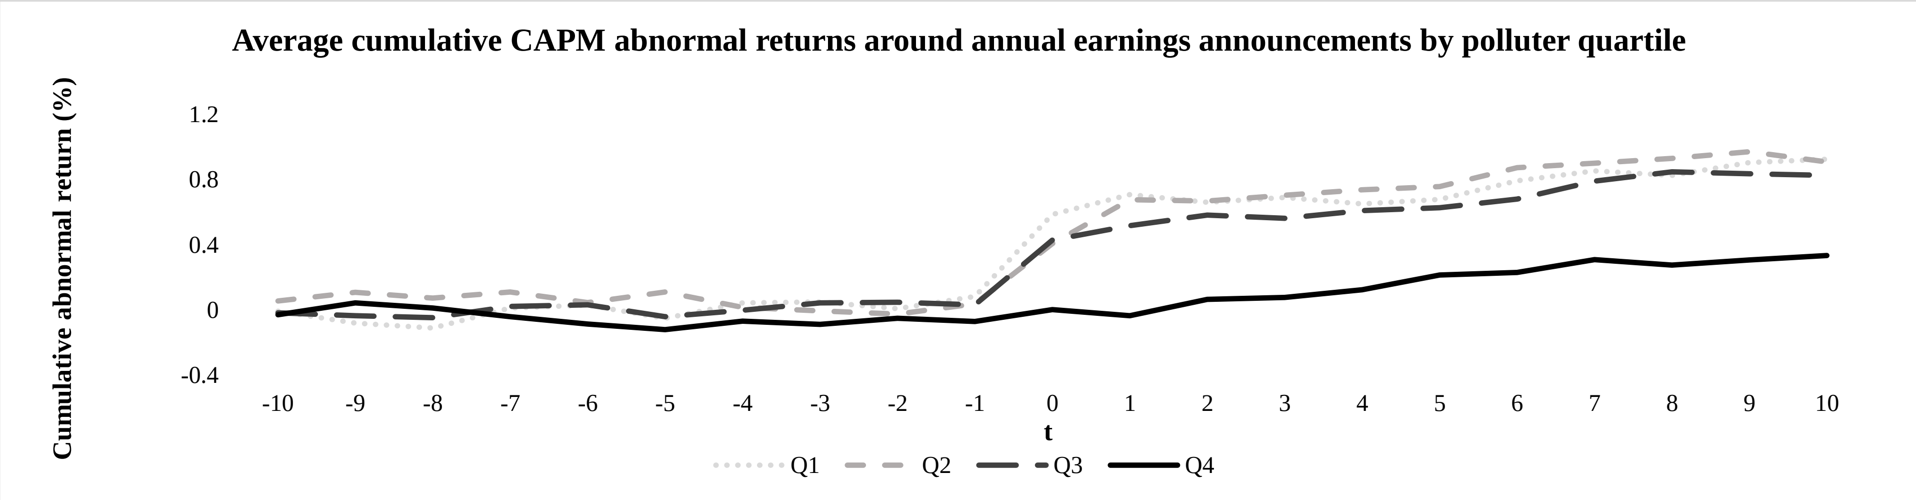

Figure C.3: Average daily abnormal returns of firms in the polluter quartiles over a window of t – 10 to t + 10 around firm quarterly earnings announcements. Abnormal returns are measured as the excess realised returns from the CAPM model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 39,400 observations in the sample. Daily abnormal returns are presented in percentage format.

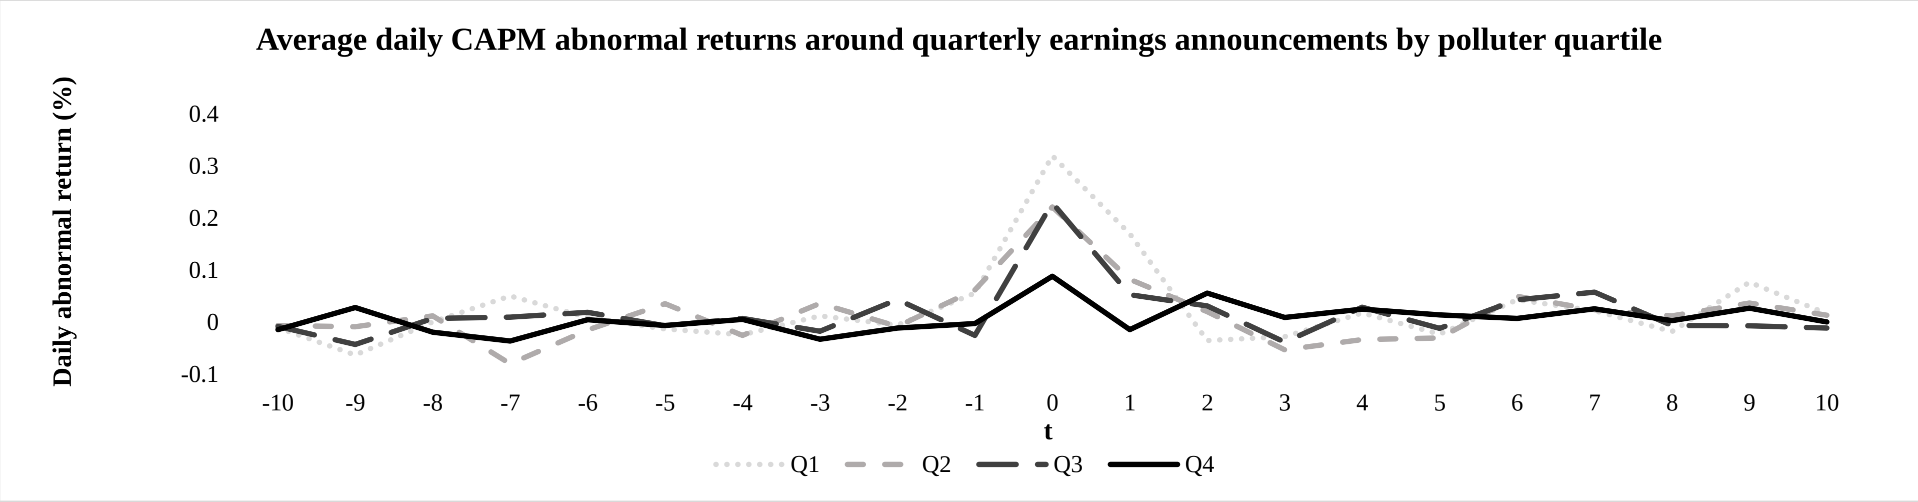

Figure C.4: Average cumulative abnormal returns of firms in the polluter quartiles over a window of *t – 10* to *t + 10* around firm quarterly earnings announcements. Abnormal returns are measured as the excess realised returns from the CAPM model. Quartile 1 represents firms that have the lowest yearly ranking of *Toxic Releases*, while quartile 4 includes firms with the highest yearly ranking of *Toxic Releases*. There are 39,400 observations in the sample. Cumulative returns are compounded daily abnormal returns and are presented in percentage format.

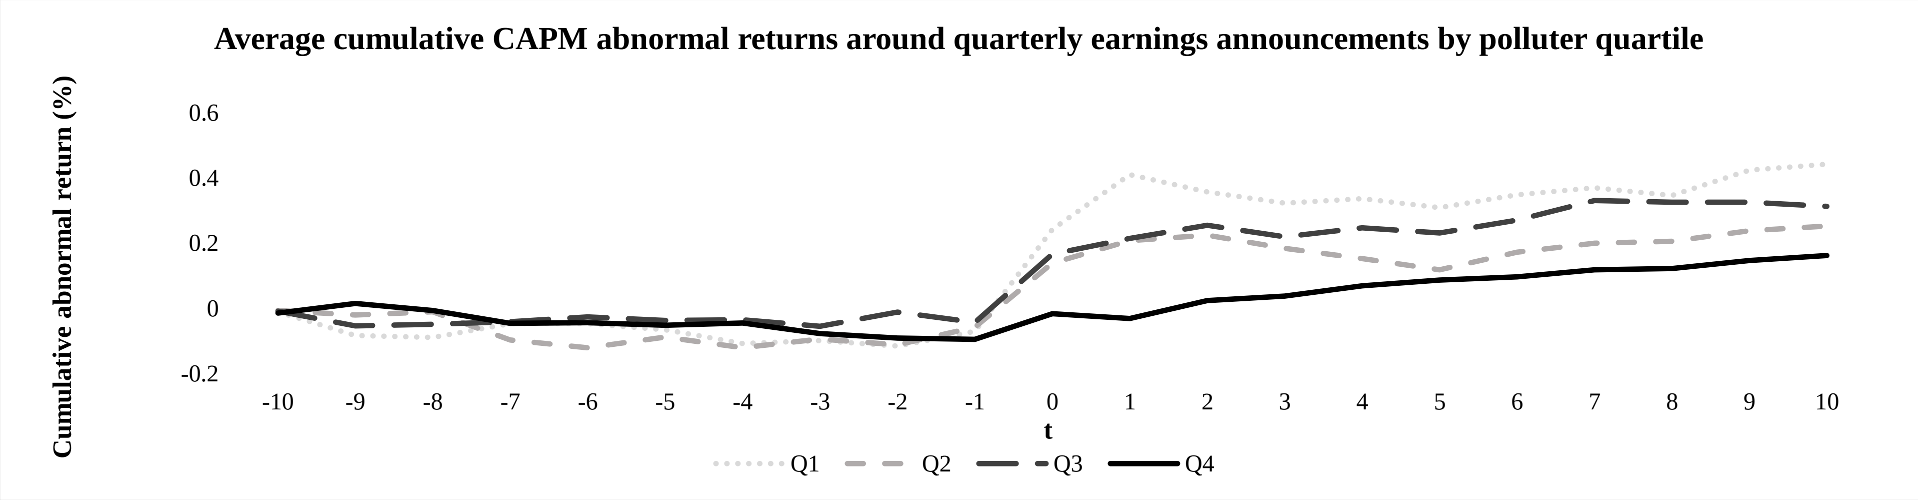

References

Abarbanell, J. S. (1991). Do analysts' earnings forecasts incorporate information in prior stock price changes? *Journal of Accounting and Economics, 14*(2), 147-165.

Abrigo, M. R., & Love, I. (2015). *Estimation of panel vector autoregression in stata: A package of programs.* Unpublished manuscript.

Aggarwal, R., Erel, I., Ferreira, M., & Matos, P. (2011). Does governance travel around the world? Evidence from institutional investors. *Journal of Financial Economics, 100*(1), 154-181.

Akerlof, G. A. (1980). A theory of social custom, of which unemployment may be one consequence. *The Quarterly Journal of Economics, 94*(4), 749-775.

Alford, A. W., & Berger, P. G. (1999). A simultaneous equations analysis of forecast accuracy, analyst following, and trading volume. *Journal of Accounting, Auditing & Finance, 14*(3), 219-240.

Ali, A., Klein, A., & Rosenfeld, J. (1992). Analysts' use of information about permanent and transitory earnings components in forecasting annual EPS. *Accounting Review, 67*(1), 183-198.

Amir, E., & Ganzach, Y. (1998). Overreaction and underreaction in analysts' forecasts. *Journal of Economic Behavior & Organization, 37*(3), 333-347.

Angel, J. J., & Rivoli, P. (1997). Does ethical investing impose a cost upon the firm? A theoretical perspective. *The Journal of Investing, 6*(4), 57-61.

Balvers, R., Du, D., & Zhao, X. (2017). Temperature shocks and the cost of equity capital: Implications for climate change perceptions. *Journal of Banking & Finance, 77*, 18-34.

Bansal, R., Kiku, D., & Ochoa, M. (2016). *Price of long-run temperature shifts in capital markets.* Unpublished manuscript.

Bansal, R., & Ochoa, M. (2011). *Temperature, aggregate risk, and expected returns.* Unpublished manuscript.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics, 121*(3), 823-866.

Barro, R. J., & Jin, T. (2011). On the size distribution of macroeconomic disasters. *Econometrica, 79*(5), 1567-1589.

Barron, O. E., Byard, D., & Liang, L. (2013). Analyst pessimism and forecast timing. *Journal of Business Finance & Accounting, 40*(5-6), 719-739.

Becker, G. S. (2010). *The economics of discrimination*. University of Chicago press.

Beckers, S., Steliaros, M., & Thomson, A. (2004). Bias in European analysts' earnings forecasts. *Financial Analysts Journal, 60*(2), 74-85.

Berkman, H. (2013). The capital asset pricing model: A revolutionary idea in finance! *Abacus, 49*(S1), 32-35.

Berkman, H., Dimitrov, V., Jain, P. C., Koch, P. D., & Tice, S. (2009). Sell on the news: Differences of opinion, short-sales constraints, and returns around earnings announcements. *Journal of Financial Economics, 92*(3), 376-399.

Berkman, H., Jacobsen, B., & Lee, J. B. (2011). Time-varying rare disaster risk and stock returns. *Journal of Financial Economics, 101*(2), 313-332.

Berkman, H., Jona, J., & Soderstrom, N. (2019). *Firm-Specific Climate Risk and Market Valuation.* Unpublished manuscript.

Berkman, H., & Truong, C. (2009). Event day 0? After-hours earnings announcements. *Journal of Accounting Research, 47*(1), 71-103.

Binder, J. (1998). The event study methodology since 1969. *Review of Quantitative Finance and Accounting, 11*(2), 111-137.

Blitz, D., & Fabozzi, F. J. (2017). Sin stocks revisited: Resolving the sin stock anomaly. *The Journal of Portfolio Management, 44*(1), 105-111.

Bradshaw, M. T. (2011). *Analysts' forecasts: What do we know after decades of work?* Unpublished manuscript.

Brav, A., Lehavy, R., & Michaely, R. (2005). Using expectations to test asset pricing models. *Financial Management, 34*(3), 31-64.

Brous, P. A., & Kini, O. (1993). A reexamination of analysts' earnings forecasts for takeover targets. *Journal of Financial Economics, 33*(2), 201-225.

Brous, P. A. (1992). Common stock offerings and earnings expectations: A test of the release of unfavorable information. *The Journal of Finance, 47*(4), 1517-1536.

Brown, L. (1998). *Managerial behavior and the bias in analysts' earnings forecasts.* Unpublished manuscript.

Brown, L. D. (1997). Analyst forecasting errors: Additional evidence. *Financial Analysts Journal, 53*(6), 81-88.

Brown, L. D., & Rozeff, M. S. (1978). The superiority of analyst forecasts as measures of expectations: Evidence from earnings. *The Journal of Finance, 33*(1), 1-16.

Burgstahler, D. C., & Eames, M. J. (2003). Earnings management to avoid losses and earnings decreases: Are analysts fooled? *Contemporary Accounting Research, 20*(2), 253-294.

Burns, Z., Chiu, A., & Wu, G. (2010). Overweighting of small probabilities. *Wiley Encyclopedia of Operations Research and Management Science*.

Bushee, B. J. (1998). The influence of institutional investors on myopic R&D investment behavior. *Accounting Review, 73*(3), 305-333.

Bushee, B. J. (2001). Do institutional investors prefer near-term earnings over long-run value? *Contemporary Accounting Research, 18*(2), 207-246.

Bushee, B. J., & Noe, C. F. (2000). Corporate disclosure practices, institutional investors, and stock return volatility. *Journal of Accounting Research, 38,* 171-202.

Butler, K. C., & Lang, L. H. (1991). The forecast accuracy of individual analysts: Evidence of systematic optimism and pessimism. *Journal of Accounting Research, 29*(1), 150-156.

Campbell, J. Y. (2003). Consumption-based asset pricing. *Handbook of the Economics of Finance, 1,* 803-887.

Capstaff, J., & Paudyal Rees, K., William. (1998). Analysts' forecasts of German firms' earnings: A comparative analysis. *Journal of International Financial Management & Accounting, 9*(2), 83-116.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance, 52*(1), 57-82.

Carney, M. (2015). Breaking the Tragedy of the Horizon–climate change and financial stability. *Speech given at Lloyd's of London by the Governor of the Bank of England*, 29.

Chhaochharia, V., Kumar, A., & Niessen-Ruenzi, A. (2012). Local investors and corporate governance. *Journal of Accounting and Economics, 54*(1), 42-67.

Choi, D., Gao, Z., & Jiang, W. (2020). Attention to global warming. *The Review of Financial Studies, 33*(3), 1112-1145.

Ciccone, S. J. (2005). Trends in analyst earnings forecast properties. *International Review of Financial Analysis, 14*(1), 1-22.

Coën, A., Desfleurs, A., & L'Her, J. (2009). International evidence on the relative importance of the determinants of earnings forecast accuracy. *Journal of Economics and Business, 61*(6), 453-471.

Copeland, L., & Zhu, Y. (2007). *Rare disasters and the equity premium in a two-country world.* Unpublished manuscript.

Crost, B., Duquennois, C., Felter, J. H., & Rees, D. I. (2018). Climate change, agricultural production and civil conflict: Evidence from the Philippines. *Journal of Environmental Economics and Management, 88*, 379-395.

Das, S., Levine, C. B., & Sivaramakrishnan, K. (1998). Earnings predictability and bias in analysts' earnings forecasts. *Accounting Review*, *73*(2), 277-294.

Dasgupta, P. (2008). Discounting climate change. *Journal of Risk and Uncertainty, 37*(2-3), 141-169.

De Bondt, W. F., & Forbes, W. P. (1999). Herding in analyst earnings forecasts: Evidence from the United Kingdom. *European Financial Management*, *5*(2), 143-163.

De Bondt, W. F., & Thaler, R. (1985). Does the stock market overreact? *The Journal of Finance*, *40*(3), 793-805.

De Bondt, W. F., & Thaler, R. H. (1987). Further evidence on investor overreaction and stock market seasonality. *The Journal of Finance, 42*(3), 557-581.

De Bondt, W. F., & Thaler, R. H. (1990). Do security analysts overreact? *The American Economic Review, 80*(2), 52-57.

Deal done. (2015). *The Economist*. Retrieved from: https://www.economist.com.

Dell, M., Jones, B. F., & Olken, B. A. (2009). Temperature and income: Reconciling new cross-sectional and panel estimates. *American Economic Review, 99*(2), 198-204.

Dell, M., Jones, B. F., & Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics, 4*(3), 66-95.

Deloitte. (2013). Incorporating environmental issues' impact into growth strategies. *The Wall Street Journal.* Retrieved from: https://www.wsj.com.

Derwall, J., & Koedijk, K. (2009). Socially responsible fixed-income funds. *Journal of Business Finance & Accounting, 36*(1-2), 210-229.

Derwall, J., Koedijk, K., & Ter Horst, J. (2011). A tale of values-driven and profit-seeking social investors. *Journal of Banking & Finance, 35*(8), 2137-2147.

Ding, D. K., Charoenwong, C., & Seetoh, R. (2004). Prospect theory, analyst forecasts, and stock returns. *Journal of Multinational Financial Management, 14*(4-5), 425-442.

Do Paço, A. M. F., Raposo, M. L. B., & Leal Filho, W. (2009). Identifying the green consumer: A segmentation study. *Journal of Targeting, Measurement and Analysis for Marketing*, *17*(1), 17-25.

Donadelli, M., Jüppner, M., Paradiso, A., & Schlag, C. (2019). *Temperature volatility risk.* Unpublished manuscript.

Donadelli, M., Jüppner, M., Riedel, M., & Schlag, C. (2017). Temperature shocks and welfare costs. *Journal of Economic Dynamics and Control, 82*, 331-355.

Dooley, R. S., & Lerner, L. D. (1994). Pollution, profits, and stakeholders: The constraining effect of economic performance on CEO concern with stakeholder expectations. *Journal of Business Ethics, 13*(9), 701-711.

Doyle, J. T., Lundholm, R. J., & Soliman, M. T. (2006). The extreme future stock returns following I/B/E/S earnings surprises. *Journal of Accounting Research, 44*(5), 849-887.

Duru, A., & Reeb, D. M. (2002). International diversification and analysts' forecast accuracy and bias. *The Accounting Review, 77*(2), 415-433.

Eames, M., & Kim, Y. (2012). Analyst vs. market forecasts of earnings management to avoid small losses. *Journal of Business Finance & Accounting, 39*(5-6), 649-674.

Easterwood, J. C., & Nutt, S. R. (1999). Inefficiency in analysts' earnings forecasts: Systematic misreaction or systematic optimism? *The Journal of Finance, 54*(5), 1777-1797.

El Ghoul, S., Guedhami, O., Kim, H., & Park, K. (2018). Corporate environmental responsibility and the cost of capital: International evidence. *Journal of Business Ethics, 149*(2), 335-361.

Elton, E. J. (1999). Presidential address: Expected return, realized return, and asset pricing tests. *The Journal of Finance, 54*(4), 1199-1220.

Embong, Z., & Hosseini, L. (2018). Analyst forecast accuracy and earnings management. *Asian Journal of Accounting and Governance, 10,* 97-108.

Engle, R. F., Giglio, S., Lee, H., Kelly, B. T., & Stroebel, J. (2020). Hedging climate change news. *The Review of Financial Studies, 33*(3), 1184-1216.

EPA. (1999). *Persistent Bioaccumulative Toxic (PBT) Chemicals; Lowering of Reporting Thresholds for Certain PBT Chemicals; Addition of Certain PBT Chemicals; Community Right-to-Know Toxic Chemical Reporting.* Office of the Federal Register, National Archives and Records Administration.

Ernst & Young. (2017). *Is your nonfinancial performance revealing the true value of your business to investors?* Retrieved from: https://www.ey.com/Publications/vwLUAssets/EY-_Nonfinancial_performance_may_influence_investors/$File/ey-nonfinancial-performance-may-influence-investors.pdf.

Espahbodi, H., Espahbodi, P., & Espahbodi, R. (2015). Did analyst forecast accuracy and dispersion improve after 2002 following the increase in regulation? *Financial Analysts Journal, 71*(5), 20-37.

Fabozzi, F. J., Ma, K., & Oliphant, B. J. (2008). Sin stock returns. *The Journal of Portfolio Management, 35*(1), 82-94.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics, 33*(1), 3-56.

Fama, E. F., & French, K. R. (1997). Industry costs of equity. *Journal of Financial Economics, 43*(2), 153-193.

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics, 116*(1), 1-22.

Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy, 81*(3), 607-636.

Fernando, C. S., Sharfman, M. P., & Uysal, V. B. (2017). Corporate environmental policy and shareholder value: Following the smart money. *Journal of Financial and Quantitative Analysis, 52*(5), 2023-2051.

Fisher-Vanden, K., & Thorburn, K. S. (2011). Voluntary corporate environmental initiatives and shareholder wealth. *Journal of Environmental Economics and Management, 62*(3), 430-445.

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

Fraj, E., & Martinez, E. (2007). Ecological consumer behaviour: an empirical analysis. *International journal of consumer studies*, *31*(1), 26-33.

Francis, J., & Philbrick, D. (1993). Analysts' decisions as products of a multi-task environment. *Journal of Accounting Research, 31*(2), 216-230.

Franzoni, F. A. (2002). *Where is beta going? The riskiness of value and small stocks.* Unpublished manuscript.

Fried, D., & Givoly, D. (1982). Financial analysts' forecasts of earnings: A better surrogate for market expectations. *Journal of Accounting and Economics, 4*(2), 85-107.

Fujiwara, I., Ichiue, H., Nakazono, Y., & Shigemi, Y. (2013). Financial markets forecasts revisited: Are they rational, stubborn or jumpy? *Economics Letters, 118*(3), 526-530.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics, 127*(2), 645-700.

Galema, R., Plantinga, A., & Scholtens, B. (2008). The stocks at stake: Return and risk in socially responsible investment. *Journal of Banking & Finance, 32*(12), 2646-2654.

Gaspar, J., Massa, M., & Matos, P. (2005). Shareholder investment horizons and the market for corporate control. *Journal of Financial Economics, 76*(1), 135-165.

Geczy, C., Stambaugh, R. F., & Levin, D. (2005). *Investing in socially responsible mutual funds*. Unpublished manuscript.

Geertsema, P. (2014). *Fire sales and contagion in financial assets*. Unpublished manuscript.

Gompers, P. A., & Metrick, A. (2001). Institutional investors and equity prices. *The Quarterly Journal of Economics, 116*(1), 229-259.

Gormley, T. A., & Matsa, D. A. (2013). Common errors: How to (and not to) control for unobserved heterogeneity. *The Review of Financial Studies, 27*(2), 617-661.

Griffin, P., Lont, D., & Lubberink, M. (2019). Extreme high surface temperature events and equity-related physical climate risk. *Weather and Climate Extremes, 26*, 100220.

Groysberg, B., Healy, P. M., & Maber, D. A. (2011). What drives sell-side analyst compensation at high-status investment banks? *Journal of Accounting Research, 49*(4), 969-1000.

Guenster, N., Bauer, R., Derwall, J., & Koedijk, K. (2011). The economic value of corporate eco-efficiency. *European Financial Management, 17*(4), 679-704.

Hamilton, J. T. (1995). Pollution as news: Media and stock market reactions to the toxics release inventory data. *Journal of Environmental Economics and Management, 28*(1), 98-113.

Hamilton, S., Jo, H., & Statman, M. (1993). Doing well while doing good? The investment performance of socially responsible mutual funds. *Financial Analysts Journal, 49*(6), 62-66.

Hart, S. L., & Ahuja, G. (1996). Does it pay to be green? An empirical examination of the relationship between emission reduction and firm performance. *Business Strategy and the Environment, 5*(1), 30-37.

Hayes, R. M. (1998). The impact of trading commission incentives on analysts' stock coverage decisions and earnings forecasts. *Journal of Accounting Research, 36*(2), 299-320.

Hitz, S., & Smith, J. (2004). Estimating global impacts from climate change. *Global Environmental Change, 14*(3), 201-218.

Hong, H., & Kacperczyk, M. (2009). The price of sin: The effects of social norms on markets. *Journal of Financial Economics*, *93*(1), 15-36.

Hong, H., & Kostovetsky, L. (2012). Red and blue investing: Values and finance. *Journal of Financial Economics, 103*(1), 1-19.

Hong, H., & Kubik, J. D. (2003). Analyzing the analysts: Career concerns and biased earnings forecasts. *The Journal of Finance, 58*(1), 313-351.

Hood, M., Nofsinger, J. R., & Varma, A. (2014). Conservation, discrimination, and salvation: Investors' social concerns in the stock market. *Journal of Financial Services Research, 45*(1), 5-37.

Horváthová, E. (2010). Does environmental performance affect financial performance? A meta-analysis. *Ecological Economics, 70*(1), 52-59.

Hou, K., Xue, C., & Zhang, L. (2014). Digesting anomalies: An investment approach. *Review of Financial Studies, 8*(3), 650-705.

Hutira, S. (2016). *Determinants of analyst forecasting accuracy.* Unpublished manuscript.

Intergovernmental Panel on Climate Change. (2014). *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Geneva, Switzerland: Author.

Irvine, P. J. (2004). Analysts' forecasts and brokerage-firm trading. *The Accounting Review, 79*(1), 125-149.

Ivković, Z., & Weisbenner, S. (2005). Local does as local is: Information content of the geography of individual investors' common stock investments. *The Journal of Finance, 60*(1), 267-306.

Jaffe, A. B., & Kerr, S. (2015). The science, economics, and politics of global climate change: A review of the climate casino by William Nordhaus. *Journal of Economic Literature, 53*(1), 79-91.

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance, 48*(1), 65-91.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263-291.

Kang, S., O'Brien, J., & Sivaramakrishnan, K. (1994). Analysts' interim earnings forecasts: Evidence on the forecasting process. *Journal of Accounting Research, 32*(1), 103-112.

Kim, I., Wan, H., Wang, B., & Yang, T. (2019). Institutional investors and corporate environmental, social, and governance policies: Evidence from toxics release data. *Management Science, 65*(10), 4901-4926.

Kim, I., & Venkatachalam, M. (2011). Are sin stocks paying the price for accounting sins? *Journal of Accounting, Auditing & Finance, 26*(2), 415-442.

King, A. A., & Lenox, M. J. (2001). Does it really pay to be green? An empirical study of firm environmental and financial performance. *Journal of Industrial Ecology, 5*(1), 105-116.

Leventis, S., Hasan, I., & Dedoulis, E. (2013). The cost of sin: The effect of social norms on audit pricing. *International Review of Financial Analysis, 29*, 152-165.

Levis, M., & Liodakis, M. (2001). Contrarian strategies and investor expectations: The UK evidence. *Financial Analysts Journal, 57*(5), 43-56.

Liesen, A. (2015). Climate change and financial market efficiency. *Business & Society, 54*(4), 511-539.

Lim, T. (2001). Rationality and analysts' forecast bias. *The Journal of Finance, 56*(1), 369-385.

Lin, H., & McNichols, M. F. (1998). Underwriting relationships, analysts' earnings forecasts and investment recommendations. *Journal of Accounting and Economics, 25*(1), 101-127.

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

Liston, D. P. (2016). Sin stock returns and investor sentiment. *The Quarterly Review of Economics and Finance, 59*, 63-70.

Liu, Y., Lu, H., & Veenstra, K. (2014). Is sin always a sin? The interaction effect of social norms and financial incentives on market participants' behavior. *Accounting, Organizations and Society, 39*(4), 289-307.

Livnat, J., & Mendenhall, R. R. (2006). Comparing the post–earnings announcement drift for surprises calculated from analyst and time series forecasts. *Journal of Accounting Research, 44*(1), 177-205.

Lobe, S., & Walkshäusl, C. (2016). Vice versus virtue investing around the world. *Review of Managerial Science, 10*(2), 303-344.

Marsden, A., Veeraraghavan, M., & Ye, M. (2008). Heuristics of representativeness, anchoring and adjustment, and leniency: Impact on earnings' forecasts by Australian analysts. *Quarterly Journal of Finance and Accounting, 47*(2), 83-102.

Maxwell, J. W., Lyon, T. P., & Hackett, S. C. (2000). Self-regulation and social welfare: The political economy of corporate environmentalism. *The Journal of Law and Economics, 43*(2), 583-618.

McNichols, M., & O'Brien, P. C. (1997). Self-selection and analyst coverage. *Journal of Accounting Research, 35*, 167-199.

Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics, 15*(2), 145-161.

Mendelsohn, R., Nordhaus, W. D., & Shaw, D. (1994). The impact of global warming on agriculture: A Ricardian analysis. *The American Economic Review, 84*(4), 753-771.

Mendenhall, R. R. (1991). Evidence on the possible underweighting of earnings-related information. *Journal of Accounting Research, 29*(1), 170-179.

Merkley, K., Michaely, R., & Pacelli, J. (2017). Does the scope of the sell-side analyst industry matter? An examination of bias, accuracy, and information content of analyst reports. *The Journal of Finance, 72*(3), 1285-1334.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The Journal of Finance, 42*(3), 483-510.

Mikhail, M. B., Walther, B. R., & Willis, R. H. (1997). Do security analysts improve their performance with experience? *Journal of Accounting Research, 35*, 131-157.

Mikhail, M. B., Walther, B. R., & Willis, R. H. (2003). The effect of experience on security analyst underreaction. *Journal of Accounting and Economics, 35*(1), 101-116.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of Finance, 32*(4), 1151-1168.

Miller, M. H. (1999). The history of finance. *The Journal of Portfolio Management, 25*(4), 95-101.

Nagel, S. (2005). Short sales, institutional investors and the cross-section of stock returns. *Journal of Financial Economics, 78*(2), 277-309.

Nordhaus, W. D. (2013). *The climate casino: Risk, uncertainty and economics for a warming world* (illustrated ed.) Yale University Press, 2015.

O'Brien, P. C. (1988). Analysts' forecasts as earnings expectations. *Journal of Accounting and Economics, 10*(1), 53-83.

Park, A., & Sabourian, H. (2011). Herding and contrarian behavior in financial markets. *Econometrica, 79*(4), 973-1026.

Petersen, M. A. (2009). Estimating standard errors in finance panel data sets: Comparing approaches. *The Review of Financial Studies, 22*(1), 435-480.

Petkova, R., & Zhang, L. (2005). Is value riskier than growth? *Journal of Financial Economics, 78*(1), 187-202.

Ranson, M. (2014). Crime, weather, and climate change. *Journal of Environmental Economics and Management, 67*(3), 274-302.

Renneboog, L., Ter Horst, J., & Zhang, C. (2008). The price of ethics and stakeholder governance: The performance of socially responsible mutual funds. *Journal of Corporate Finance, 14*(3), 302-322.

Richardson, S., Teoh, S. H., & Wysocki, P. (1999). *Tracking analysts' forecasts over the annual earnings horizon: Are analysts' forecasts optimistic or pessimistic?* Unpublished manuscript.

Rietz, T. A. (1988). The equity risk premium a solution. *Journal of Monetary Economics, 22*(1), 117-131.

Ritter, J. R. (2003). Behavioral finance. *Pacific-Basin Finance Journal, 11*(4), 429-437.

Romer, D. (1984). The theory of social custom: A modification and some extensions. *The Quarterly Journal of Economics, 99*(4), 717-727.

Salaber, J. M. (2007). *The determinants of sin stock returns: Evidence on the European market.* Unpublished manuscript.

Salaber, J. M. (2009). *Sin stock returns over the business cycle.* Unpublished manuscript.

Schaeffer, R., Szklo, A. S., de Lucena, André Frossard Pereira, Borba, Bruno Soares Moreira Cesar, Nogueira, L. P. P., Fleming, F. P., . . . Boulahya, M. S. (2012). Energy sector vulnerability to climate change: A review. *Energy, 38*(1), 1-12.

SEC. (2010). *Commission guidance regarding disclosure related to climate change. Final rule.* (Vol. 75, No. 25).

Shumway, T. (1997). The delisting bias in CRSP data. *The Journal of Finance, 52*(1), 327-340.

So, E. C. (2013). A new approach to predicting analyst forecast errors: Do investors overweight analyst forecasts? *Journal of Financial Economics, 108*(3), 615-640.

Starks, L. T., Venkat, P., & Zhu, Q. (2017). *Corporate ESG profiles and investor horizons.* Unpublished manuscript.

Statman, M. (2000). Socially responsible mutual funds (corrected). *Financial Analysts Journal, 56*(3), 30-39.

Statman, M., & Glushkov, D. (2009). The wages of social responsibility. *Financial Analysts Journal, 65*(4), 33-46.

The Forum for Sustainable and Responsible Investment. (2016). *Report on US sustainable, responsible and impact investing trends.* Retrieved from http://www.ussif.org/files/SIF_Trends_16_Executive_Summary(1).pdf.

Trigeorgis, L. (1993). Real options and interactions with financial flexibility. *Financial Management, 22*(3), 202-224.

Tversky, A., & Kahneman, D. (1981). *Evidential impact of base rates.* (No. TR-4). Stanford University Dept of psychology.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 184,* 1124-1131.

Von Wallis, M., & Klein, C. (2015). Ethical requirement and financial interest: A literature review on socially responsible investing. *Business Research, 8*(1), 61-98.

Wachter, J. A. (2013). Can Time-Varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance, 68*(3), 987-1035.

Wellington, F., & Sauer, A. (2005). *Framing climate risk in portfolio management*. World Resources Institute, Washington, DC.

Wright, W. F. (1980). Cognitive information processing biases: Implications for producers and users of financial information. *Decision Sciences, 11*(2), 284-298.

Zhang, P., Zhang, J., & Chen, M. (2017). Economic impacts of climate change on agriculture: The importance of additional climatic variables other than temperature and precipitation. *Journal of Environmental Economics and Management, 83,* 8-31.